Department of Defense

High Performance Computing Modernization Program

2007 Annual Report

Solving the Hard Problems

For more information about the
DoD High Performance Computing Modernization Program
and the
DoD High Performance Computing Modernization Program Office
visit our website at
http://www.hpcmo.hpc.mil

HPC DoD Modernization Program

Department of Defense
High Performance Computing Modernization Program
2007 Annual Report

A report by the Department of Defense
High Performance Computing Modernization Program Office

March 2008

Message from the Director

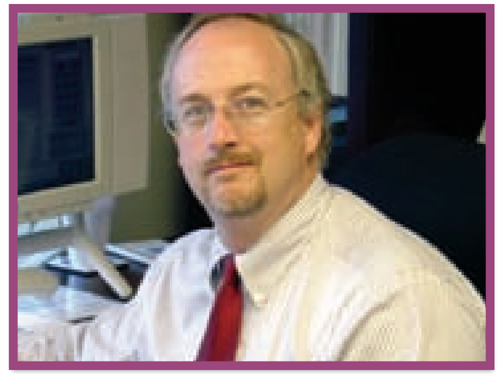

The High Performance Computing Modernization Program (HPCMP) continues to delivers to the Department of Defense (DoD) one of the world's top supercomputing infrastructures. By the end of 2007, the program had eight of the world's largest supercomputers deployed at four Major Shared Resource Centers (MSRCs) and two of the Allocated Distributed Centers (ADCs). Today, the Department of Defense has access to some of the world's most powerful supercomputers and to a variety of computing architectures, chosen to best meet the Department's identified requirements. This allows our scientists and engineers to match software applications to supercomputers. Our Defense Research and Engineering Network provides connections to over 155 sites with connection speeds ranging from 45 to 2,400 megabits per second.

Perhaps, most importantly, our program provides key computer and computational science expertise to scientists and engineers across the Department. We successfully completed several software development projects that introduced parallel, scalable production software now in use across the Department and the broader national community. This past year, we delivered 45 training events attended by 561 people and coordinated technology sharing projects between the defense laboratories and over two dozen universities.

These activities and the daily work performed by members of the HPCMP community have a positive impact on our national defense posture. Our scientists and engineers are now in possession of tools once only imagined. Today, we support nearly 500 individual projects covering the spectrum from basic research through applied engineering and test and evaluation support—each project contributing to our national defense posture.

As the program completes another year, we continue to improve the Department's supercomputing environment supported by leading edge networking capabilities. The DoD HPCMP community is working in concert with other federal agencies to identify future trends and requirements.

While the state of commercial computing hardware continues to advance at a rapid pace, driven by a huge commercial market and Defense Advanced Research Projects Agency (DARPA) investments, advances in parallel software driven primarily by federal investment have not kept pace with the hardware. During FY 2007, we began planning the execution of a new effort, starting in FY 2008, to focus on high performance software. The new effort, Computational Research and Engineering Acquisition Tools and Environments (CREATE), will focus on transitioning and improving research applications and moving these applications into mainstream engineering. Our planning considers "lessons learned" from similar scale projects carried out by the nuclear weapons, the climate modelling, and other computational communities. It builds on the DoD experience with our institutes and portfolios as well. The CREATE project is planned as a 12-year $365M effort with 70% of the funding coming from Office of the Secretary of Defense and 30% from the Military Services. The tri-Service Test and Evaluation communities will provide validation experiments and data for the projects.

Our Department's needs for science and engineering to speed its transformation goals continue to accelerate. The High Performance Computing Modernization Program team is dedicated to deploying and operating superior supercomputing environments and productivity enhancing services allowing DoD's scientists and engineers to develop the best technological solutions for our nation's defense. As President Bush has said, "science and technology have never been more essential to the defense of the nation…"

Cray J. Henry
Director
High Performance Computing Modernization Program

Acknowledgements

This is a very special annual report. First, it documents the start-up of an important new initiative called Computational Research and Engineering Acquisition Tools and Environments. An outstanding team and a great plan have been put together to execute this initiative. Second, the Institutes have shown great success, and a targeted call for proposals will result in increasing the number by three to a total of nine during 2008. Finally, this is my last year as your editor of the annual report, at least for the next several years. Effective 1 January 2008, my assignment is that of Technical Director, Maui High Performance Computing Center. I am looking forward to new challenges and a new environment, having been assigned in the Washington area for nearly a decade.

Many professional colleagues in the field and in the program office— far too many to acknowledge— have contributed to this annual report. I thank them for their assistance and support. Their efforts have been made especially difficult by a move of our offices from Arlington to Lorton. I greatly appreciate their attention to duty amidst considerable distraction. All the efforts described in this document were performed using resources provided by the Department of Defense High Performance Computing Modernization Program. Additional sponsorship has come from the Offices of Research, the Defense Laboratories, and the Test and Evaluation Centers.

Clifford E. Rhoades, Jr.
High Performance Computing Modernization Program

Table of Contents

Section 1

Program Overview

Program Overview

Introduction

Since the founding of our nation, it has been the duty of our military to ensure our survival, defend our lives and property, and promote our vital interests at home and abroad. To provide for the common defense, our military strives for excellence in organizing, training, and equipping our forces to fight and win our nation's wars.

No modern war has been won without technological superiority. No future war will be won without it. History teaches us that victory goes to those who properly prepare. The Department of Defense (DoD) faces many challenges. We must defend America by maintaining a military second to none. We must do so affordably. The High Performance Computing Modernization Program (HPCMP) provides many capabilities and tools the Department needs to address defense problems. These tools include modern high performance computing (HPC) hardware and software and the expertise to use them.

Our military strength depends on many factors. Our people are our greatest asset for they are our intellectual capital. They include active Service members, Reserves, National Guard, civil servants, political appointees, and contractors. High performance computing hardware and software help our people make our military the best in the world.

Many military problems are complicated and often require very powerful tools to be solved. Some problems are too expensive for experiment to address. Others are too difficult to be solved with paper and pencil.

Introduction—Continued

The Department uses high performance computing tools to help solve some of these hard problems. Sophisticated hardware and software tools give us advantage over potential adversaries.

Many modern weapons systems present hard problems. Early in system development, we must make trade-offs to balance performance, time, and available resources. How do we determine cost, schedule, and performance? How do we take into account technical and management risks? HPC hardware and software contribute to answering these questions.

HPC hardware and software help us answer other important questions as well. They can be used to address a wide spectrum of issues, including: protecting our bases of operations through the mitigation of toxic threats; modeling to support certification of new aircraft-store combinations before deployment to conflicts in Afghanistan and Iraq; supporting US supremacy in air and space; conducting climate, weather, and ocean modeling that provides valuable information for countermine warfare; and preparation for emergency operations and humanitarian relief operations throughout the world. These are but a few examples.

For more than a decade, the HPCMP has supported a workforce that routinely uses HPC resources to solve many of the Department's most challenging scientific and engineering problems. This, in turn, helps the United States ensure military advantage and war-fighting superiority on the 21st century battlefield.

The Program enables scientists and engineers to further the Department's objectives through research, development, test, and evaluation (RDT&E) activities that support science and technology (S&T), and test and evaluation (T&E). These efforts focus on the most complex, and highest priority defense challenges. This annual report highlights a small portion of the work performed to support the Department.

HPCMP Mission

Accelerate development and transition of advanced defense technologies into superior warfighting capabilities by exploiting and strengthening U.S. leadership in supercomputing, communications, and computational modeling.

HPCMP Vision

A pervasive culture existing among DoD's scientists and engineers where they routinely use advanced computational environments to solve the most demanding problems.

Introduction—Continued

Defense scientists and engineers, using resources provided by the HPCMP, address multi-disciplinary scientific and engineering challenges. These include problems of interest across the Services and to joint force commanders. Designing materials for specific purposes such as body armor or agile laser eye protection, modeling complex flow fields around air systems to increase performance, and improving the accuracy of ocean and weather prediction models are examples. Today's work will:

- enhance force protection against terrorist threats;
- advance dynamic signal intelligence mission planning;
- improve detection of targets based on their spectral or spatial/spectral signatures; and
- address the critical need to develop new high energy density materials for explosives and rocket fuels.

Beginning in fiscal year 1994, Congressional investments in HPC have spurred cultural changes in the fundamental way S&T and T&E are pursued. In 1993, the Department had just over 180 gigaFLOPS (10^9 FLoating-point OPerations per Second or GF) of computational power to support the S&T community. By the end of 2007, the program had over 437 teraFLOPS (10^3 gigaFLOPS) of computing capacity, a nearly 2,500 times improvement!

Figure 1 illustrates the growth in computational capabilities at our High Performance Computing Centers. This vast increase in capability was obtained by applying sound management practices and good investment strategies. Similarly, we transitioned our communications network linking the laboratories from a government-owned, government-operated asset to a commercial environment with secure, high bandwidth capability.

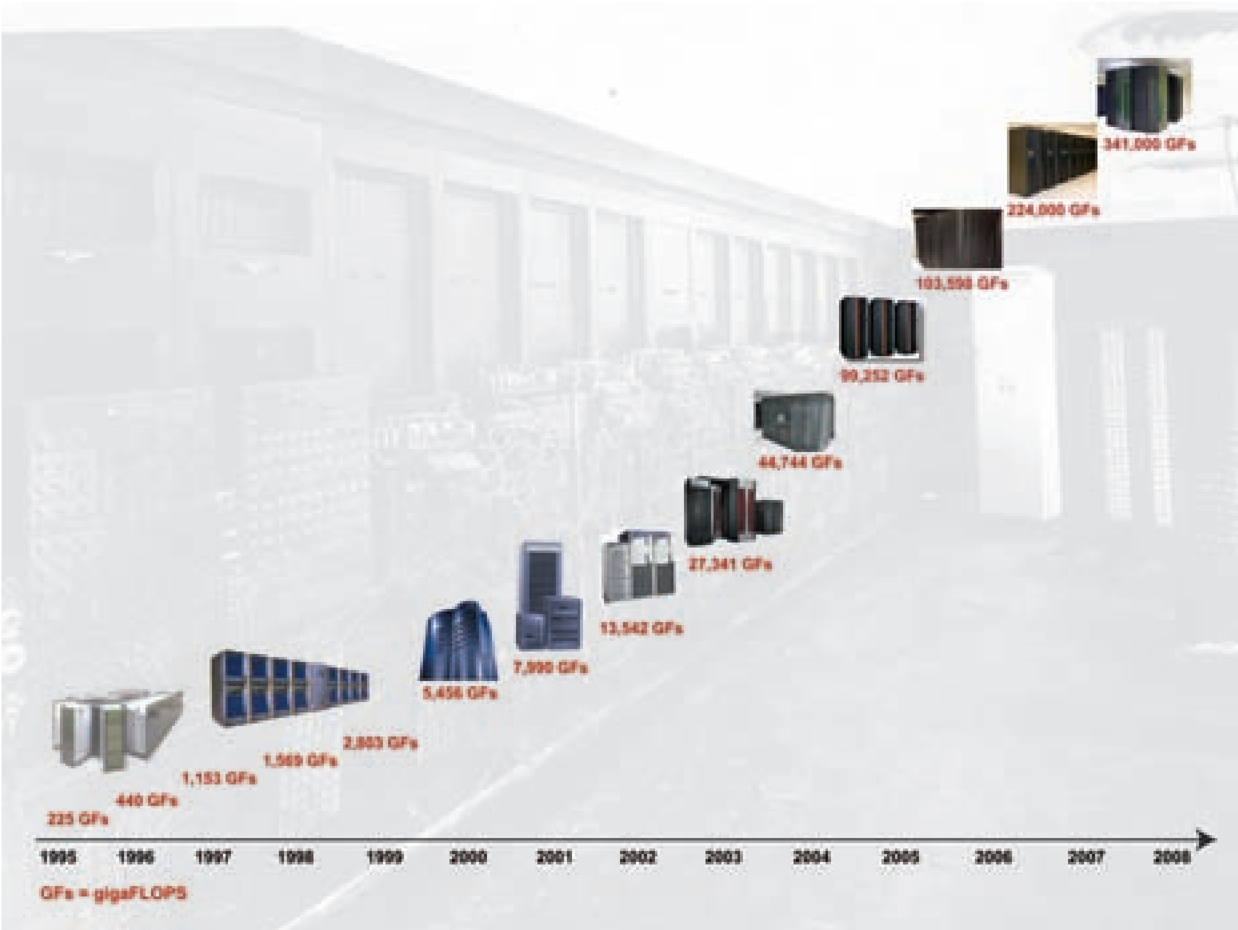

Figure 1. Computational capability at the Major Shared Resource Centers at the beginning of the fiscal year

HPCMP GOALS

The HPCMP achieves the Program's mission and vision (as stated on page 6) by focusing on five specific goals. Each activity within the program supports one or more of these goals, with progress tracked and successes delineated. These five goals are:

- acquire, deploy, operate, and maintain best-value supercomputers;
- acquire, develop, deploy, and support software applications and computational work environments that enable critical DoD research, development, and test challenges to be analyzed and solved;

Goals—continued

- acquire, deploy, operate and maintain a communications network that enables effective access to supercomputers and to distributed S&T/T&E computing environments;
- promote collaborative relationships among the DoD computational science community, the national computational science community, and minority serving institutions; and
- continuously educate the research, development, test, and evaluation workforce with the knowledge needed to employ computational modeling effectively and efficiently.

The progress the HPCMP has made in meeting these goals is discussed in detail in Section 2.

The HPCMP Community

The HPCMP community consists of nearly 4,000 scientists, engineers, computer specialists, networking specialists, and security experts working throughout the United States. All three Military Departments and several Defense Agencies participate in the program. These users execute nearly 500 projects—each validated by the Military Services and Defense Agencies. Most projects are in one or more of the program's ten computational technology areas (CTAs) (see Table 1 on page 10). Figure 2 shows the locations of people using the program's resources. The user base is diverse, drawing from the government workforce, academia, and industry. The demographics by type of workforce as well as by the DoD organizations are shown in Figure 3. Figure 4 includes a breakdown of users by CTA.

Figure 2. The beige color represents states with people using HPCMP resources

Table 1. Computational Technology Areas (CTAs)

Computational Technology Area	Acronym	Description
Computational Structural Mechanics	CSM	Covers the high resolution, multi-dimensional modeling of materials and structures subjected to a broad range of loading conditions, such as quasi-static, dynamic, electro-magnetic, shock, penetration, and blast.
Computational Fluid Dynamics	CFD	Provides accurate numerical solution of the equations describing fluid and gas motion.
Computational Chemistry, Biology, and Materials Science	CCM	Predicts properties and simulates the behavior of chemicals and materials for DoD applications. Methods ranging from quantum mechanical, atomistic, and mesoscale modeling, to multiscale theories that address challenges of length- and time-scale integration, are being developed and applied. Of recent emerging interest in the CCM CTA are methodologies that cover bioinformatics tools, computational biology, and related areas, such as cellular modeling.
Computational Electromagnetics and Acoustics	CEA	Provides high-resolution multidimensional solutions of electromagnetic and acoustic wave propagation, and their interaction with surrounding media.
Climate/Weather/Ocean Modeling and Simulation	CWO	Involves accurate numerical simulation and forecast of the Earth's atmosphere and oceans on those space and time scales important for both scientific understanding and DoD operational use.
Signal/Image Processing	SIP	Extracts and analyzes key information from various sensor outputs in real-time; sensor types include sonar, radar, visible and infrared images, signal intelligence, and navigation assets.
Forces Modeling and Simulation	FMS	Focuses on the research and development of HPC-based physical, logical, and behavioral models and simulations of battlespace phenomena in the correlation of forces.
Environmental Quality Modeling and Simulation	EQM	Involves the high-resolution modeling of hydrodynamics, geophysics, and multi-constituent fate/transport through the coupled atmospheric/land surface/subsurface environment, and their interconnections with numerous biological species and anthropogenic activities.
Electronics, Networking, and Systems/C4I	ENS	Focuses on the use of computational science in support of analysis, design, modeling, and simulation of electronics from the most basic fundamental, first principles physical level to its use for communications, sensing, and information systems engineering; activity ranges from the analysis and design of nano-devices to modeling systems-of-systems.
Integrated Modeling and Test Environments	IMT	Addresses the application of integrated modeling and simulation tools and techniques with live tests and hardware-in-the-loop simulations for the testing and evaluation of DoD weapon components, subsystems, and systems in virtual and composite virtual-real environments.

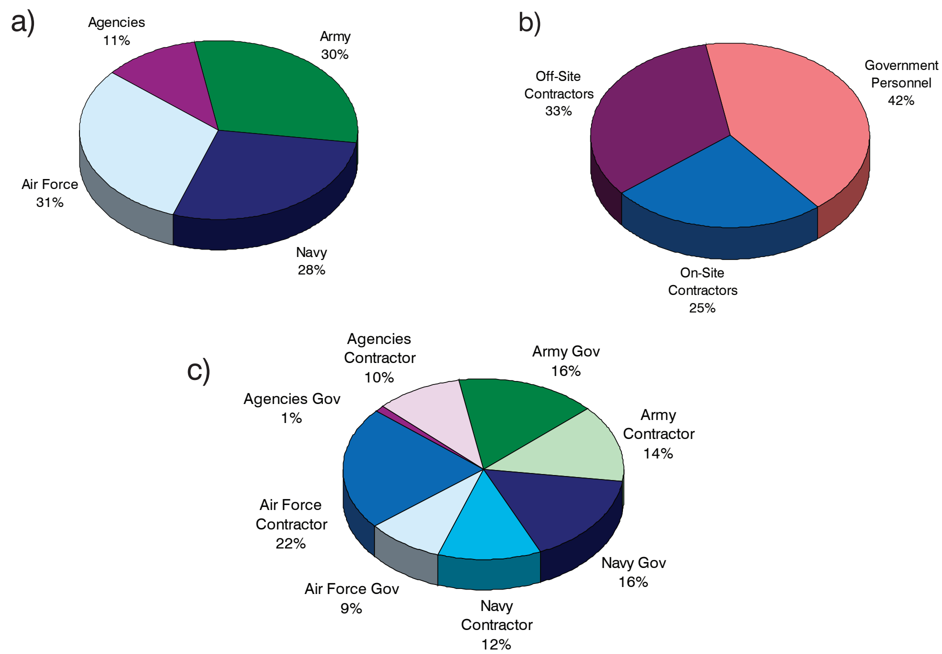

Figure 3. FY 2007 HPC statistics: a) Categorization of DoD HPC Users by Service/Agencies; b) Categorization of DoD HPC Users by Government or Contractor; and c) Distribution of DoD HPC Users as Government or Contractor by Service/Agency

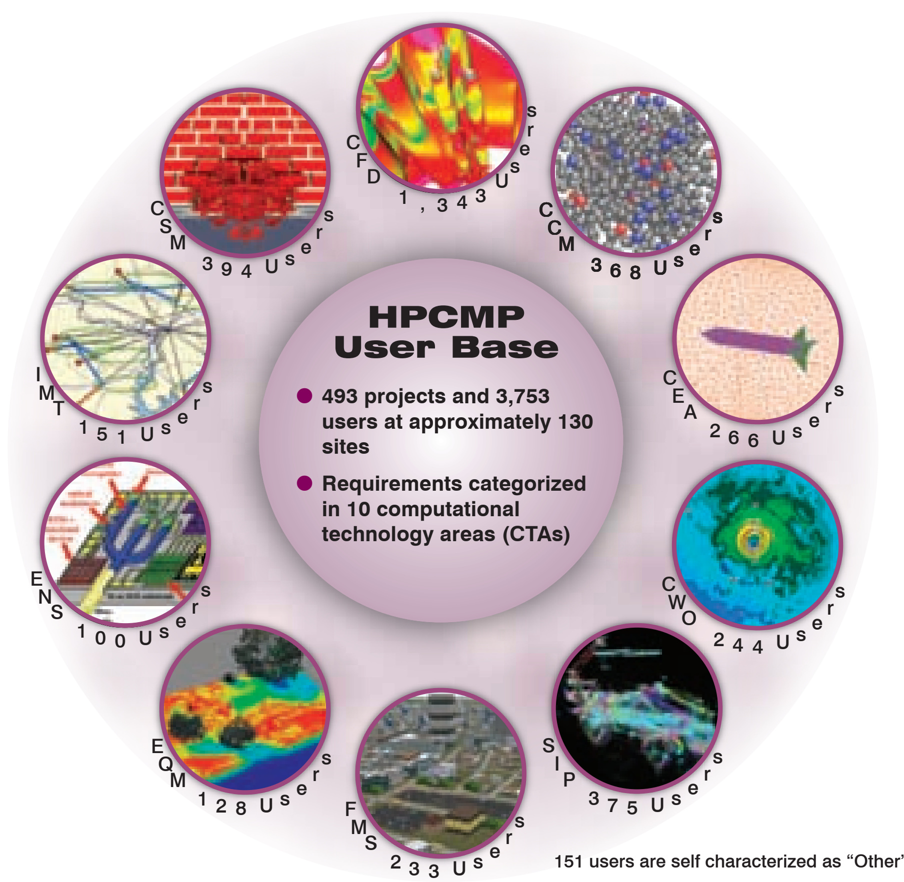

Figure 4. HPCMP serves a large, diverse, Department of Defense, user community

HPCMP Community Computational Requirements

Validated requirements serve as the basis for HPCMP investments and operational decisions. Each year, the program gathers, assesses, and validates user community requirements. This requirements determination includes all aspects of HPCMP activities and capabilities: system hardware, software, networking, and training. In the past year, overall requirements increased and are projected to continue to increase at a steady, consistent rate. Total requirements in any given year are approximately two and one-half times total program capability, ramping up from approximately 771 teraFLOPS or 967 Habus in FY 2008 to 2,370 teraFLOPS or 2,926 Habus in FY 2012 (see Figure 5). (Habus are a measure of computational performance. See callout on this page for a definition.) Once collected, the Services' and Agencies' S&T and T&E Executives review, correct, validate, and approve their requirements. HPCMP conducts requirements analyses as a fundamental part of an overall systems engineering process that collects and analyzes information to make investment decisions.

The general conclusion of that requirements analysis is that a complete HPC environment must be provided to support the DoD's S&T and T&E communities. A spectrum of computational platforms, both at the unclassified and classified levels, must be provided to support a wide range of DoD applications. These platforms must be balanced with respect to computational power, central memory, and data storage capabilities. Diverse systems and applications software that enable DoD computational scientists and engineers to perform their mission are required. A reliable high-speed network that connects the users to these resources and to each other is necessary, as is the continuation of an aggressive training program that broadens and educates DoD's HPC users. Progress must be balanced across all program activities to optimize the impact of HPC on the DoD S&T and T&E programs' support of the warfighting mission.

Habu—a Measure of Computational Performance

The HPCMP rates computer systems in terms of the speed at which DoD computational applications run on the systems. For the past six years, the HPCMP has run a suite of applications on existing and new systems to obtain performance comparisons. By comparing the timing results for these applications, the HPCMP is able to compare the performance of any system relative to the others. In 2002, a large IBM system located at the Naval Oceanographic Office MSRC named Habu, was designated the baseline system. Hence, performance measures are all in "Habu" equivalent units. For example, if a new system is rated at two Habus, that system is roughly two times more capable than a system rated at one Habu. That is, the new system executes the suite of applications at roughly twice the performance of the old. Of course, any individual application may run faster or slower. The line in Figure 5 shows the growth in computational requirements in Habu units of system performance.

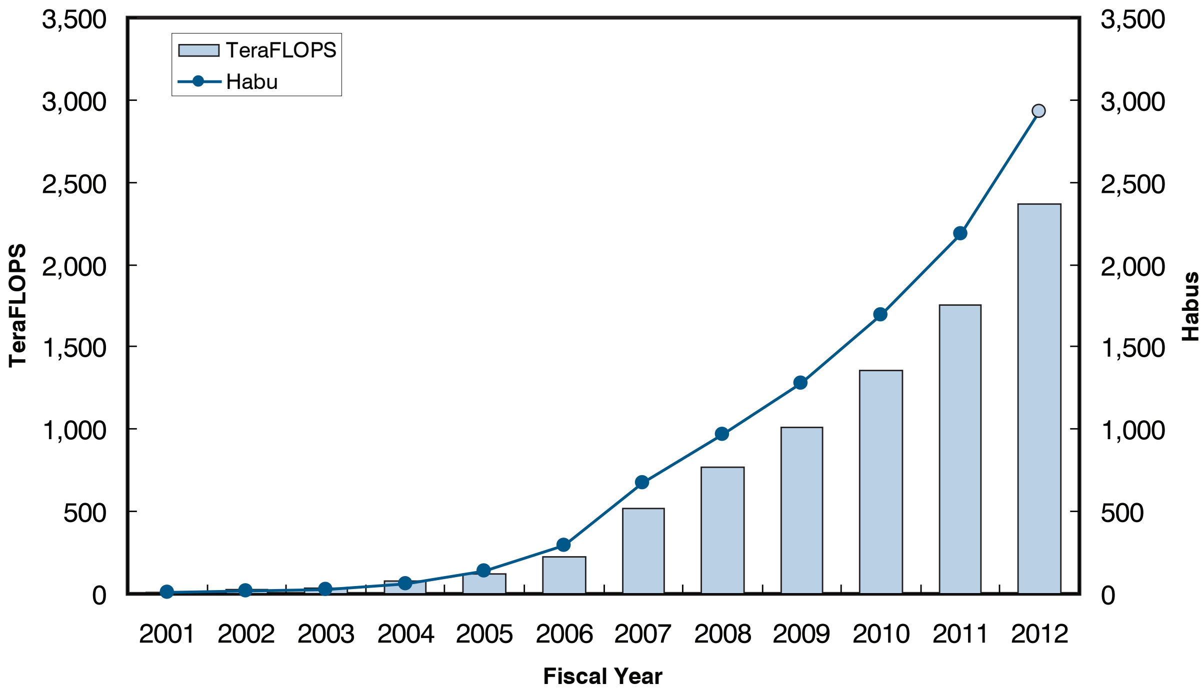

Figure 5. Total computational requirements of the HPCMP community

HPCMP Organization

Efforts of the HPCMP fall into five general areas. These are: 1) Networking, 2) Access Control and Information Assurance, 3) HPC Centers, 4) Software Applications Support, and 5) User Applications. As shown in Figure 6, they form the basis of an integrated program strategy to provide a technologically advanced computational environment to support the ongoing and emerging needs of the Department's laboratories and test centers.

Organizationally, the HPCMP has three major components: HPC Centers, Networking and Information Assurance, and Software Applications Support. These components are interdependent, with distinct business practices and community relationships.

The HPC Centers component includes four major shared resource centers (MSRCs) and four allocated distributed centers (ADCs). These computer centers provide DoD scientists and engineers with the resources necessary to solve the most demanding computational problems. While not part of the Centers component, dedicated

Figure 6. HPCMP integrated program strategy

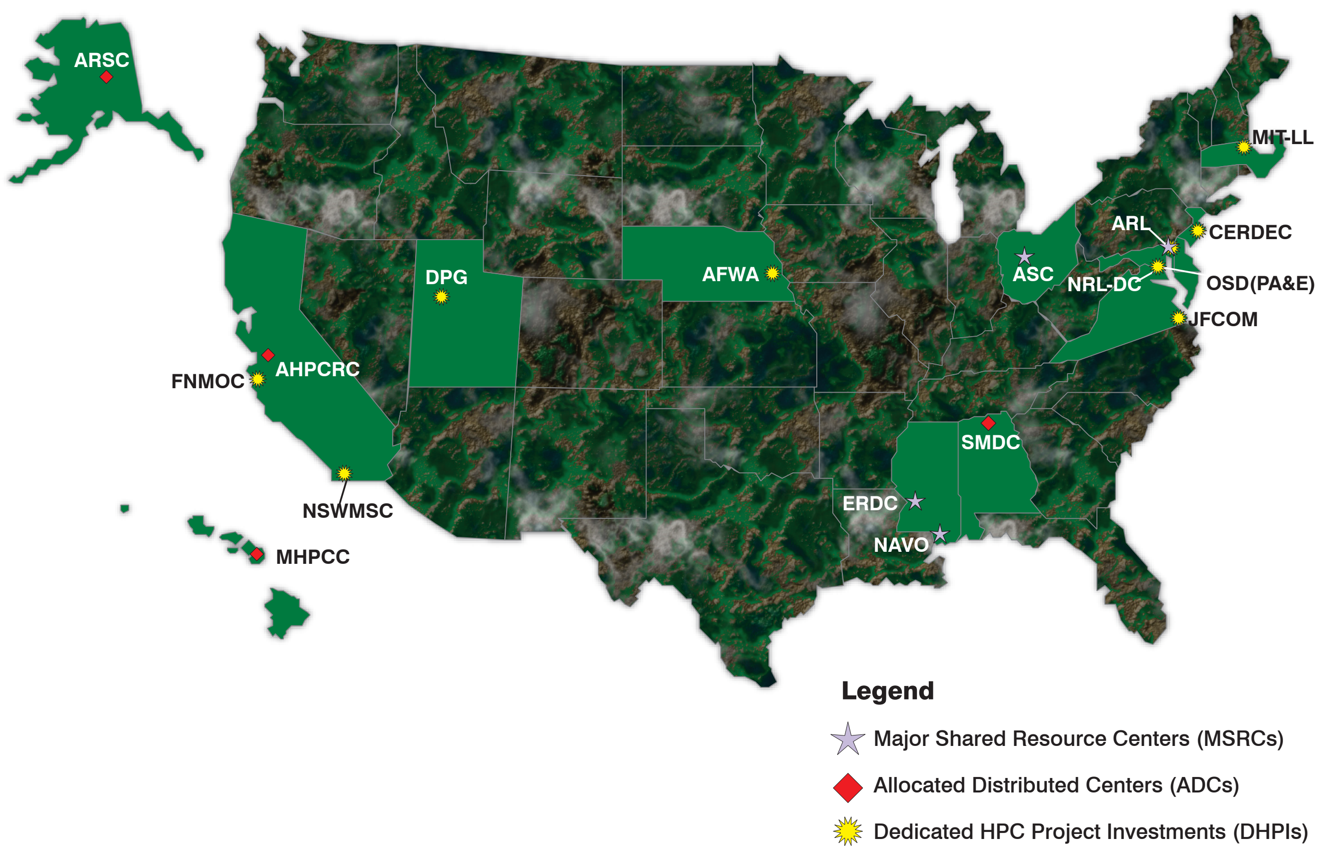

Figure 7. Location of HPCMP resources (MSRCs, ADCs, and DHPIs)

HPCMP Organization—continued

HPC project investments (DHPIs) provide resources to support specific projects not readily addressed in the HPC Centers' shared resource environment. Figure 7 shows the locations of the MSRCs, ADCs, and DHPIs.

The Networking and Information Assurance component of the HPCMP includes the Defense Research and Engineering Network (DREN). DREN provides advanced communication capabilities at faster speeds to a larger user community than previously possible, while addressing our security requirements.

The Software Applications Support (SAS) component provides expert services to assist our customers in most effectively using the HPC systems, provides investments in human capital across the DoD to facilitate the application of HPC tools, and supports a modest investment in a few high-need HPC software applications.

DoD Challenge Projects

Because resources fall short of requirements, the Program ensures resources are allocated to the highest priority requirements first. DoD Challenge Projects address high priority computational work that can be done at the shared resource centers. These projects are the DoD's highest-priority, highest-impact computational work, both from technical and mission-relevance standpoints. During FY 2007, these projects account for approximately 37% of the allocations of resources at the HPC Centers. They range from predicting the radar signature of new ground vehicles to understanding the aerodynamic flow around and through jet turbine engines. There were 40 active DoD Challenge Projects in FY 2007—26 continuing projects and 14 new projects. The 14 new projects were selected from 23 proposals submitted by the Services and Agencies in response to the HPCMP's annual call for Challenge Project proposals. Selections were made by peer review with a panel consisting of Service, Agency, DoD, and external reviewers. Table 2

Table 2. FY 2007 DoD Challenge Projects

Project Title	Project Leader/Organization
Advanced Chemical Oxygen-Iodine Laser Technology Development Using 3-D Navier-Stokes Simulation	Timothy Madden, Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, NM
Applications of Time-Accurate CFD in Order to Account for Blade-Row Interactions and Distortion Transfer in the Design of High Performance Military Fans and Compressors	Steven E. Gorrell, Air Force Research Laboratory, Propulsion Directorate, Wright-Patterson AFB, OH
Applied Computational Fluid Dynamics (ACFSD) in Support of Aircraft-Store Compatibility and Weapons Integration	Jacob Freeman, Air Force SEEK EAGLE Office, Eglin AFB, FL
Characterization and Prediction of Stratospheric Optical Turbulence for DoD Directed Energy Platforms	Frank H. Ruggiero, Air Force Research Laboratory, Hanscom AFB, MA
Computational Chemistry Modeling of the Atmospheric Fate of Toxic Industrial Compounds	Steven Bunte, and Margaret Hurley, Army Research Laboratory, Aberdeen Proving Ground, MD and Daniel Curran, Air Force Technical Applications Center, Patrick AFB, FL
Computational Simulations of Combustion Chamber Dynamics and Hypergolic Gel Propellant Chemistry for Selectable Thrust Liquid/Gel Rocket Engines	Michael Nusca, Army Research Laboratory, Aberdeen Proving Ground, MD
Computational Studies of Naval SONAR and NVRAM Devices	Andrew M. Rappe, University of Pennsylvania, Philadelphia, PA
Computer Design and Simulation of Molecular Devices and Energy Sources for Naval Applications	Mark R. Pederson, Naval Research Laboratory, Washington, DC
Coupled Aircraft/Ship Performance Prediction for Dynamic Interface	Susan Polsky, Naval Air Warfare Center, Aircraft Division, Patuxent River, MD
Coupled CFD/CSM/DPM Modeling of Structure Response to Blast Loading	Joseph D. Baum, Science Applications International Corporation, McLean, VA (Defense Threat Reduction Agency)
DARPA Underwater Express	Robert Kunz, Pennsylvania State University, State College, PA (Office of Naval Research)
Decision Support for Seismic and Acoustic Sensors in Urban Terrain	Stephen Ketcham, USACE Engineer Research and Development Center, Hanover, NH
Design of Energetic Ionic Liquids	Jerry Boatz, Air Force Research Laboratory, Propulsion Directorate, Edwards AFB, CA
Design of Materials for Laser Protection Applications	Ruth Pachter, Air Force Research Laboratory, Materials and Manufacturing Directorate, Wright-Patterson AFB, OH
Direct Numerical Simulations of Active Control for Low-Pressure Turbine Blades	Herman Fasel, University of Arizona, Tucson, AZ (Air Force Office of Scientific Research)
Direct Quantum Mechanical Simulations of Shocked Energetic Materials Supporting Future Force Insensitive Munitions (IM) Requirements	William Mattson, Army Research Laboratory, Aberdeen Proving Ground, MD
Explosive Structure Interaction Effects in Urban Terrain	James T. Baylot, USACE Engineer Research and Development Center, Vicksburg, MS
First-Principle Predictions of Crystal Structure of Energetic Materials	Krzysztof Szalewicz, University of Delaware, Newark, DE (Army Research Office)

Table 2. FY 2007 DoD Challenge Projects—continued

Project Title	Project Leader/Organization
Global Ocean Prediction with HYCOM	Alan Wallcraft, Naval Research Laboratory, Stennis Space Center, MS
High Fidelity Electromagnetic Target Signatures for Combat Identification	Mary Ann Gualtieri, Air Force Research Laboratory, Sensors Directorate, Wright-Patterson AFB, OH
High Resolution Simulation of Full Aircraft Control at Flight Reynolds Numbers	Scott Morton, US Air Force Academy, Colorado Springs, CO
Integrated Analysis of Scramjet Flowpath with Innovative Inlets	Datta Gaitonde, Air Force Research Laboratory, Air Vehicles Directorate, Wright-Patterson AFB, OH
Integration of Simulation and Test for Strike Aircraft Design and Development	Bradford E. Green, Naval Air Systems Command, Patuxent River, MD
Large-Scale Deterministic Predictions of Nonlinear Ocean Wave-Fields	Dick K.P. Yue, Massachusetts Institute of Technology, Cambridge, MA (Office of Naval Research)
Millimeter-Wave Radar Signature Prediction Improvement for Ground Vehicles	William Coburn, Army Research Laboratory, Aberdeen Proving Ground, MD
Modeling Breaking Ship Waves for Design and Analysis of Naval Vessels	Dick K.P. Yue, Massachusetts Institute of Technology, Cambridge, MA (Office of Naval Research)
Modeling of Mine Countermeasure Dart Dispense: Multiple-Body Six Degree-of-Freedom (6-DOF) Computational Fluid Dynamics (CFD) with Collisions	Gary Prybyla, Naval Surface Warfare Center, Indian Head, MD
Molecular Rotors for Nanotechnology	Josef Michl, University of Colorado, Boulder, CO (Army Research Office)
Multidisciplinary Computational Terminal Ballistics for Weapons Systems	Kent Kimsey and David Kleponis, Army Research Laboratory, Aberdeen Proving Ground, MD
Multi-Scale Predictability of High-Impact Weather in the Battlespace Environment	James Doyle, Naval Research Laboratory, Marine Meteorology Division, Monterey, CA
Multiscale Simulations of Nanotubes and Quantum Structures	Jerry Bernholc, North Carolina State University, Raleigh, NC (Office of Naval Research)
Polynitrogen/Nanoaluminum Surface Interactions	Jerry Boatz, Air Force Research Laboratory, Propulsion Directorate, Edwards AFB, CA
Prediction Capability for High-Speed Surface Ships	Joseph Gorski, Naval Surface Warfare Center, Carderock Division, West Bethesda, MD
Protein Design Processes	Kenneth Dill, University of California, San Francisco, (Air Force Research Laboratory, Wright-Patterson AFB, OH)
Signature Evaluation for Thermal Infrared Countermine and IED Detection Systems	John Peters, USACE Engineer Research and Development Center, Vicksburg, MS
Simulation of a Dynamically Maneuvering Unmanned Combat Air Vehicle	Raymond Gordnier, Air Force Research Laboratory, Air Vehicles Directorate, Wright-Patterson AFB, OH
Simulation of Enhanced Explosive Devices in Chambers	John B. Bell, Lawrence Berkeley National Laboratory, Berkeley, CA (Defense Threat Reduction Agency)

Table 2. FY 2007 DoD Challenge Projects—continued

Project Title	Project Leader/Organization
Solidification of Complex High Temperature Structural Analysis	Christopher Woodward, Air Force Research Laboratory, Materials and Manufacturing Directorate, Wright-Patterson AFB, OH
Statistical Fatigue and Residual Strength Analysis of New and Aging Aircraft Structure	Scott Fawaz, US Air Force Academy, Colorado Springs, CO
Time-Accurate Coupled CFD/RBD Simulations of Free Flight Aerodymanics of Guided Weapons	Jubaraj Sahu, Army Research Laboratory, Aberdeen Proving Ground, MD
Toward a High-Resolution Global Coupled Navy Prediction System	Julie McClean, Scripps Institution of Oceanography, La Jolla, CA (Office of Naval Research)
Unsteady, Multidisciplinary Rotorcraft Simulations for Interactional Aerodynamics	Mark Potsdam, Army Research, Development and Engineering Command, Moffett Field, CA
Virtual Prototyping of Directed Energy Weapons	Keith Cartwright, Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, NM

lists the FY 2007 DoD Challenge Projects. Almost all Challenge Project Leaders presented the results of their work at the annual Users Group Conference held in Pittsburgh, PA in June 2007.

Dedicated HPC Project Investments (DHPIs)

The HPCMP also supports high priority computational work requiring dedicated HPC resources. These projects typically need for quick turnaround of the computational work, either real-time or near-real-time calculations, often in support of a specific test event. DHPIs meet such requirements. These small to medium-sized projects require HPC resources that have one or more of the following attributes:

- require access to data or computational resources under time critical constraints that can not tolerate network latency or shared computing;
- real-time signal image processing;
- embedded systems applications;
- dedicated computational resources available immediately as needed;

DHPIs—continued

- early access technology evaluation; or
- require special considerations, including security requirements or unconventional operations.

Examples of the types of projects supported by DHPI resources include:

- real-time analytic and decision support in test and evaluation of land combat systems;
- platform for conducting operational tests of weather research and forecast models;
- real-time global-scale computer-generated forces experimentation;
- real-time hardware-in-the-loop avionics and weapon systems simulations for test and evaluation;
- modeling and simulation of command, control, communication, computers, intelligence, surveillance and reconnaissance (C4ISR) electronic systems under realistic tactical conditions; and
- real-time data imaging of aerial objects.

The HPC systems are procured either through the annual technology insertion process (TI-XX) that acquires new computational capability for the HPCMP Centers, or by providing procurement funding for the systems directly to the user site that proposed the dedicated HPC project.

Four DHPIs awards were made in FY 2007; these include projects to be hosted at: 1) Joint Forces Command (JFCOM), Suffolk, VA, to exploit HPC for use in developing, exploring, testing, and validating 21st century battlespace concepts, 2) Air Force Weather Agency (AFWA), Offutt AFB, NE and Fleet Numerical Meteorology and Oceanography Center (FNMOC), Monterey, CA, to develop a joint ensemble forecast system across a homogeneous computing environment, 3) Naval Research Laboratory (NRL), Washington, DC and Naval Special Warfare Command, Coronado, CA to develop and then provide access to an ultra-fast high-accuracy software model which provides predictions for airborne contaminant transport, and 4) the Office of the Secretary of Defense (Program Analysis and

DHPIs—continued

Evaluation) [OSD (PA&E)], Washington, DC, to provide infrastructure for the seamless execution of the joint warfare system.

User organizations that were awarded HPC resources to support their projects in prior years and reached milestone completion presented reports at the annual Users Group Conference in June 2007. A typical DHPI has a life-cycle of two to three years, depending on the project's established milestones. In FY 2007, three DHPI projects from the following sites were transitioned from HPCMP oversight: Arnold Engineering Development Center, Space and Naval Warfare Systems Center, and Naval Undersea Warfare Center. Accordingly, the following DHPIs are currently under HPCMP oversight:

- MIT-Lincoln Labs, Lexington, MA;
- Command Electronics Research, Development and Engineering Center, Ft. Monmouth, NJ;
- Dugway Proving Grounds, Dugway, UT;
- Air Force Weather Agency, Offutt AFB,NE/Fleet Numerical Meteorological and Oceanographic Center, Monterey, CA;
- Joint Forces Command, Suffolk, VA (J7/J9);
- Naval Research Laboratory, Washington, DC/Naval Special Warfare Command, Coronado, CA; and
- Office of the Secretary of Defense (Program Analysis and Evaluation), Washington, DC.

One DHPI award is planned for FY 2008 using FY 2008 funds. As a result, there are ten sites with active projects. Table 3 lists the DHPIs for FY 2005, FY 2006, FY 2007, and FY 2008.

Table 3. Dedicated HPC Project Investments for fiscal years 05 through 08

FY 2008
Program-Restricted Hydrocode Calculations of Active, Reactive, and Combined Armor Packages, Army Research Laboratory, Aberdeen Proving Ground, MD
FY 2007
Application of HPC to Support Operational Use of CT-Analyst, Naval Research Laboratory, Washington, DC and Navy Special Warfare Command Mission Support Center, Coronado, CA
Distributed Continuous Experimentation Environment, Interactive Joint Features Laboratory, Joint Forces Command, Suffolk, VA
Joint Ensemble Forecast System, Air Force Weather Agency, Offutt AFB, NE and Fleet Numerical Meteorology, Oceanography Center, Monterey, CA
Joint Warfare Systems Server, Office of the Secretary of Defense (Program Analysis & Evaluation), Washington, DC
FY 2006
Interactive Algorithm Development for WMD Defense, Massachusetts Institute of Technology-Lincoln Laboratory, Lexington, MA
Operational, Probabilistic, Numerical Weather prediction at High Resolution, Dugway Proving Ground, UT
Virtual Electronic Battlefield, Command Electronics Research, Development and Engineering Center, Ft. Monmouth, NJ
FY 2005
Aircraft Store Certification, Arnold Engineering Development Center, Arnold AFB, TN and Air Force Seek Eagle Office, Eglin AFB, FL
Concurrent Computation and Utilization Environment, Naval Surface Warfare Center, Bethesda, MD
Space Situational Awareness, Air Force Maui Optical & Supercomputing Site, Kihei, HI
Torpedo Hardware-in-the-Loop Modeling & Simulation, Naval Undersea Warfare Center, Newport, RI

Capability Applications Projects (CAPs)

Starting in FY 2005, the HPCMP made available newly acquired systems for capability applications projects, to test key DoD application codes on a substantial portion of entire HPC systems and to solve large problems quickly. Thus, the goals of capability applications projects are:

- to quantify the degree to which important application codes scale to thousands of processors; and
- to enable new science and technology by applying these codes in dedicated, high-end, capability environments.

The process is an extension of pioneer usage of new HPC systems, but it focuses much more heavily on true capability use for a short time before those systems are put into allocated operational use. It is implemented

CAPs—continued

in two phases: Phase I, which focuses on scalability testing of applications codes proposed for CAPs, and Phase II, during which a subset of successfully tested codes and projects have dedicated access to the newly acquired systems for production work to solve a large, significant problem. The period of time dedicated to this capability workload typically lasts from one to three months.

In FY 2007, CAPs were executed on the NAVO 3,000-processor IBM P5+ (Babbage), the ARL 4,500-processor Linux Networx Xeon cluster (MJM) (both TI-06 acquisitions), the MHPCC 5,100-processor Dell Xeon cluster (Jaws), the ARSC 2,300-processor Sun Opteron cluster (Midnight), and the ASC 9,200-processor SGI Altix (Hawk), a TI-07 acquisition. CAPs were not run on another TI-07 acquisition, the upgrade of the ERDC Cray XT3 from single- to dual-core capability since that system was already in operational use as a single-core system before the upgrade. CAPs are planned on the final TI-07 acquisition, the ERDC Cray XT4, when that system arrives in mid-FY 2008.

For the TI-06 CAP activity executed in FY 2007, 17 proposals were submitted. Of these, 14 were chosen for Phase I projects across the four systems (nine on Babbage, five on MJM, two on Midnight, and four on Jaws). Of these, three were selected for Phase II CAP work on Babbage, two on Jaws, and two on Midnight. These are listed in Table 4. No Phase II CAP work was performed on MJM since it was already quite late for production workload. For the TI-07 CAP activity, 11 proposals were submitted to be run on Hawk in FY 2007 and the ERDC XT4 (to be run in FY 2008). Of these, eight ran Phase I projects on Hawk. Three of these eight were selected for Phase II CAP work on Hawk, and the Phase II work was nearly complete at the end of FY 2007. Reports on the seven Phase II CAPs completed on Babbage, Jaws, and MJM were made at the FY 2007 HPCMP Users Group Conference.

Quantum Algorithms for Computational Fluid Dynamics Simulations of Turbulence

Jeffrey Yepez, Air Force Research Laboratory, Space Vehicles Directorate, Hanscom AFB, MA

Aircraft designs require an understanding of highly nonlinear jet flows, for example accurate and rapid computations of nozzle geometry trust, or flow structures produced by the interaction and merger of co-rotating wing-tip vortices generated in flight. Revolutionary quantum device designs require a related understanding of nonlinear superfluid flows. The similarity of vortex dynamics in quantum and classical fluids is striking. The quantum computing team at the Air Force Research Laboratory and the College of William & Mary continued record setting achievements for the fastest simulations to date using a unique parallel computing approach.

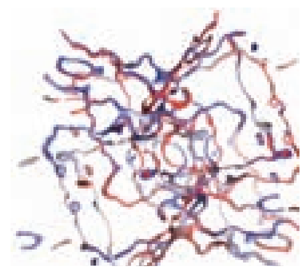

Figure A. Turbulent vortex solitons in a Bose-Einstein condensate (BEC) computed using a quantum lattice gas algorithm on a 512^3 grid on HAWK (SGI Altix 4700) at the AFRL MSRC. The BEC wave function vanishes along the filamentary center of a vortex soliton. Transverse vortex waves (helical Kelvin wave perturbations of a vortex tube away from its cylindrical shape) drive an instability fundamental to a wide class of fluidic morphological transformations.

Strongly correlated fluidic systems were tested through two CAP II quantum algorithm projects, modeling a Bose-Einstein condensate superfluid in Figure A, a near-inviscid Navier-Stokes fluid in Figure B, and magnetohydrodynamics. These large simulations provide a better understanding of the morphological evolution and structural development of turbulence in fluids. A single simulation run takes only a few hours on a supercomputer on a dedicated block of thousands of high performance processors, generating terabytes of data per day.

These codes are the fastest known in the computational physics community and have the best scaling behavior on parallel supercomputers (highest teraflop rates ever achieved by any code). They are the forerunners of quantum codes that will be used in the near future on powerful quantum computers for physical modeling and simulation. With near optimal scaling on parallel computers, unconditional stability (even with complex boundaries), second order accuracy, and less memory and processor time than with other models of turbulence, one can consequently achieve higher nonlinearity (e.g., Reynolds number) than other methods.

This novel class of algorithms was discovered under the Air Force Office of Scientific Research Quantum Computation for Physical Modeling initiative.

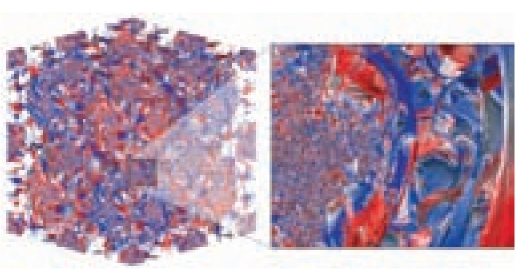

Figure B. Navier-Stokes free shear computed with an entropic lattice gas on a 1600^3 grid on BABBAGE (IBM P5+) using 2,048 processor elements at the NAVO MSRC. Nonlinear dynamics over disparate spatial scales is faithfully captured (inset).

References

1. Yepez, Jeffrey, "Quantum lattice-gas model for computational fluid dynamics." *Physical Review E*, 63(4): 046702, Mar 2001.

2. Chen, Zhiying, Jeffrey Yepez, and David G. Cory, "Simulation of the Burgers equation by NMR quantum-information processing." *Physical Review A*, 74(4): 042321, Oct 2006.

3. Yepez, Jeffrey, "Open quantum system model of the one-dimensional Burgers equation with tunable shear viscosity." *Physical Review A*, 74(4): 042322, Oct 2006.

4. Boghosian, Bruce M., Peter J. Love, Peter V. Coveney, Iliya V. Karlin, Sauro Succi, and Jeffrey Yepez, "Galilean-invariant lattice-Boltzmann models with H theorem." *Physical Review E*, 68(2): 025103, Aug 2003.

5. Keating, Brian, George Vahala, Jeffrey Yepez, Min Soe, and Linda Vahala, "Entropic lattice Boltzmann representations required to recover Navier-Stokes flows." *Physical Review E*, 75(3): 036712, 2007.

Table 4. FY 2007 Phase II Capability Applications Projects

Project Title	Project Leader/Organization
A High-Resolution Simulation of Three Directed Energy Concepts of HPM Munitions: Magnetron with Pulsed Power and Antennas; Airborne Active Denial System (ADS): 95GHz Gyrotron; and RF Counter-Improvised Explosive Device (C-IED)	Keith Cartwright, Air Force Research Laboratory, Kirtland AFB, NM
Size Dependent Plasticity: Search for Strengthening Mechanisms	Christopher Woodward, Air Force Research Laboratory, Wright-Patterson AFB, OH
Self-consistent Modeling of Nonequilibrium Properties of Josephson Junctions	James Freericks, Georgetown University, Washington, DC (Office of Naval Research)
Submarine Crashback Maneuvers	Scot Slimon, Electric Boat Corporation, Groton, CT (United States Navy)
Quantum Algorithms for MHD and Turbulent Flows	George Vahala, College of William and Mary, Williamsburg, VA (Air Force Office of Scientific Research)
Characterization and Prediction of Stratospheric Optical Turbulence	Joseph Werne, NorthWest Research Associates, Boulder, CO (Air Force Research Laboratory, Hanscom AFB, MA)
Understanding and Predicting Shockwave and Turbulent Boundary Layer Interactions	Maria Pino Martin, Princeton University, Princeton, NJ (Air Force Office of Scientific Research)

Computational Research and Engineering Acquisition Tools and Environments (CREATE)

Computer power has grown exponentially during the last 60 years. By 2012, the DoD computational science and engineering community will have access to computers with peak processing speeds in the petaFLOPS (10^{15} FLOPS) range. Computer power of this magnitude offers the potential for DoD computational applications to employ highly accurate numerical methods, include all physical effects known to be important, predict the behavior of a complete system like an entire airplane, and obtain results quickly enough that users can make parameter and sensitivity studies. Such computer power has the potential to give DoD engineers the opportunity to produce improved designs and detect and fix design flaws before major schedule and funding commitments have been made.

Realizing this potential will require development of new software applications to incorporate complete sets of accurate models and exploit the power of these new, massively parallel and very complex supercomputers.

CREATE—
CONTINUED

In the past, DoD relied heavily on application codes developed by other agencies such as the Department of Energy (DOE), National Aeronautics and Space Administration (NASA), academia, and commercial vendors. However, in the future, DoD will need to develop more of its engineering design and analysis tools. The reasons for this are:

- other government agencies are focusing on their own mission requirements;
- the market niche is too small to allow vendors to be commercially successful; and
- developing new large-scale, massively parallel tools is an expensive, long-term effort.

The recent Quadrennial Defense Review (2006, pp. 4, 67–71), and the 2006 Government Accountability Office report (GAO-06-585T, 2006) highlight the need to improve the DoD acquisition process. Consequently, during FY 2006, the HPCMP community proposed the CREATE program as an initiative in the budget process.

Present acquisition programs largely follow a "build, test and break, fix, build, …" methodology. This results in late discovery of design flaws, immature technologies issues, and system integration problems, causing costly rework and redesign that contribute substantially to cost overruns and schedule delays. Optimized engineering designs developed early in the acquisition process using the CREATE tools will substantially reduce costs, shorten schedules, and increase design and program flexibility and agility. Above all, CREATE will improve acquisition program performance by reducing design flaws, allowing development of sound engineering designs quickly and flexibly, and enabling the systems integration engineering process earlier in the acquisition process.

The CREATE program will develop and deploy three sets of advanced computational engineering design tools for acquisition programs:

- military aircraft design,

CREATE—
CONTINUED

- military ship design, and
- antenna design and integration with platforms.

These are illustrated in Figure 8.

The military aircraft design project will develop a design optimization tool to simulate unsteady, separated flow, initially for individual aircraft components and ultimately for an entire aircraft. The military ship design project will develop and deploy accurate physics-based models for navy vessels to address ship shock response and hydrodynamics, and an improved early-stage ship design process, when life-cycle costs are locked in. The antenna design project will build an efficient electromagnetic design code that incorporates modern physics and computational algorithms for high performance computers. This new generation of computational design tools will enable acquisition system engineers to rapidly produce optimized designs for complete systems and make better design decisions than previously possible. A fourth project will provide software development support to these three projects for problem generation, software engineering, collaboration tools, data assessment and analysis, and computational mathematics.

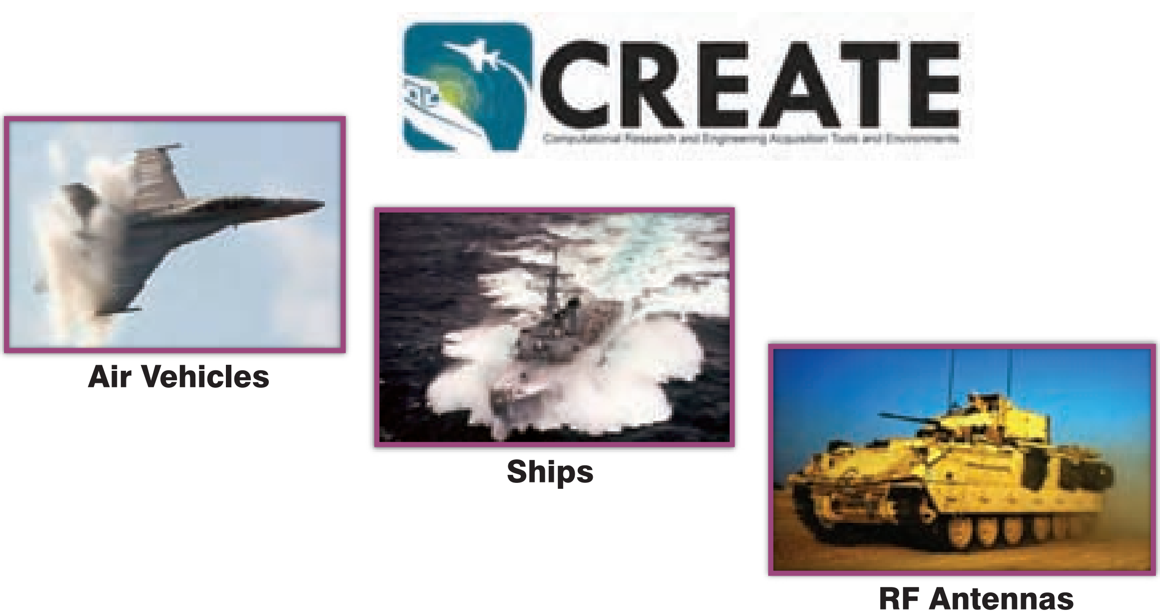

Figure 8. CREATE addresses aircraft design, ship design, and antenna design

CREATE—
CONTINUED

The CREATE projects will be managed to provide enhanced engineering design tools within three to four years after the projects begin. Fully mature tools will be delivered at the end of the 12-year project schedule including two years for planning and project design. The initial tool sets will use supercomputers coming online in 2008–2010. The final tool sets will be designed to be able to exploit the next generation of supercomputers available in 2020.

While CREATE does not officially begin until FY 2008, planning for CREATE began in FY 2007. Planning is based on the "lessons learned" from similar scale projects carried out by nuclear weapons, climate modeling, and other computational communities, and will build on the DoD experience with the HPCMP's institutes and portfolios. The projected funding level is $365M for FY 2008 through FY 2019, with approximately 70% contribution from OSD and 30% from the Services. The tri-Service T&E communities will provide validation experiments and data for the projects as well.

In FY 2007, a formal planning group for CREATE was put together. Sub-teams were created for air vehicles, ships, and radio frequency (RF) antennas, and a fourth team was formed for computational infrastructure. To develop project plans, the teams:

- surveyed the acquisition communities for aircraft, ship, and antenna systems to identify the capability gaps that computational engineering should fill;
- began the development of requirements for computational engineering tools to fill the identified gaps;
- began identifying the computational engineering tool sets that should be developed with the resources available to the CREATE program within the 12-year program life that could meet those requirements and would have the most impact on the acquisition programs; and
- began the development of plans to develop and deploy those tool sets.

CREATE—
CONTINUED

The computational infrastructure team gathered requirements from the three weapon systems projects, assessed the available HPCMP infrastructure, and surveyed the infrastructure that other agencies and industry employ to develop and deploy large-scale computational engineering tools.

The general strategy for tool development for each area has been first to assess the existing computational tool base, augment and enhance those tools, and then support the transition of those tools to the acquisition community. The experience of the users with those tools and the "lessons learned" will then be used to design a new set of tools for use with next generation computers. This will ensure that the development team is connected to the acquisition programs and that the acquisition programs will use the tools as they are improved and developed.

The air vehicle team identified two initial focus areas for FY 2008 starts:

- development of a design code for subsonic and supersonic fixed wing aircraft (Kestrel project, based at the Seek Eagle, Eglin AFB, FL), and
- deployment and testing of existing and new codes in support of existing acquisition programs (Shadow-Ops, led by Naval Air Systems Command (NAVAIR), Patuxent River, MD).

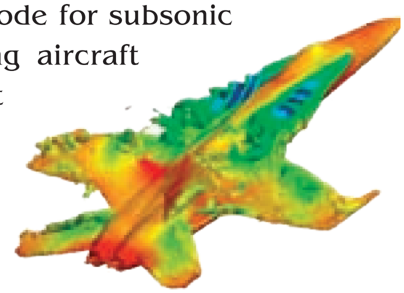

In addition, two focus areas for potential FY 2009 starts are being assessed:

- airframe-propulsion Integration, and
- advanced Concept Development (focusing on early design development).

The air vehicles project will be located at several institutions. The headquarters is at the NAVAIR at Patuxent River, MD. The air vehicles project will be closely coordinated with the Rotorcraft and Air

CREATE—
CONTINUED

Munitions Institutes. The project will work closely with the two institutes to enhance the value of CREATE and institute tools, as well as exploit the close connection of the Institutes to the rotorcraft community and Seek Eagle program.

The ship team identified two initial sub-projects for FY 2008 and a third for a potential FY 2009 start:

- underwater explosion survivability, and
- hydrodynamic performance (stability and slamming)

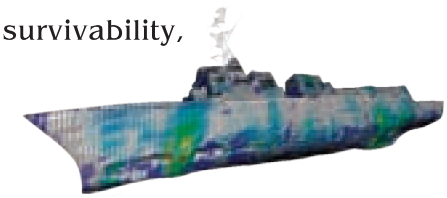

A third project focused on developing an early-stage rapid design capability is being assessed for a FY 2009 start. The Naval Surface Warfare Center, Carderock Division, West Bethesda, MD, is the headquarters for the Ship Design Project.

The RF antenna team initially organized into three sub-projects focused on:

- requirements gathering, monitoring, and tracking,
- testing, and
- product development.

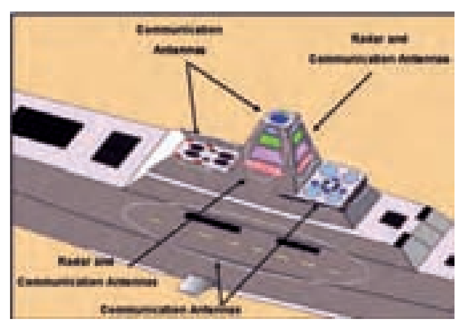

The project headquarters is in the Sensors Directorate, Air Force Research Laboratory, Wright-Patterson AFB, OH. The activities of each RF antenna team will be carried out at a number of performing institutions.

The computational infrastructure project is gathering requirements for the software tools and hardware environment needed by the three major projects. The importance of being able to generate rapidly problems for analysis quickly emerged as the most important infrastructure need, and work has been ongoing to develop a strategy for problem generation. The ability

CREATE—
CONTINUED

to develop rapidly geometry descriptions and generate meshes are key elements of that strategy.

Since no single DoD institution has the staff with the broad collection of skills needed to execute all elements of one project, each project is multi-institutional. Success will require that the non-collocated staff work together as a closely integrated team. To facilitate this, the CREATE program has made an early effort to identify and deploy collaboration tools, including good video and audio conferencing tools and servers.

Much of the effort in FY 2007 was devoted to identifying and filling the key leadership positions. Nominations were sought from the senior service acquisition executives as part of a national search for the four project managers. The applicants and nominees were assessed and four project managers were selected (Figure 9).

The search for key program office staff and deputy project managers was completed in FY 2007. CREATE is now actively filling out the initial development teams. The CREATE OSD budget ramps up from $10M in FY 2008 to $26M in FY 2013, and the Services are beginning to provide their share of support.

During the latter half of FY 2007, CREATE effort began preparation of three major planning documents for the program and for each project:

- initial capabilities document (ICD) that identifies the gaps in present acquisition processes that computational engineering can fill. This document defines the high level requirements for the CREATE projects;
- technical planning document (TPD) that specifies the software tool sets to be developed by the CREATE projects, and describes the plan for developing the software tools and the project management structure and processes; and

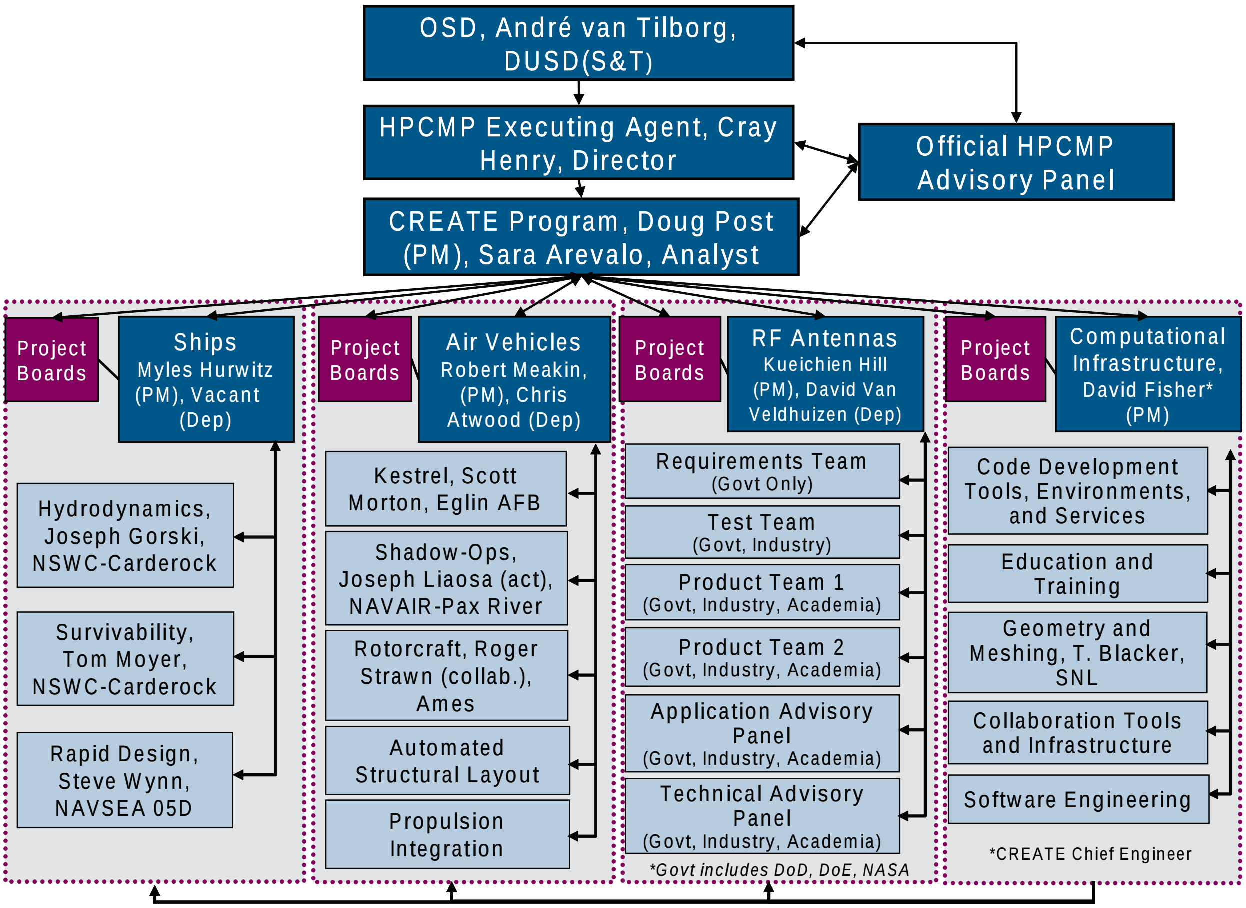

Figure 9. CREATE program organization

CREATE—
CONTINUED

- financial resources plan that describes the financial and staffing resources for the project and the organization and processes for distributing the resources, managing the staff, and placing and funding contracts. These will include developing memorandums of agreement and memorandums of understanding between CREATE and the DoD Services and other government agencies that will be collaborating with the CREATE program.

Initial versions of these documents have been developed and are being reviewed.

CREATE does not have the resources to provide the entire computational and organizational infrastructure needed. To provide the needed infrastructure, the

CREATE—
CONTINUED

project is establishing collaborations with other government agencies and with industry to obtain the needed infrastructure support. Initial, very promising discussions have been held with NASA, DOE (both National Nuclear Security Administration and the Office of Science), and the National Science Foundation. Fruitful interactions with industry are essential. A key goal is for industry to adopt CREATE tools as part of their design process. Outreach to industry has also begun and the initial reception by industry has been very positive and welcoming.

CREATE goals for FY 2008 include:

- completion of the initial plans for the CREATE program and projects;
- completion of the software design for the initial CREATE tool sets;
- validation of the ICDs by the relevant acquisition and service organizations;
- complete establishment of the initial project teams;
- complete establishment of the oversight, review and advisory processes for the CREATE program and projects;
- establishment and implementation of processes for program and project management, and distribution and control of financial resources;
- completion of support staffing for the program and project offices;
- definition of a strategy for problem generation and establishment of sub-project to begin implementation of that strategy;
- initial deployment of a functional set of collaboration tools, including servers for program documents, configuration management of codes, and project communication;
- establishment of collaborative agreements with other government agencies for software development infrastructure; and

CREATE—CONTINUED

- definition of the verification and validation strategy and establishment of collaborative agreements with the relevant service T&E communities to execute the needed validation experiments.

Highlights of Impact in FY 2007

The High Performance Computing Modernization Program provides some of the tools the Department needs to address defense problems.

Some problems are of immediate concern, while others are of longer-term interest. Hence, program investments impact both short-term and long-term issues. The following vignettes serve as overviews of some highlights that occurred in FY 2007.

Army-led Rotorcraft Institute Helps Early Deployment of CV-22 Osprey

Thanks in part to analyses performed by the High Performance Computing Institute for Advanced Rotorcraft Modeling and Simulation (HI-ARMS), the CV-22 Osprey may be able enter service in Iraq far in advance of the originally planned date of FY 2009.

Accelerated deployment of the Bell Boeing CV-22 Osprey may be necessary because the Sikorksy MH-53 Pave Low, presently used by Air Force Special Operations, is nearing the end of its service life.

“Three weeks ago we lost an Air Force Special Operations Command MH-53 Pave Low helicopter. Thankfully, there were no serious injuries to the crew, but there easily could have been a different outcome,” said *Gen. T. Michael Moseley, USAF Chief of Staff*, on Wednesday, September 26, 2007 at the Air Force Association’s 2007 Air & Space Conference and Technology Exposition. Gen. Moseley added,

> *“Now the MH-53 is a great airplane; … But we’ve been retiring a few of them each year; we’ve lost six in combat, and now we have 21 left in the inventory. And over the course of the next year-and-a-half, we’ll eventually retire the remaining MH-53s. … These helicopters were to be replaced by CV-22s, but delays in appropriation, production, and delivery leave us currently with only six CV-22s. In light of the slow production rate, we are looking at options to deploy the CV-22 pre-initial operational capability, in fact, to prevent a mission impacting loss on US Special Ops Command in its vertical lift and covert and searching capability.”*

The CV-22 is the special operations variant of the Marine MV-22. The CV-22 and MV-22 are about 85 percent common. For Special Operations Forces (SOF) missions, the CV-22 has special equipment, such as terrain-following and terrain-avoidance radar, an advanced electronic warfare suite, high-capacity fuel tanks, and state-of-the-art avionics.

The CV-22s have a number of special operational requirements. Among these is a requirement to dump fuel safely not only in emergencies but also as standard operating procedure for mission management. This is one of the issues addressed by HI-ARMS. When the first fuel dump maneuver

a) Fuel dump maneuver　　b) CFD simulation　　c) Shortened tube

CV-22: This chart depicts a fuel entrainment (into the empennage) problem associated with a required fuel dump maneuver (a). computational fluid dynamics (CFD) was used to demonstrate the ability to correctly predict the phenomenon (b). Visualization of the flow-field suggested to the analyst a novel and very inexpensive solution to the problem—shorten the dump-tube so that the release point occurs upstream of flow separation and parallel to the freestream (c). Simulations show that no fuel entrainment occurs with this release point.

was attempted in flight test, fuel impinged on the fuselage ramp and intruded into the horizontal tail, where hot electronics gear is located. HI-ARMS used CFD to demonstrate the ability to correctly predict the phenomenon, and to suggest a very simple and inexpensive remedy, namely, shorten the fuel dump tube. A wide range of potential influences across the flight envelope were able to be investigated using CFD (e.g., speed, angle of attack, flaps, sideslip, power effects, and CV-22 specific electronics) at a fraction of the cost of flight testing.

The CV-22 Osprey has twice the top speed, three times the payload, and up to five times the range of existing SOF rotorcraft. This results from the tilt-rotor design that enables the aircraft to take off and land like a helicopter but rotate its engine nacelles forward while in flight to achieve the speed of a turboprop airplane.

The Air Force's first operational CV-22 Osprey tilt-rotor aircraft hovers upon arrival at Kirtland Air Force Base, NM. The aircraft will be flown by Airmen of the 71st Special Operations Squadron (71 SOS), commanded by Lt Col James Cardoso.

It promises to bring major advances in combat capability, and may well revolutionize the way SOF missions are conducted. With higher top speed, longer range, and greater carrying capacity than today's helicopters, the Osprey would permit SOF to undertake more missions during "one period of darkness." Moreover, Osprey would still be able to take off and land vertically in tight spots.

The CV-22 may also be assigned to handle other combat missions. With its speed, range, and internal cargo capacity, the Osprey could be adapted to meet numerous other missions, combat and noncombat. These missions include Combat Search and Rescue, disaster relief, aerial refueling, air medical evacuation, and executive transport. "I feel that, once we're successful [deploying the CV-22], … this airplane could compete for some [other] missions," said *Gen. (now retired) Charles R. Holland, (former) Commander of United States Special Operations Command*.

High Performance Computing Improves Anti-Submarine Warfare by a Factor of Two

Anti-submarine warfare (ASW) planners use decision aids to help decide where best to place acoustic sensors to achieve the desired level of confidence in detecting enemy submarines. Improvements in these decision aids—a direct result of high performance computing—permit a reduction by as much as a factor of two in the number of sensors required, permitting redeployment of resources for other critical missions.

Recent exercises in the Pacific have shown that these decision aids—coupled with improved instrumentation—greatly benefit ASW missions. The Operational Navy is very enthusiastic about these new capabilities. (See sidebar.)

The Commander Task Force 74 sent the following testimonial to Commander, Navy Meteorology and Oceanography Command

"Your teams' detailed oceanographic analysis significantly improved ASW planning and execution throughout this exercise. The coordinated efforts demonstrated by your teams were superb. The use of latest technologies (i.e., NCOM and Performance Surfaces) proved a remarkable new capability in environmental analysis and ASW planning."

Message from RADM McAneny to NOAC Yoko, NOAC Stennis, and NAVO – CTF 74 140310ZAUG 07

Adversary submarines are a threat to the Fleet and commerce alike. Finding such submarines in the vast ocean is a key element of anti-submarine warfare. Sonar is an important tool for detecting submarines acoustically. Sound propagation in the ocean is a very complicated phenomenon. It depends on the state of the ocean. These state variables include temperature, turbulence, salinity, and density.

Dynamic ocean processes produce small thermal variations that induce spatial and temporal variability in the ocean's index of refraction and in the spatial scale along an acoustic propagation path.

Because of ocean variability, the probability that a submarine will be detected in a specific location depends on the location (direction and distance) of the detector. The probability of detection also depends on the quality and sensitivity of the sensor.

Thanks to supercomputer-enabled advances in ocean modeling, the Navy weather community provides better support to the ASW community. The Global Navy Coastal Ocean Model (NCOM) now gives more accurate vertical resolution resulting in better 3D forecasts of temperature, salinity, and currents. High-resolution versions of NCOM are embedded within lower resolution, wider-coverage versions to provide accurate speed of sound profiles at almost any location in the world. This is shown in Figure A.

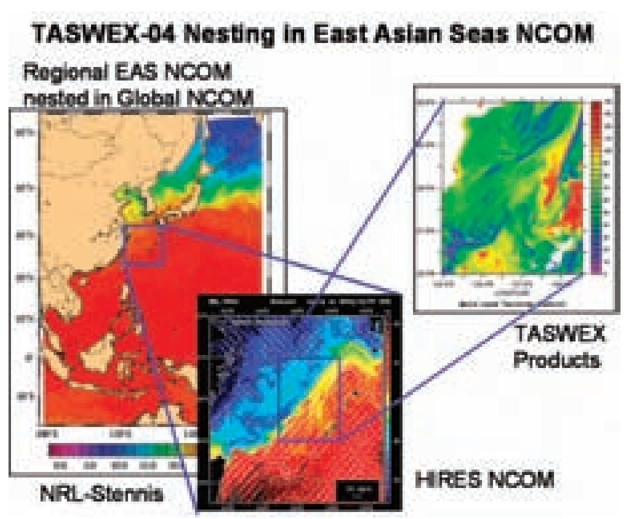

Figure A. Nesting NCOM for high local resolution[1]

Probability of Detection (Pd) Map, Varying Detector Depth Plan View (Op Area, 0 to 200 m) for 17 July 06, Threat at 200 ft, Detector between 0 and 650 ft

Probability of Detection (Pd) Map Varying in Time Volume View (Op Area by 200 m in depth) for 5–25 July 06 Threat at 200 ft, Detector between 0 and 650 ft

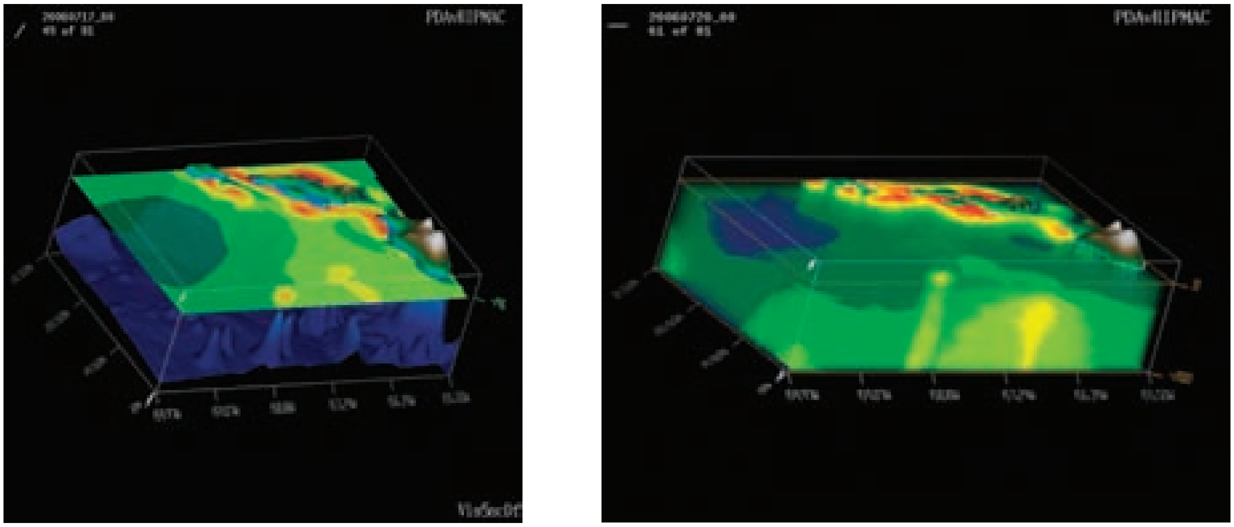

Figure B. Probability of detection map[2]

To do ocean forecasting, three elements are needed:

- the fundamental equations describing the physics,
- the geometry of the problem, and
- initial conditions and/or boundary conditions.

In ocean modeling, boundary conditions are obtained from instrumentation, other models, and/or climatology. The best available boundary conditions for regional models are obtained from the results of the global computations performed on HPC systems.

The speed of sound profiles that are necessary for ASW modeling tools are in turn calculated from the computations. From the sound profiles, 3D probability of detection maps are obtained as shown in Figure B.

References

1. MSRC Ocean Modeling Requirements at the Naval Oceanographic Office, (NAVO), Dr. Frank L. Bub, Technical Lead, Ocean Modeling (NP1M)

2. High performance CWO model at NAVO, Dr. John Harding, Chief Scientist, Navy Meteorology and Oceanography Command

Key Models Used to Produce Ocean Forecasts

- Navy Layered Ocean Model (NLOM),
- Navy Operational Global Atmospheric Prediction System (NOGAPS)
- Navy Coastal Ocean Model (NCOM)
- Hybrid Coordinate Ocean Model (HYCOM)
- Modular Ocean Data Assimilation (MODAS)
- Relo NCOM-Relocatable NCOM

Section 2

Performance Results

Performance Results

FY 2007 Operations and Performance

The work accomplished in FY 2007 continues to assists the DoD S&T and T&E communities to provide support to the warfighter. Benefits are both near-term, that is, within the present fiscal year, and far-term.

This section begins with a discussion of a recent return on investment study. The rest of the section is organized by the program goals given on pages 8 and 9.

Determining the DoD HPCMP Value to the Warfighter—Return on Investment (ROI)

The DoD HPCMP resources are a fundamental enabling technology for developing future capabilities and for responding to immediate combat threats. Over the years, there have been many examples that confirm that the HPCMP resources provide high value for the warfighter. Recently, this value was quantitatively demonstrated.

In FY 2007, the HPCMP continued its process to quantify the program's value. It conducted an ROI study of several important projects. This study examined the climate/weather/ocean modeling and simulation (CWO) portfolio to validate further the ROI process and methodology, and to determine the value of this portfolio to the DoD.

The Joint Interoperability Test Command was the lead in interviewing program users, acquisition officers, and theater warfighters to collect and analyze the data. The National Defense University and the Defense

ROI—CONTINUED

Acquisition University validated the ROI methodology and made recommendations to improve future ROI analyses.

For this phase of the ROI process, the following projects were evaluated:

- anti-submarine warfare,
- tropical cyclone prediction,
- aircraft routing,
- reduction in personnel levels,
- dust/sand storm prediction,
- Global Positioning System (GPS) accuracy improvement, and
- divers/mine warfare.

Validated quantitative benefits were obtained for the first four of these areas.

Approximately 70 people were interviewed. They included principal investigators, computational scientists and engineers, project managers, and military personnel who have used real-time tactical aids developed with HPCMP support.

The investment cost for each project using HPCMP Centers was determined by using the program budget allocated for each major component of the program: HPC Centers, Networking and Information Assurance, and Software Applications Support divided by the number of central processing unit (CPU) hours used for that year, and then allocated proportionally to each project's CPU usage. For projects using dedicated systems, the investment was the funding that the Program provided for that system.

Since 1994, the HPCMP has invested more than $400M in providing computational resources for CWO. Using these resources, DoD scientists and engineers have developed more accurate climate, weather, and ocean models. These new models, combined with better sensors, are worth many times more than the investment. For example, today's four-day weather forecast is as accurate as a three-day forecast was 10

ROI—CONTINUED

years ago. The following examples are specific CWO benefits determined by the ROI study thus far.

Anti-Submarine Warfare

Models developed on HPCMP resources generate the acoustic databases employed by acoustic sensors used to detect enemy submarines. The combination of more accurate ocean models and underwater instrumentation has improved the probability of detecting an enemy submarine. This improvement means that as many as two of the four submarines now required to perform certain barrier missions can be redeployed for other missions.

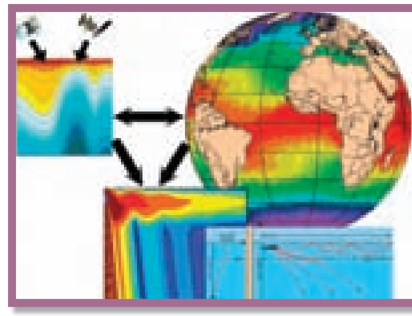

Tropical Cyclone Prediction

Navy Operational Global Atmospheric Prediction System (NOGAPS) is generally the most accurate predictor of tropical cyclone paths and is one of the cornerstones of the consensus models used in official National Hurricane Center forecasts. The HPCMP resource-derived applications help predict the path, intensity, areas for evacuation, and landfall of tropical cyclones. The cyclone prediction models offer the potential to save lives, avoid unnecessary evacuations, effectively identify threatened areas, and save billions of dollars in cyclone preparations and precautions.

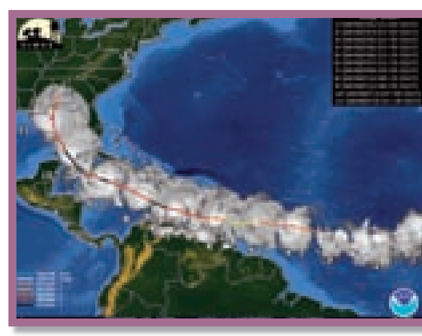

ROI—CONTINUED

"'Owning' the weather became more of a reality with the arrival of new computer modeling tools, enhanced high-resolution satellite imagery, and reachback communications."—Dr. Frank Bub, NAVO Ocean Predictions Group

Aircraft Routing

Weather forecast data from the Navy's numerical weather prediction applications, NOGAPS and Coupled Ocean/Atmosphere Mesoscale Prediction System, have been greatly enhanced using HPCMP resources. The forecast data from these applications provide input to the Optimum Path Aircraft Routing System. Considerable DoD fuel savings have been documented as a result of using the weather data in planning for optimal aircraft routing.

Reduction in Personnel Levels

The combination of HPC models developed on HPCMP resources with centralized reachback cells, improved wide-band information technology, and interactive, tailored meteorological information has allowed significant manpower reduction in the Commander, Naval Oceanography and Meteorological Center (CNMOC) Organization. CNMOC management attributes this reduction to the availability of HPCMP resources at the Naval Oceanographic Office Major Shared Resource Center along with the HPC models.

Dust/Sand Storm Prediction

During Desert Storm (1991), dust storms in Iraq and Kuwait greatly hampered the ability of coalition forces to perform their mission. In subsequent years, the HPCMP resources were used to develop aerosol and dust transport models that can be embedded in weather

ROI—CONTINUED

prediction models to forecast accurately dust storms. The Naval Research Laboratory's Aerosol Analysis and Prediction System and the Air Force Weather Agency's Dust Transport Application are successfully used to predict dust storms in Iraq today. These predictions increase the probability of optimal weapons selection, increase the number of successful sorties, and save fuel and time by avoiding aborted missions.

Global Positioning System (GPS) Accuracy Improvement

Calculating the earth rotation rate is essential for the Global Positioning System (GPS) accurately to determine position. The earth rotation changes as a function of the distribution of the total mass of the earth and its atmosphere. Since the distribution of the atmospheric mass varies, it can have a significant impact on the rate of the earth's rotation. Recently, the US Naval Observatory has been receiving data from the NOGAPS application, providing the distribution of mass in the atmosphere for computing the earth's rate of rotation. This correction to the earth's rotation rate has resulted in a 10–15 percent improvement in GPS accuracy and should allow for improvements in the performance of precision-guided weapons and space-based systems.

Divers/Mine Warfare

The Naval Oceanographic Office used HPCMP resources to provide tidal information for divers performing special missions in Operation Iraqi Freedom. Models

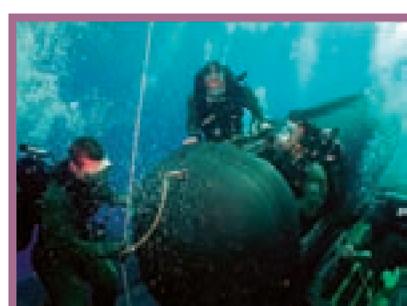

ROI—CONTINUED

that were run on HPCMP resources include the Advanced CIRCulation Model for Shelves, Coastal Seas, and Estuaries and the PCTides model. The models determine the optimum time to dive by predicting the minimum current and maximum depth of the tides.

Summary

Although these results will be refined as more data are collected, enough information has been gathered to make it clear that 1) HPCMP is more than paying its way in providing benefits to the warfighter, and 2) without HPCMP resources, some high-priority warfighter needs could not be satisfied.

HPCMP's support to CWO yields a $9–$18 return for every dollar invested. The breakout of the benefits is shown in the following table.

	Based on 10-year projections:	
CWO Benefit Areas	**Lower Bound**	**Upper Bound**
Anti-Submarine Warfare	$3,236M	$6,471M
Cyclones: Civilian	$419M	$448M
Cyclones: Military	$25M	$74M
Personnel Savings	$559M	$624M
Aircraft Routing	$33M	$462M
Total Return	$4,272M	$8,079M
Investment Total	**$417M**	
ROI	**925%**	**1,839%**

When the previous FY 2006 Armor/Anti-Armor study is combined with the CWO study, the return is $7–$15 for every dollar invested.

		Benefits	
Application	**Investment**	**Lower Bound**	**Upper Bound**
Armor/Anti-Armor	$136M	$487M	$935M
CWO	$417M	$4,272M	$8,079M
Total	**$553M**	**$4,758M**	**$9,014M**
ROI		**760%**	**1,530%**

GOAL 1: ACQUIRE, DEPLOY, OPERATE AND MAINTAIN BEST-VALUE SUPERCOMPUTERS

The program provides high performance computing capabilities to the DoD S&T and T&E communities through three modes:

- major shared resource centers (MSRCs);
- allocated distributed centers (ADCs); and
- dedicated HPC project investments (DHPIs).

Major Shared Resource Centers (MSRCs)

The MSRCs are very large centers that provide leading-edge, high performance computational resources, data storage, data interpretation, and HPC technical expertise to the defense community. The MSRCs are "purple", that is, they serve all DoD Services and Agencies without regard to their location or supporting organization. This policy was adopted (and continues today) in the FY 1993 Implementation Plan, published in December 1992. One of the "core functions required" included:

"A fair-share job scheduling type of algorithm to implement Service/Agency shares. The scheduler must guarantee the agreed DoD percentage access, while making unclaimed time intended for any party available proportionally to all other parties. Periodic reports showing usage requests and usage allocations for local and DoD-wide researchers will be submitted for verification of fair allocation of the agreed DoD percentage, and for corrective action as required."

The fair-share allocation of a 30/30/30/10 dates from January 1993. The 30/30/30/10 split means that the three Services each receive 30 percent of the resources, and the defense agencies 10 percent.

Goal 1—continued

The MSRCs are located at four government installations; they are listed below and shown in Figure 10:

- US Army Research Laboratory (ARL), Aberdeen Proving Ground, MD;
- Aeronautical Systems Center (ASC), Wright-Patterson AFB, Dayton, OH; (In second quarter FY 2008, the MSRC at Wright-Patterson AFB will be assigned to the Air Force Research Laboratory)
- US Army Corps of Engineers (USACE) Engineer Research and Development Center (ERDC), Vicksburg, MS; and
- Naval Oceanographic Office (NAVO), Stennis Space Center, MS.

At the beginning of FY 2007, the HPC systems at the four MSRCs had a total computational capability of 224 teraFLOPS (i.e., the capability to perform 224 trillion mathematical operations per second). During FY 2007, the HPCMP procured three very large systems for deployment at two of the MSRCs (ASC and ERDC). These new systems have a computational capability of 132 teraFLOPS. At the end of FY 2007, the total capability of the HPC systems at the four MSRCs stands at 341 teraFLOPS. The bars in Figure 11 show the computational growth in teraFLOPS as well as Habus at the four centers over the past seven years, as of the end of the fiscal year. (See callout on page 12 for a definition of a Habu.)

Allocated Distributed Centers (ADCs)

To complement the computational capacity of the MSRCs, the HPCMP also supports four "mid-sized" centers that provide additional computational resources to DoD researchers. These centers are the ADCs. From the DoD's perspective, ADCs function like smaller scale MSRCs but, in addition to the DoD, may serve other customers as well. The four centers are listed below and shown in Figure 10:

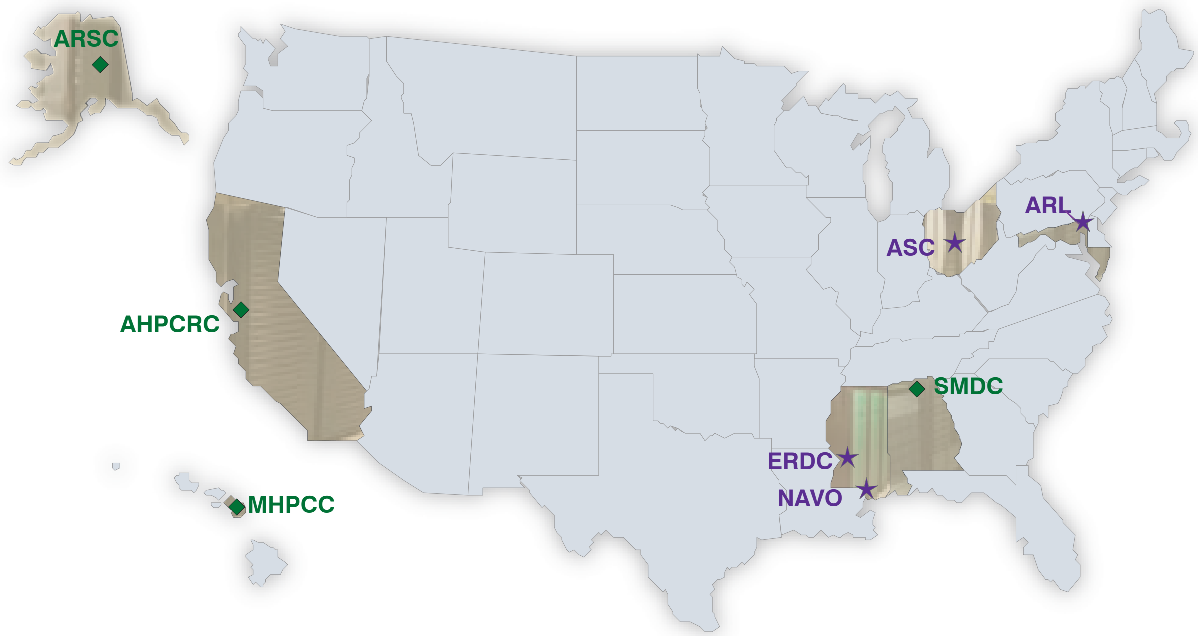

Figure 10. Location of major shared resource centers (purple) and allocated distributed centers (green)

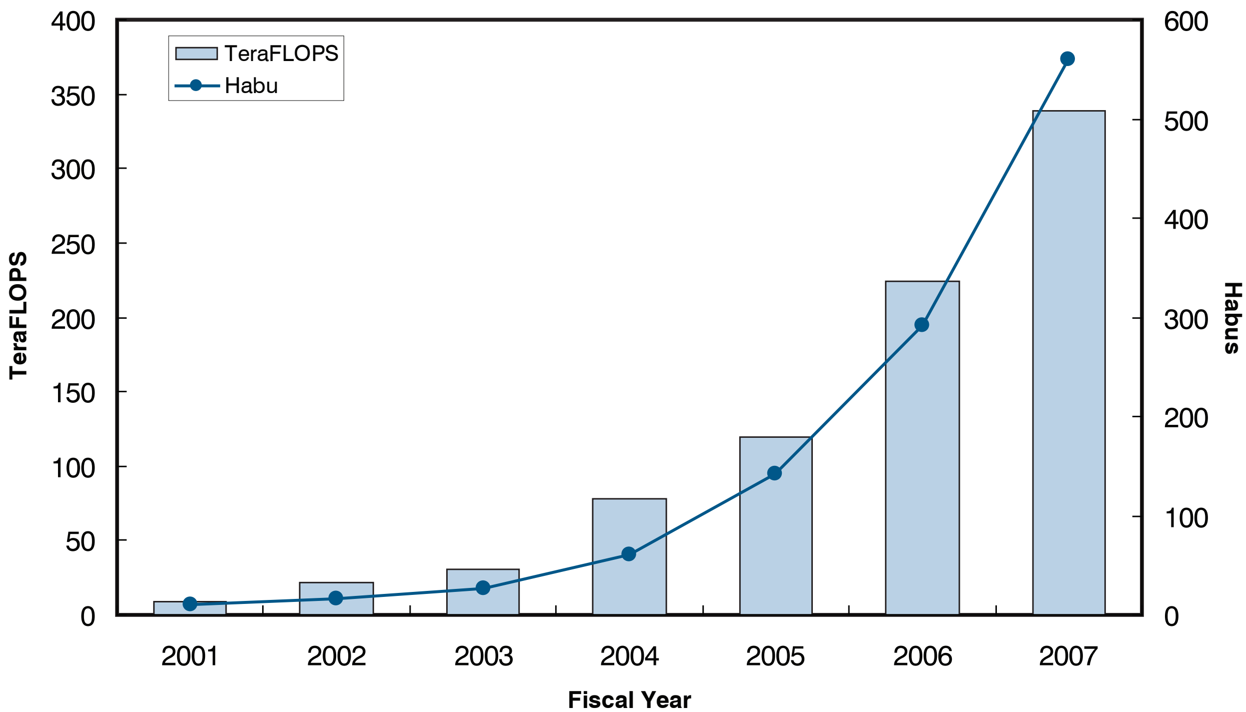

Figure 11. Growth in capability of the MSRCs as of the end of the fiscal year

Goal 1—continued

- Arctic Region Supercomputing Center (ARSC), Fairbanks, AK;
- Maui High Performance Computing Center (MHPCC), Air Force Maui Optical and Supercomputing Site (AMOS), Kihei, HI;
- Army High Performance Computing Research Center (AHPCRC), Moffett Field, CA; and
- Army Space and Missile Defense Command (SMDC), Huntsville, AL.

In FY 2007 the ARSC ADC continued to support open literature, DoD basic research. The academic community of users, whose research is supported by the Offices of Research in the Defense Services, has difficulty quickly obtaining the access clearances needed to use the systems located at the MSRCs. This operational model allows the ARSC to mix non-DoD

Overarching Governing Infrastructure for Centers

The HPCMP community is a geographically distributed one with a valuable diversity of local skills and capabilities that must be captured in support of community ends. To achieve this transfer of best practices and innovations, an overarching Centers' governing infrastructure was created:

- the Center directors of the four MSRCs, plus ARSC, and MHPCC are members of the Centers Board of Directors (CBoD) for the HPCMP Centers capability;
- the technical specialists that design, build, and implement the solutions comprise the Engineering Design and Process Management Board (ED&PMB); and
- the Distributed Implementation and Operations Team (DIOT), are a group of individuals at each of the Centers positioned to sustain the capability.

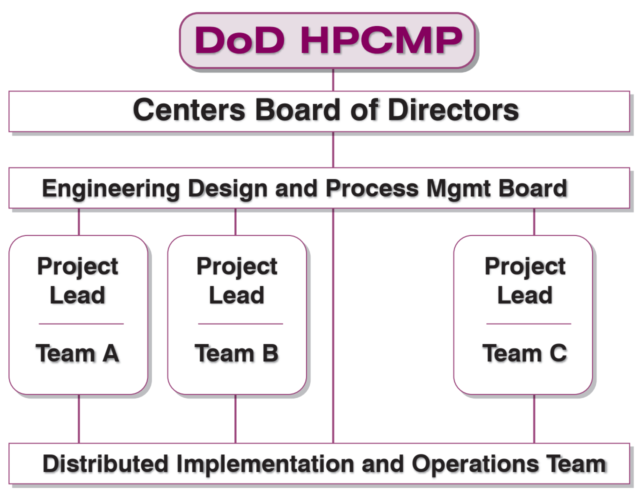

This governing infrastructure was developed in early FY 2004 and chartered in March 2004. The adjoining figure shows the organization of the infrastructure. The CBoD meets several times per year to initiate new investigations and to monitor ongoing cross-center initiatives. The other two teams address the initiatives. Unified direction from the CBoD has helped to keep the ED&PMB and the DIOT on focus.

GOAL 1—CONTINUED

university related work and DoD open literature work on the same systems; a win-win example of how the DoD leverages the use of ADCs.

Collectively, the ADCs have several large HPC systems, including a 952 dual-core Opteron processor SUN system at ARSC and a 5,120 dual-core Xeon processor Dell system at MHPCC, both added in FY 2006. The ADCs provide a total of 91.6 teraFLOPS (131.0 Habus) of computational capability to the HPCMP. Adding this computational power to the capability located at the MSRCs at the end of FY 2007, the HPCMP total capability increases from 341 teraFLOPS (564 Habus) to 433 teraFLOPS (695 Habus).

Summary

The hardware and software acquisition budget for the MSRCs over the last eight years has had zero real growth. In FY 2001, the Program implemented an acquisition process whereby all MSRC systems are acquired through consolidated, large competitively-awarded contracts. The leveraging of volume purchasing power combined with technology advances commensurate with Moore's law (an observation made by the former Chief Executive Officer of Intel Corporation that the number of transistors on a silicon chip doubles every 18 months) has provided the HPCMP with computational capabilities that exceed traditional growth curves.

Goal 2: Acquire, develop, deploy and support software applications and computational work environments that enable critical DoD research, development, and test challenges to be analyzed and solved

The Software Applications Support (SAS) component supports Goal 2 above. SAS presently consists of three major efforts: portfolios, software applications institutes, and user productivity enhancement and technology transfer (PET), formerly known as programming environment and training. In FY 2007, we planned a fourth effort, called CREATE, previously discussed in Section 1 of this report, to build engineering design analysis tools for aircraft, ships, and antennas. Execution of CREATE will start in FY 2008.

The ultimate aim of SAS is to provide DoD scientists and engineers with the capability for modeling and simulating the physical world to facilitate the design, development, test, and deployment of superior weapons systems, thereby allowing our soldiers, sailors, marines, and airmen to be prepared better through training, tactics, and support systems.

Portfolios

The DoD's increasing reliance on technology as a force-multiplier translates into increasingly complex multidisciplinary problems seeking more exact and faster times to solution. Also, on the rise, is the Services' dependence upon immediate technological solutions to new enemy tactics and threats. Today's problems require the integration of multiple computational disciplines. The HPCMP initiated tightly and loosely-integrated portfolios to focus on such multidisciplinary challenges. Portfolios provide efficient, scalable, portable software codes, algorithms, tools, and models and simulations that are needed by a large number of

GOAL 2—CONTINUED

S&T and T&E scientists and engineers to accelerate technological solutions. Portfolio development teams span DoD Services and Agencies and include algorithm developers, applications specialists, computational scientists, computer scientists and engineers, and end users.

Developing software for scalable HPC systems remains challenging and labor intensive. The HPCMP helps the DoD take advantage of existing and future computing and communications capabilities by building software with an emphasis on reusability, scalability, portability, and maintainability. The resultant applications greatly advance the design, acquisition, and utilization of military technologies and aid in the development of improved military capability for the 21st century. In addition, this initiative is producing a new generation of world-class scientists and engineers trained in scalable software techniques to reduce the future costs of doing business and increase our defense capabilities.

The portfolios, illustrated in Figure 12, focus on specific themes that encompass multiple CTAs and cross Service and Agency boundaries.

Collaborative Simulation and Testing (CST)

The CST portfolio is inserting high fidelity modeling and simulation technologies into rotary wing, fixed wing and airdrop design, analysis, and test and evaluation throughout the product's life-cycle. The portfolio provides the warfighter with fast, state-of-the-art analysis by introducing web-based plug-and-play integration of multi-disciplinary codes on HPC systems.

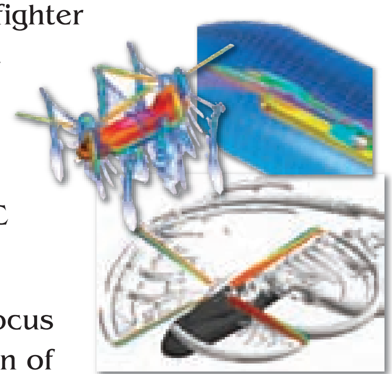

The portfolio's rotary wing focus area is optimizing the design of the CH-47 and other helicopters through higher fidelity predictions of airloads and performance. Aerodynamic and

Goal 2—continued

structural codes have been integrated on HPC platforms, providing a state-of-the-art aeroelastic analysis capability, leading to aircraft weight reductions.

The fixed wing focus area is increasing the compatibility of engines with airframes, including Joint Strike Fighter, through simulations on HPC platforms. High fidelity modeling of the external aircraft flow is integrated with flow predictions inside the engine to provide enhanced performance through an improved design and analysis tool for modern fan systems.

Finally, the airdrop focus area is transforming the design of parachute systems by providing HPC-based analysis capability. Computational fluid dynamics and computational structural dynamics codes have been integrated to enable a low-cost redesign of an existing parachute to increase payload weight with a minimal increase in parachute system weight.

This portfolio is expected to complete work in early FY 2008.

Multiphase Flow Target Response (MFT)

The MFT portfolio is integrating state-of-the-art, physics-based codes from DoD and the DOE to provide the DoD acquisition community a common and supported toolset that delivers predictions for multiphase blast and enhanced blast weapons effects in complex military operations on urbanized terrain environments. During FY 2006, the portfolio integrated multiphase flow physics into the DOE codes, CTH, and ALE3D. A framework for turbulent mixing, gas phase chemistry and metal combustion modeling has been developed and implemented into CTH and ALE3D. The portfolio is also developing tools for insensitive munitions characterization by integrating particle methods and material models

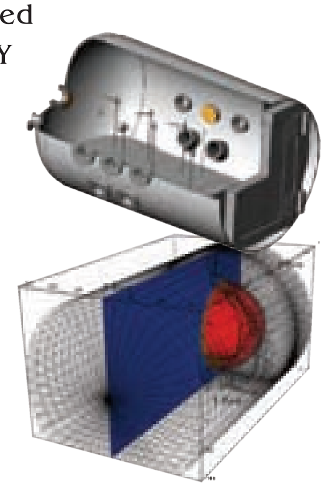

GOAL 2—CONTINUED

into PRESTO, a DOE code designed for problems with large deformations, nonlinear material behavior, and contact.

In FY 2007, the portfolio completed alpha testing of inert multiphase flow functionality for CTH and ALE3D. Work will proceed into the first quarter of FY 2008. By the end of the quarter, two-way coupling of CTH and Presto will be completed as will demonstrations of reacting multiphase flow and structural dynamics/ multi-axial loading. We expect the portfolio will wrap-up with beta testing by the end of that fiscal quarter.

Virtual Electromagnetic Design (VED)

This highly integrated portfolio gives DoD the ability to design from first-principle electromagnetics in-situ wideband, multi-functional antennas and rough surface scattering solutions for a wide range of DoD functions including communication, acquisition, target identification, surveillance, and electronic attack. The VED portfolio and other new improved concurrent electromagnetic particle in cell code developments permit end-to-end integrated system modeling, from pulsed power to antenna and platform. This enables time-critical counter-improvised explosive device simulation, avoiding $12M and 18 months experimental cost. It further impacts airborne counter electronics efforts, avoiding expensive flight tests and reducing electromagnetic compatibility/electromagnetic interference-related problems.

During FY 2007, beta testing was successfully completed on approximately 85% of the VED codes. We expect that all testing will wrap-up by the end of the first quarter of FY 2008.

Evolution of Software Applications Support within the HPCMP

A timeline is shown below that graphically depicts how the DoD software applications programs are transitioning from CTA focused activities to ones that will lead to tightly integrated, multidisciplinary codes that tackle some of the most comprehensive and complex problems facing the DoD warfighters today. Both paradigms have evolved from individual software projects for applications codes from the mid-nineties, where the efforts focused on enhancing DoD applications codes originating up to several decades earlier. These codes were enhanced to become more robust and execute efficiently on scalable hardware coming on line in the mid-to-late nineties. From the beginning of the software applications efforts in 1998 until today, the DoD has completed over 100 projects involving many hundreds of codes; this was a great boon to the weapons development, testing, and warfighting communities. These efforts improved the speed, complexity, and accuracy of military simulations in materials for combat platforms, space and earth weather prediction, littoral environments, weapons systems, and simulations for the battlefield. Codes released within the last few years: predict the weather with forecasting and nowcasting; model radar-based sensing of surface and subsurface targets, including land mines, unexploded ordnance, and vehicles; model 3D rectangular arrangements, such as the pulsed plasma micro-thruster for microsatellite propulsion; model and simulate large-scale military communications and tactical signal intelligence platforms, weather forecasting model improvements; and simulate large-scale, heterogeneous, and communication networks.

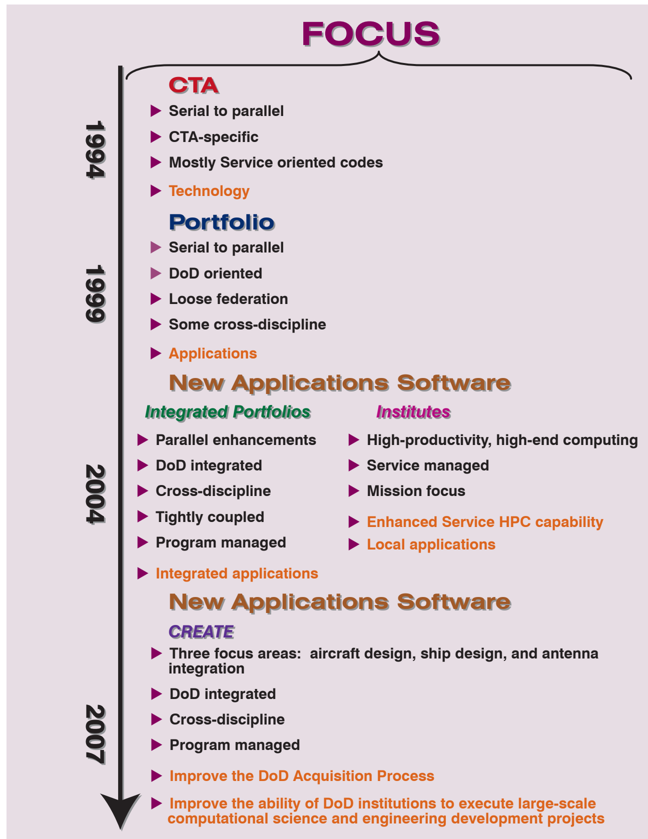

Goal 2—continued

Physics-based Environment for Urban Operations (PEUO)

The PEUO portfolio focuses on integrating an entity-level urban combat model with a command and control module and a chem-bio dispersion model. The integrated software will be used by the warfighter training community to provide realistic scenarios for urban combat operations. The command and control features furnish physics-based radio propagation

Goal 2—continued

models with non-stationary networks while the chem-bio dispersion models provide real-time, weather-dependent, concentration bands of toxic cloud movement. The portfolio has integrated CT-Analyst software with the One Semi-Automated Forces (OneSAF) Objective System (OOS). CT-Analyst is an efficient graphical user interface for instantaneous (50ms or less) plume rendition and situational analysis. OOS is a next-generation computer generated forces simulation. A demonstration of the integrated software has been provided for stakeholders from the US Army program Executive Office for Simulation, Training and Instrumentation who were quite impressed with the demonstration and scenario, which is running with a company moving through Baghdad and halting when they encounter a smoke cloud.

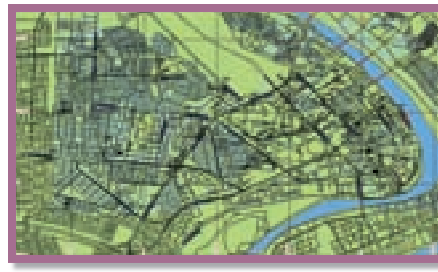

PEUO is scheduled for beta testing in the first quarter of FY 2008 and subsequently will be capable of being deployed.

Insensitive Munitions (IM)

In February 2007, the HPCMP and executive agents for IM within the Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics organized an IM computational science workshop. Participants at the workshop identified a set of IM modeling and simulation technical objectives which were within the scope, resources, and timeframe of the HPCMP. The final outcome of the workshop was an outline of an execution plan for a feasible IM portfolio. The portfolio will increase the capability, of the two Department

Portfolios

Collaborative Simulation and Testing (CST)

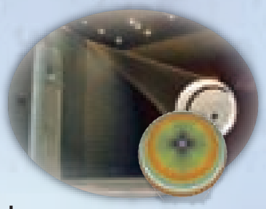

This portfolio provides scalable software for military applications focused on collaborative simulation and testing to reduce risk in weapon system development and to provide information to senior decision makers throughout the life cycle of the systems. The products of the CST portfolio will allow for simulations and resultant data prior to a test, capture test results, perform real-time data validation of test results, and provide for data interrogation and comparisons after the test.

Multi-Phase Flow Target Interaction (MFT)

This portfolio delivers integrated, state-of-the-art, physics-based codes for the design and performance prediction of munitions. The approach is a spiral development, incorporating and integrating current understanding of phenomenology into existing, validated DoD and DOE component codes resulting in an integrated toolkit. The industrial base will use this toolkit to efficiently conceive, design and qualify new systems, considering the full spectrum of energetics technologies and applications to the United States.

Physics-based Environment for Urban Operations (PEUO)

This portfolio focuses on integrating an entity-level, urban combat model with a command and control model and a chem-bio dispersion model. The integrated portfolio is used to provide realistic training scenarios for urban combat operations. The command and control features provide physics-based radio propagation models with non-stationary networks while the chem-bio dispersion models provide real-time, weather-dependent, concentration bands of toxic cloud movement. This portfolio will provide enhanced combat training as well as domestic disaster event, recovery and relief training.

Virtual Electromagnetic Design (VED)

This highly integrated portfolio provides the DoD the ability to design, from first-principle electromagnetics, in situ wide-band, multi-functional antennas and rough surface scattering solutions for a wide range of DoD activities including communication, acquisition, target identification, surveillance, and electronic attack.

Insensitive Munitions (IM)

Munitions must reliably fulfill their performance, readiness and operational requirements on demand and minimize the probability of inadvertent initiation and severity of subsequent collateral damage to weapon platforms, logistic systems and personnel when subjected to unplanned stimuli.

Figure 12. Portfolios

Goal 2—continued

of Energy developed multi-physics codes currently in use by a majority of DoD users, to model bullet and fragment impacts (BFI) against rocket motors and confined energetic warheads. The deliverable product of this effort will be an improved modeling toolkit, consisting of two enhanced multi-physics codes, with the ability to more accurately model bullet and fragment impacts against rocket motors and confined energetic warheads.

This toolkit is being developed with the intent of providing it to government and industrial munitions designers currently engaged in the field of designing and developing more IM compliant munitions. Its acceptance and use by government laboratories and industry will advance the state-of-the-art in the fidelity of the tools that US ordnance engineers use to develop bullet and fragment impact compliant munitions' systems. It will be the first step, for the DoD and its industrial base, to be able to model BFI against rocket motors and confined energetic warheads with a higher degree of fidelity than was ever previously possible.

For this portfolio, FY 2007's focus was on planning and organizing the effort.

Institutes

Institutes address Service/Agency high priority, high value technology or materiel RDT&E mission priorities and augment traditional processes with computational insight by using legacy or newly-developed computational techniques. Additional information about the six institutes is contained in Figure 13 and in the following paragraphs.

Institute for Maneuverability and Terrain Physics Simulation (IMTPS)

The mission of the IMTPS is to foster a culture within DoD of using high-fidelity simulation to attack problems hindering maneuverability. The IMTPS focuses on simulating near-surface environmental processes

Goal 2—continued

to support: 1) detection of landmines, improvised explosive devices, and unexploded ordnance; 2) use of seismic and acoustic unattended ground sensor networks; 3) analysis of maneuverability and trafficability; and 4) remote sensing of denied areas. The IMTPS vision integrates physics-based geotechnical, geophysical, hydrogeologic, and hydrologic analyses into a virtual-testing facility for resolving terrain-related warfighter problems.

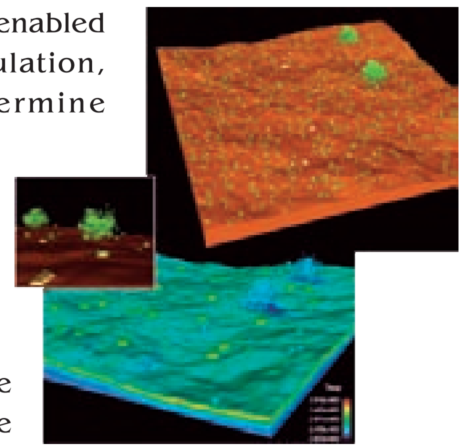

In FY 2007, the IMTPS enabled high-resolution simulation, using the countermine computational testbed, to explore the effects of weather conditions on the performance of infrared sensors in detecting landmines and other explosive devices. These simulations were driven by observed meteorological data including natural rainfall events, transient cloud cover, and observed wind speed. The goal was to provide operational guidance on the best and worst conditions in which to search for targets. These same computational tools were used to help understand the phenomena behind the detectability of disturbed soil and to guide field testing of sensors common in theater by predicting optimal 'hunting' times.

To accommodate minefield and other large-area detection from airborne platforms, our computational domains must increase from the current surface size of 100–400 m^2 to at least 10^4 m^2, while maintaining enough pixels on individual targets. Such simulations could quickly top 1 billion elements. Over the next two years, a domain-decomposition approach is being built to tile multiple, 5–10 million element simulations and produce seamless, synthetic, high-resolution

GOAL 2—CONTINUED

imagery comparable to that from an entire over-flight. This technology also has immediate use in synthetic environments for unmanned ground vehicles.

Biotechnology High-Performance Computing Software Applications Institute for Force Health Protection (BHSAI)

The BHSAI serves as an interdisciplinary, tri-Service resource to develop and apply HPC software that will accelerate research and development of militarily necessary medical products for DoD's Force Health Protection strategy. In force protection, the ability to rapidly differentiate between benign and pathogenic organisms is vital for battlefield treatment. The differences between a benign and pathogenic organism are impossible to distinguish by eye, necessitating chemical and biological diagnostic assays. The problem to be solved is how to distinguish between closely related species where one, Yersinia pestis causes plague, but another Yersinia pseudotuberculosis is less harmful and does not require the same intense medical intervention.

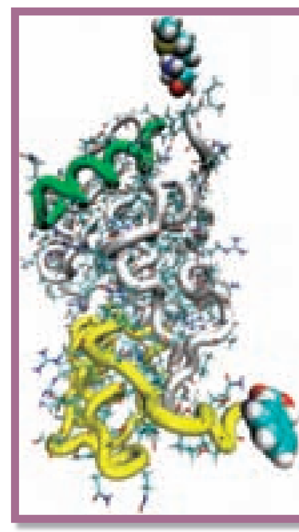

During FY 2007, the Institute enhanced the tools that where previously developed to examine sequence fragments and identify short, specific and unique sequences to an organism, termed "fingerprints". The software is now able to search for fingerprints common to multiple genomes in addition to searching single genomes, and is faster than the previous year's version. The US Army Medical Research Institute of Infectious Diseases (USAMRIID) has used the BHSAI's diagnostic assays to create printed custom arrays incorporating in-silico-determined fingerprints able to differentiate *Y. pestis* from *Y. pseudotuberculosis.*

GOAL 2—CONTINUED

The software is now entirely transparent to end users through the use of a graphical user interface. For the first time, life scientists at USAMRIID are now able to design diagnostic probes directly from their desktop computers, seamlessly utilizing the powerful HPC resources of MSRCs. The probes are used in diagnostic assays to identify pathogens that the warfighter may be exposed to in the field.

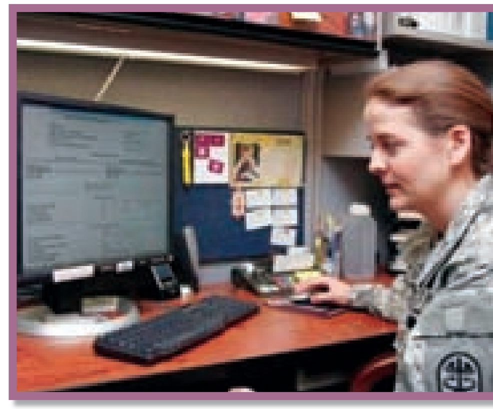

MAJ Jeanne Geyer of USAMRIID using software on her desktop to identify probes for pathogen diagnostic that effortlessly interfaces with the MSRCs HPC resources. The high-throughput software, developed by the Biotechnology HPC Software Applications Institute (BHSAI), Ft. Detrick, MD, identifies unique pathogen "fingerprints" (short DNA sequences) by comparing the pathogen's DNA sequence against all known sequences in large genome databases. The software is entirely run through a Web-browser, making it easier for life scientist without HPC skills to use the system.

Aiding in the design of pathogen probes, the software provides the flexibility of finding fingerprints that are common to multiple genomes as well as fingerprints that are unique to a given pathogen. Schematically, the sequences indicated in grey, green, and red are common to all three Burkholderia pathogen species. The software significantly reduces the time and number of experiments needed for the design of diagnostic assays.

Burkholderia suis	GTGGATGCAAGCTGGAGATA
Burkholderia abortus	GATATAATGGATAAATCGCAAAGCGTGG
Burkholderia melitensis	GGATCGATACGCAAGCACGGTGGGC

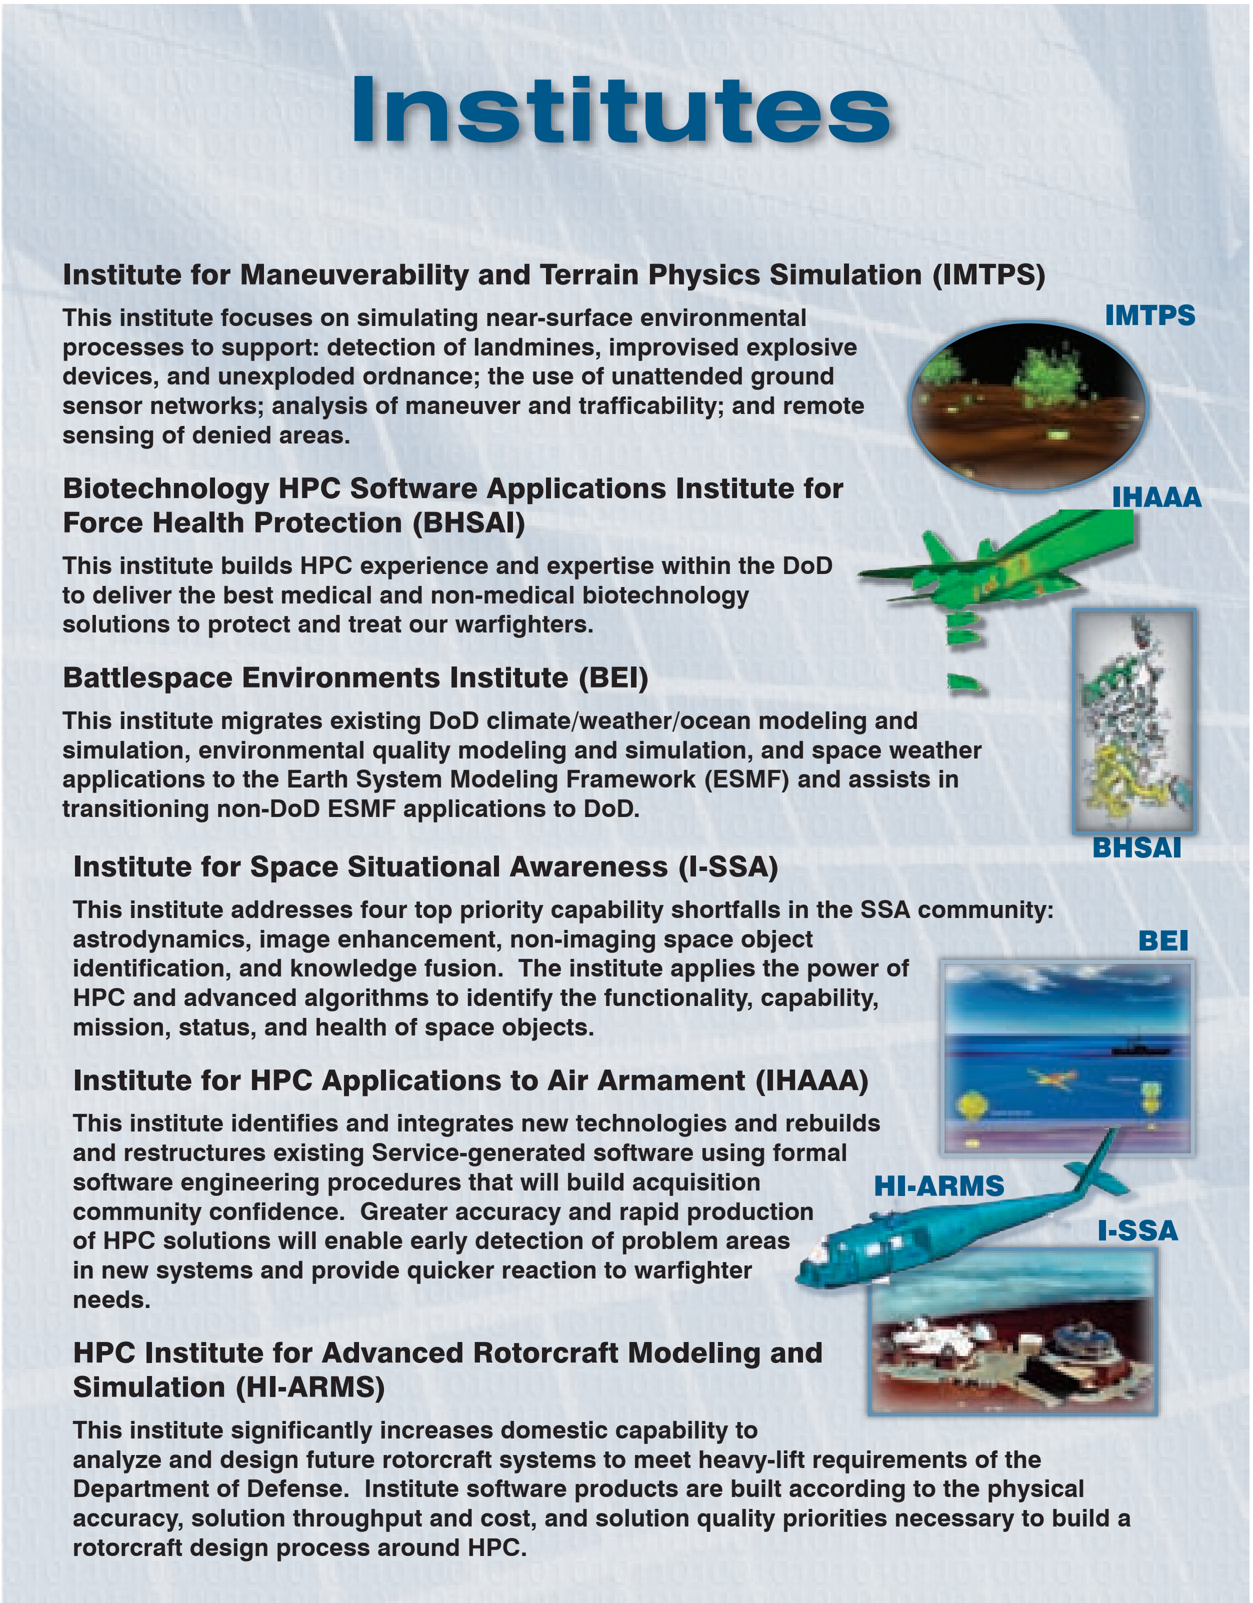

Figure 13. Institutes

GOAL 2—CONTINUED

Battlespace Environments Institute (BEI)

The BEI migrates existing DoD existing climate/weather/ocean modeling and simulation, environmental quality modeling and simulation, and space weather applications to the Earth System Modeling Framework (ESMF). During FY 2007, BEI completed a beta test of a two-way coupled ice-ocean prediction system that is planned for operational transition to the Naval Oceanographic Office in 2008. Testing in the Bering Sea demonstrated improved skill in predicting the location of the ice edge versus the uncoupled model. The figure below shows improvement as compared with the satellite derived ice edge. The ocean model HYCOM (Hybrid Coordinate Ocean Model), exports sea surface temperature, sea surface salinity, surface ocean currents and available freeze/melt heat fluxes to the ice model CICE (sea ice model), which in turns exports ice concentration, ice-ocean stresses, heat and salt fluxes and solar radiation at the ice base.

A two-way coupled atmosphere-ocean prediction system has also completed successful development and beta testing. The two-way coupled ocean-atmosphere prediction system will provide realistic feedback at the air-sea interface; leading to more accurate predictions for the warfighter in littoral and deep-water areas. It may also be used for improved hurricane prediction, ocean prediction, optimum track ship routing, search and rescue, anti-submarine warfare and tactical planning. This application will be transitioned as part of the Coupled Ocean Atmosphere

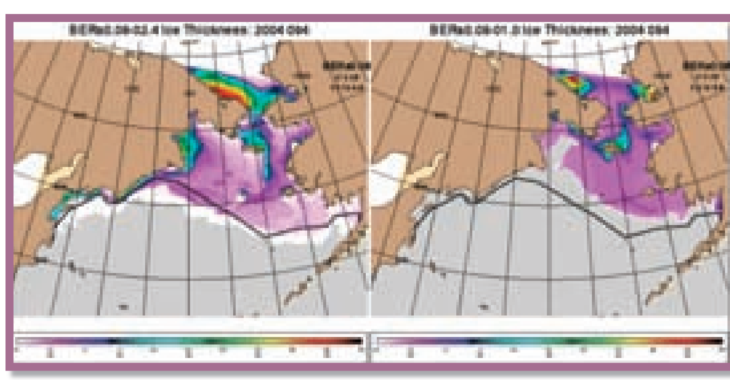

GOAL 2—CONTINUED

Mesoscale Prediction System—On Scene system at the Fleet Numerical Oceanography and Meteorology Command in Monterey, CA. Additional applications under development and testing include a coupled watershed nearshore application for riverine interests and a space weather application designed to improve the lead forecast time of coronal mass ejections which affect satellite communications.

Institute for HPC Applications to Air Armament (IHAAA)

The IHAAA integrates HPC with the need to field new weapons and new weapon configurations in a rapidly changing warfare environment. Through its various projects, the IHAAA delivers faster analysis capability that translates to faster development and certification of air armament systems. This directly provides improved war fighting capability for combatant commanders.

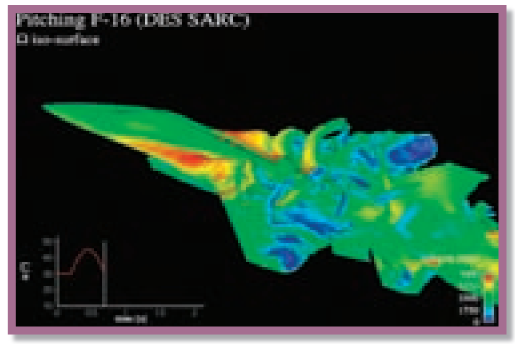

Accomplishments during FY 2007 include: F/A-18C clearance for GBU-12, GBU-38, MK-82, Mk-83, and Mk-84 with 300 gallon tank and LITENING pod configurations, clearance for the Higher Speed Anti-Radiation Demonstration on a QF-4, and integration analysis of Miniature Air Launch Decoy on a B-52. The total number of simulation runs performed exceeded 6,000, with an achieved cost avoidance of over $2.0M. These accomplishments shorten delivery time for new war fighting capability, reduce costs to deliver this capability, and help minimize limitations to system employment.

IHAAA is also preparing modeling and simulation capability for developmental weapons systems, such as the F-35 GBU-12 simulated weapon bay deployment employing moving control surfaces. Another significant

GOAL 2—CONTINUED

developing effort is the Aircraft and Armament Stability & Control project (AAS&C) (F-16 example pictured on page 65). The AAS&C computationally derives stability derivatives based on prescribed aircraft simulated maneuvers. The project has cultivated a pronounced interest from major aerospace industry contractors, such as Boeing, Lockheed, and Raytheon, and provides an opportunity to impact not only existing test requirements but to extend into the early phases of the acquisition design and development process as well.

Institute for Space Situational Awareness (I-SSA)

The Space Situational Awareness (SSA) mission area is very broad, spanning the space support and mission support foundation tiers of United States military space power and directly enabling all other Space Control operations. I-SSA supports the SSA needs of Space Control operational and acquisition forces by developing HPC versions of important existing applications and by developing new parallel scalable applications where there are gaps in warfighting capabilities. For example, this Institute is increasing combat capability by providing sharply enhanced high-resolution images of space objects that must be viewed through the turbulent, blurring atmosphere. The enhanced images are produced by the Institute-developed Physically Constrained Iterative Deconvolution (PCID) software suite. Dramatic decreases in execution times have been realized through software engineering and employment of HPC, thus making PCID practical for users.

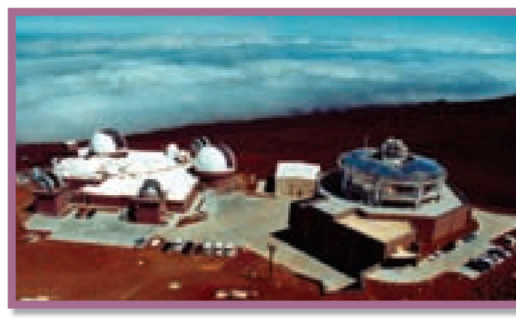

Goal 2—continued

During FY 2008, the Institute will make PCID available in a user-friendly interactive operational mode for further evaluation and actual use by SSA personnel at the Maui Space Surveillance System atop Mount Haleakala in Hawaii (pictured on page 66).

During FY 2007, this Institute has also provided enhanced capability to analyze and re-architect the global Space Surveillance Network (SSN) force structure, considering both optical and radar sensors (ground-based and space-based); satellite catalog completeness and accuracy; sensor tasking modes; and event detection coverage and timeliness by using HPC to reduce simulation times an order of magnitude with the Air Force's SSN Analysis Model. The FY 2007 improvements provided rapid SSN performance evaluations, both current and projected, to operational and acquisition managers, and, also enabled the assessment of the effect of individual sensor data quality on the overall SSN operation and products. The latter type of analysis, which was not possible until now, is crucial for making cost-effective system engineering tradeoffs in high-value acquisition programs for space surveillance sensors.

In other activities during FY 2008, the Institute will demonstrate a new scalable multiple-hypothesis tracking software application for associating widely separated observations of thousands of space objects in near-real time. This demonstration will be used and evaluated by key joint-service SSA personnel in an operational environment. Additionally, the Institute will begin designing, building, and hosting a comprehensive SSA data repository, high fidelity simulation suite, and exploitation testbed to support DoD-wide research, development, and testing of advanced SSA data fusion techniques.

HPC Institute for Advanced Rotorcraft Modeling and Simulation (HI-ARMS)

The mission of HI-ARMS is to transform the analysis-test paradigm that currently exists within the rotorcraft

Goal 2—continued

industry and government laboratories in the United States into one built around HPC, which will provide domestic manufacturers the means to create effective designs (or upgrades) of rotorcraft systems required by DoD and to minimize development cost and risk. Simultaneously, the DoD will have the means to accurately predict mission capability, to improve the effectiveness of vehicle test programs, and to effectively conduct rotorcraft source selection processes, including analyses required to support airworthiness qualification.

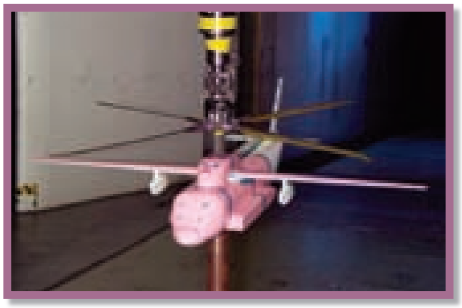

a)

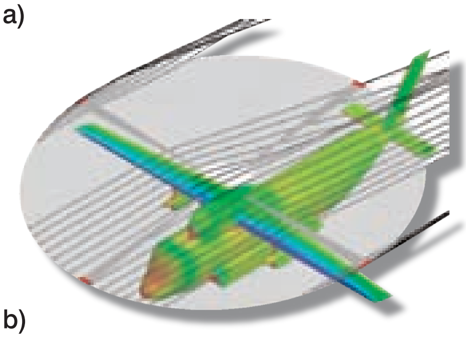

b)

Currently, HI-ARMS's work is supporting investigation of the rotor-fuselage interactions of the slowed-rotor compound, a candidate for the future Joint Heavy Lift (JHL) rotorcraft (Figure a). Improved prediction and understanding of rotor-fuselage interactions can reduce the possibility of aircraft problems and lead to mitigation through aerodynamic design. This JHL application highlights the importance of getting HPC tools into the design process as early as possible, in order to avoid costly problems in future flight tests, especially for untested, novel configurations. For example, computational simulations of a slowed-rotor compound using an actuator disk rotor model (Figure b) show potentially adverse empennage interactions with the rotor wake. These HI-ARMS computations compare favorably with available wind tunnel test measurements of surface and flowfield data.

Asynchronous Parallel Optimization of Relativistic Magnetron Cathode for Counter Electronics

Peter J. Mardahl, Matthew T. Bettencourt, and Keith L. Cartwright, Air Force Research Laboratory, Directed Energy Directorate, Kirtland AFB, NM

Producing the most suitable high power microwave (HPM) source for weapons applications requires understanding many complicated issues. Suitability includes many physical and performance characteristics, such as size, weight, form factor, output power and frequency, energy and power efficiency, operator safety, reliability, and maintainability.

To make the problem tractable and less human labor intensive, we have created an automated HPM weaponry source design method using an integration of ICEPIC[1] (Improved Concurrent Electromagnetic Particle-In-Cell)used for simulating HPM sources) and APPSPack[2] (Asynchronous Parallel Pattern Search) used for nonlinear optimization.

With the new design methodology, humans still decide which parameters to vary and reasonable ranges for the parameters. However, the process of actually finding the optimum operation within the parameter space is now automated.

Designs resulting from the computer-aided process are superior to human-only designs, because the computer can automate the process of making effective design tradeoffs in multiple dimensions, a difficult proposition for an unaided human. The end result is a more effective HPM weapon system designed at less expense and in less time.

For this effort, we applied the optimizer to four parameters which described the shape of the magnetron cathode, a component from which electrons are emitted into the magnetron. In particular, we considered the benefits which could be obtained by inserting a shaped cathode in place of the standard circular-cross-section cathode: The figure shows a shaped cathode in the A6-3 magnetron.

The expected improvements were 1) increased mode purity, 2) increased power and energy efficiency, 3) faster device start-up, 4) increased power output, and 5) better reliability of operation when non-ideal voltage and magnetic field were applied. We did not expect changes in device size or frequency of operation from changing the cathode; in fact, none resulted.

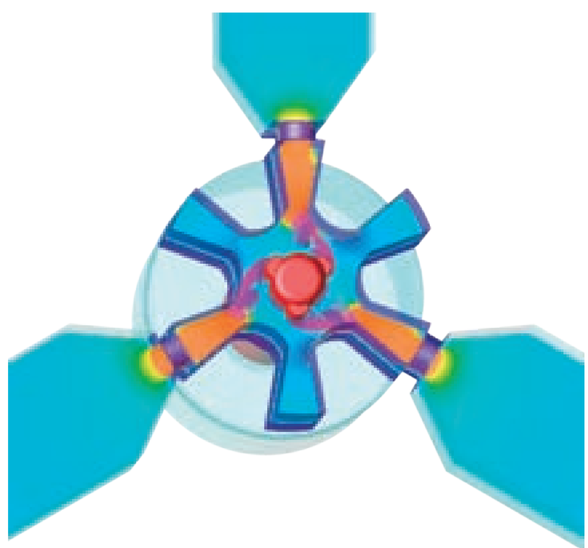

Shaped cathode in the A6-3 relativistic magnetron

The cathode shape parameters varied were:

- Cathode core size
- Cathode bump group 1 size
- Cathode bump group 2 size
- Cathode clocking

These parameters form a 4-dimensional parameter space which had to be explored. A straightforward exploration with 8 points in each dimension would have required 4,096 designs be evaluated to produce an optimum. However, the optimization process only required 84 designs to be evaluated before discovering an optimum.

During design, a maximum of 3,456 CPUs of Jaws at MHPCC (a Dell PowerEdge system), were used to evaluate six separate designs simultaneously, with each design evaluation requiring nine runs of 64 CPUs each. For this number of CPUs and for this ICEPIC input, parallel efficiency was approximately 80%.

There were two especially significant results of this effort. One was the application of this new design technique to the design of an optimal cathode, which led to a great improvement in the simulated performance of the magnetron. The other was the creation and successful testing of an automated technique for converting design ideas into optimized HPM designs, i.e., the successful integration of APPSPack and ICEPIC.

The new optimal cathode improved the power output 1.4x, averaged over a magnetic field and voltage scan, with each point in the scan weighted equally. Peak power and energy efficiency were both increased 1.5x for a similar average. Likewise, impedance remained within acceptable ranges, while mode purity increased 1.2x.

Of more significance than the actual improvements to the magnetron design is the new design method. A similar cathode optimization for a previous magnetron design (the A6-3) took a year or more to develop—not a man-year, the problem was not worked full-time, but a calendar year. The new design method produced an optimal cathode in less than six calendar days, at a cost of 300,000 CPU-hours. Further, this design brought improvements in several objectives: not just peak power, but energy and power efficiency, mode purity, and impedance; only peak power was considered previously in human-driven optimization.

The new technique speeds the development of HPM weaponry, reduces development cost, and yields superior weapons. The 1.4x power output increase quoted above will increase effective weapon range. The 1.5x efficiency increase will reduce the size and weight of the pulsed power system required to drive the HPM weapon, and will reduce the amount of fuel used. The decrease in HPM source size will in turn reduce the size of the entire HPM system, allowing more possible weapon transport modes, giving the warfighter more battlefield options.

References

1. Peterkin, Jr., Robert E. and John W. Luginsland, "Theory and design of a virtual prototyping environment for directed energy concepts." *Computing in Science and Engineering*, 4, no. 2, p. 42, March–April 2002.

2. Kolda, Tamara G., "Revisiting Asynchronous Parallel Pattern Search for Nonlinear Optimization". *SIAM J. Optimization*, 16(2), pp. 563–586, Dec 2005, (doi:10.1137/040603589) (http://epubs.siam.org/sam-bin/dbq/article/60358).

GOAL 2—CONTINUED

Institutes to be Awarded in FY 2008

On October 19, 2007, the Deputy Under Secretary of Defense for Science and Technology (DUSD(S&T)) solicited recommentions for additional Institutes. This call was targeted to the following focus areas:

- modeling human behavior (e.g., cultural, ethnic, religious, societal, etc.) at individual, group and population levels for political, military, economic, social, infrastructure, and information course of action analysis and forecasting or forces modeling;
- modeling dynamically reconfigurable networks in adverse environments;
- shock and vibration survivability of large, multi-component ship structures;
- reactive chemistry of explosives and propellants for use by the insensitive munitions modeling community;
- high power microwave systems modeling for counter improvised explosive devices, non-lethal, and information operations/electronic warfare applications; and,
- dynamic, time-accurate, thermal management system modeling, allowing insertion of models of contractor components.

The services and agencies submitted seven proposals. A highly qualified team of subject matter experts from within the Department of Defense evaluated the proposals. Based on Department priorities, these evaluations and funding contraints, DUSD(S&T) approved the following awards to begin in FY 2008:

- Institute for Dynamically Reconfigurable Networks in Adverse Environments, Army Research Laboratory, Aberdeen Proving Ground, MD;
- Institute for High Power Microwave Employment, Integration, Optimization, and Effects, Air Force Research Laboratory, Kirtland AFB, NM; and,

Goal 2—continued

- Institute for Multi-Scale Reactive Modeling and Simulation of Insensitive Munitions, Army Research Laboratory, Aberdeen Proving Ground, MD.

User Productivity Enhancement and Technology Transfer (PET)

PET enables the DoD HPC user community to make the best use of the computing capacity the program provides and to extend the range of DoD technical problems solved on HPC systems. PET is enhancing the total capability and productivity of the program's user community through HPC-related science and technology support, training, collaboration, tool development, support for software development, technology tracking, technology transfer, and outreach to users.

PET is responsible for gathering and deploying the best ideas, algorithms, and software tools emerging from the national HPC infrastructure into the DoD user community. The PET activities are conducted through two separate contracts; one to MOS University Consortium, led by Mississippi State University, and the second to High Performance Technologies, Incorporated. The teams from both contracts involve academic leaders to serve as points of contact for each of the areas covered by PET and experienced Ph.D.-level personnel located at DoD sites to provide HPC and one-to-one scientific assistance to HPCMP users. The teams are comprised of experts from a broad range of universities and companies highly regarded in the HPC field (see Figure 14). In addition, PET personnel lead short-term projects that focus on delivering capabilities for specific needs.

PET supports all ten HPCMP computational technology areas, and the following four crosscutting areas, with a broad HPCMP-wide management approach.

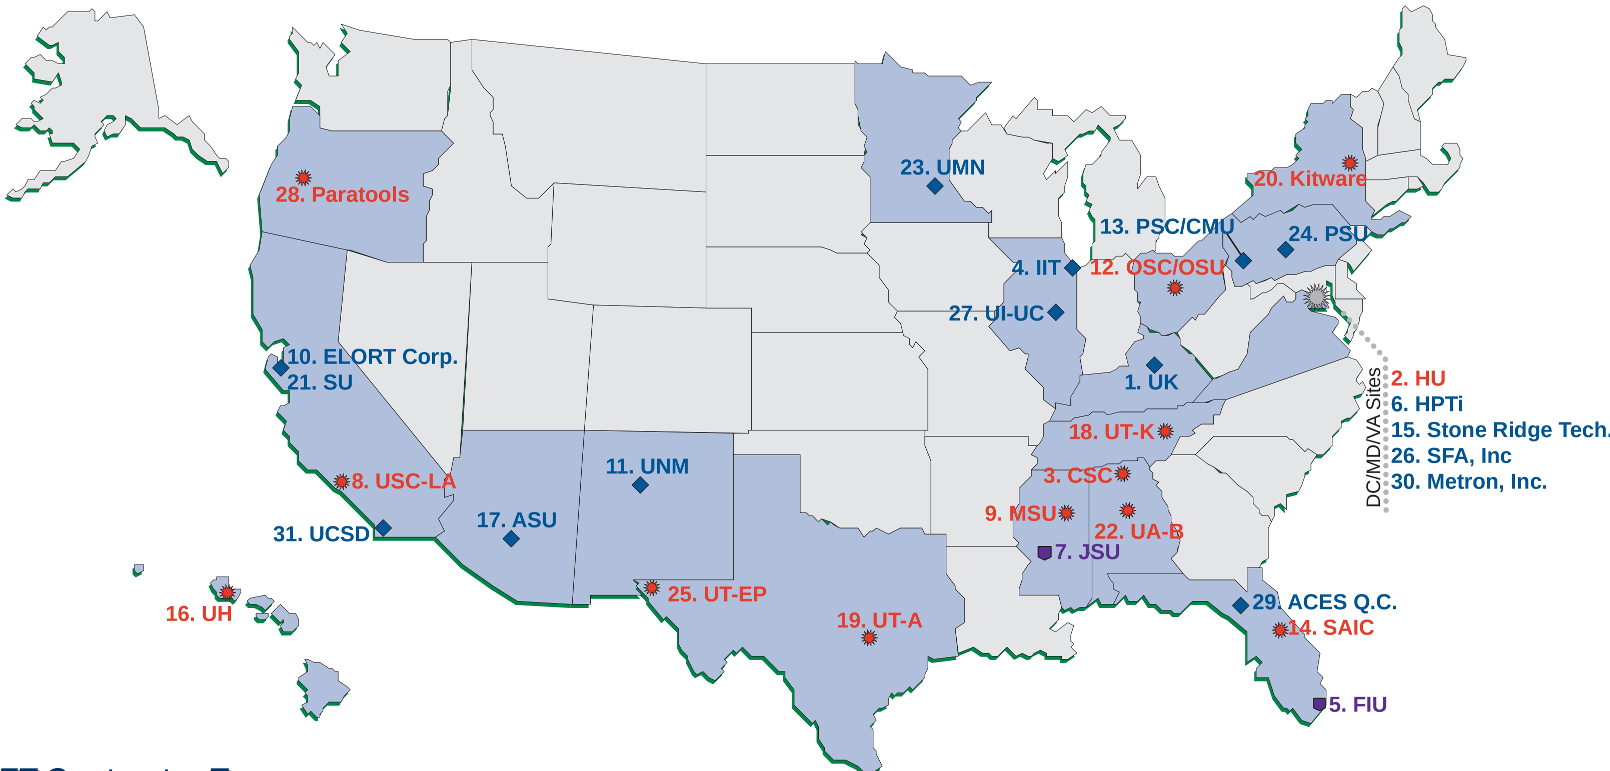

PET Contractor Teams

1. University of Kentucky, Lexington, KY
2. Howard University*, Washington, DC
3. Computer Sciences Corporation, Huntsville, AL
4. Illinois Institute of Technology, Chicago, IL
5. Florida International University*, Miami, FL
6. High Performance Technologies, Inc., Reston, VA
7. Jackson State University*, Jackson, MS
8. University of Southern California, Los Angeles, CA
9. Mississippi State University, Starkville, MS
10. ELORT Corp., Sunnyvale, CA
11. University of New Mexico, Albuquerque, NM
12. Ohio Supercomputer Center/The Ohio State University, Columbus, OH
13. Pittsburgh Supercomputing Center/Carnegie-Mellon University, Pittsburgh, PA
14. Science Applications International Corporation, Orlando, FL
15. Stone Ridge Technology, Bel Air, MD
16. University of Hawaii*, Honolulu, HI
17. Arizona State University, Tempe, AZ
18. University of Tennessee, Knoxville, TN
19. University of Texas, Austin, TX
20. Kitware, Clifton Park, NY
21. Stanford University, Palo Alto, CA
22. University of Alabama, Birmingham, AL
23. University of Minnesota, Minneapolis, MN
24. Pensylvania State University, University Park, PA
25. University of Texas*, El Paso, TX
26. SFA, Inc., Crofton, MD
27. University of Illinois, Urbana-Champaign, IL
28. Paratools, Eugene, OR
29. ACES Q.C., Gainesville, FL
30. Metron, Inc., Reston, VA
31. University of California at San Diego

** denotes Minority Serving Institution (MSI)*

As of: February 2008

Legend

- HPTi Team
- MOS University Consortium
- HPTi and MOS

Figure 14. PET contractor teams

GOAL 2—CONTINUED

Enabling Technologies (ET)

The ET functional area provides tools, algorithms, and standards for pre- and post-processing large datasets. Such processing includes the following technologies: mesh generation, visualization (both local and remote), data mining and knowledge discovery, image analysis, and problem solving environments.

Computational Environment (CE)

Improving the usability of the computational environments at the HPCMP centers is critical for easily and effectively using the program's resources. CE includes all aspects of the user's interface to high performance computing resources, such as programming environments (debuggers, libraries, solvers, higher order languages; performance analysis, prediction, and optimization tools), computing platforms (common queuing, clusters, distributed data, and metacomputing), parallel algorithms, user access tools (portals and web-based access to high performance computing resources), and consistency across the centers for locating these capabilities.

Collaborative and Distance Learning Technologies (CDLT)

This functional area focuses on supporting HPCMP users who are unable to attend HPC-based events, such as training classes and meetings. CDLT is responsible for webcasting and video-capturing events and post-processing the material to create high quality instructional content. After approval, such content is available for downloading from the PET Online Knowledge Center. CDLT also provides support for video teleconferencing services. Strong interactions with the DREN component and with Centers' staffs ensure that CDLT activities are coordinated and incorporated into the program's networking and security infrastructure.

Goal 2—continued

User Training Coordination (UTC)

UTC, formerly known as Education, Outreach, and Training Coordination (EOTC), coordinates formal and informal knowledge delivery to the DoD HPCMP current and potential user community. The change of name and mission reflects the creation of a new and separate HPCMP area to focus on HPC-related activities and outreach to minority students and minority serving institutions. This new focus area, Joint Educational Opportunities for Minorities (JEOM), described on page 84, also includes former EOTC elements such as student summer internships, summer visiting faculty, and summer institutes.

UTC functions include: coordination of on-site training at the Program's shared resource centers and remote sites; selecting optimal training delivery methods and media; and coordinating outreach forums, such as conferences, workshops, seminars, and symposia. The training program results from information gathered from multiple sources, including the HPCMP User Requirements Database, discussions with users, and formal requests from users via the PET Online Knowledge Center (OKC). Together with CDLT, UTC conducts a broad-based DoD HPCMP user training program at HPCMP Centers and at other user sites, as well as via webcasting and online (via video-captured events on the PET OKC). Training content targets both commercial and government software, software profiling tools, traditional and high level software languages, new hardware features, such as Graphical Processing Units, and more.

PET Highlights

HPCMP technical and program management has emphasized and encouraged our entire team of functional experts, on-site personnel, principal investigators, and business administrators to focus on the key goals of PET program: technology transfer, user productivity, and DoD mission impact. The following examples show such achievements.

Radio Frequency Radiation Effects on the Warfighter

The understanding of the biophysical and biological impact of electromagnetic (EM) fields on humans is an important field for DoD researchers. Many DoD personnel work in close proximity to intense and possible pervasive EM fields generated by equipment, such as field radios and cell phones. The Frequency Radiation Branch at Air Force Research Laboratory in San Antonio, TX has developed a virtual simulation of the human based upon data from the Visible Man project. They conduct simulations of the effect of various EM fields on the human body with this technology, however, the computational demands can be significant, reducing the number of iterations and hypotheses they can test.

The PET SIP team helped deliver a week-long course at the Air Force Research Laboratory, Human Effectiveness Directorate, Directed Energy Bioeffects Division, Radiofrequency Radiation Branch (AFRL/RHDR) on MATLAB. The team provided assistance with optimizing the existing MATLAB-based electromagnetics simulation code. The team was able to vectorize significant portions of the code, which resulted in a speedup of 30X.

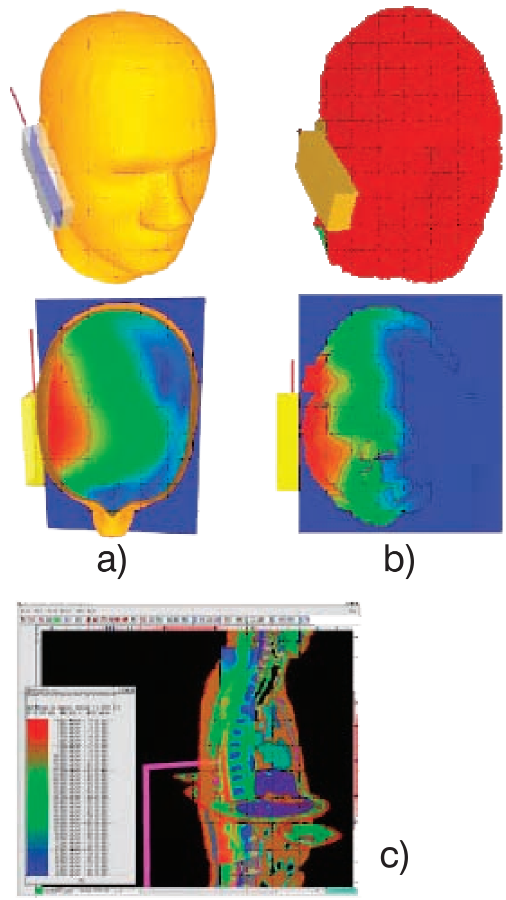

a) SAR at SAM phantom; b) SAR at human head; and c) ACD inside a human body exposed to the feld from a Loop antenna at 50 khz.

Jason Payne (AFRL/RHDR) provided the following comment: "I would like to take this opportunity to thank you for deriving a vectorization process for our two-dimensional Finite Difference Time Domain (FDTD) code. Our research group at AFRL relies heavily on computational methods such as FDTD in order to predict energy absorption in biological tissues due to RF energy. However, these methods can be very computationally expensive, and simulation run times may vary anywhere from hours to days or longer. . . I recently brought this to your attention during a MATLAB training class, and you were able to very quickly turnaround a process for vectorizing and further optimizing our 2D FDTD simulation code. Initial runs indicate that implementing this process may decrease our simulation run time by up to a factor of 30. This type of performance enhancement will greatly increase the quality and efficiency of work that our modeling team can output."

The PET SIP team decreased simulation run times and thus increased quality and efficiency of the modeling of EM fields. The end result will be better designed communications devices for the Warfighter that have been shown to have minimal chance of impacting their health due to electromagnetic field.

Urban Acoustic Array Processing

Problem

Localizing and detecting acoustic signals in a realistic urban environment is crucial for decision-support and is technically challenging due to the complex propagation environment. A loud acoustic signal may be easy to detect in an urban environment, but localizing that signal is much more difficult. In particular, the desired signal and its location are both unknown ahead of time. Array processing works by coherent direction-based summation of signals across the array aperture to infer signal location. To achieve localization, signals must be related in a way that separates "desired" signals from "noise."

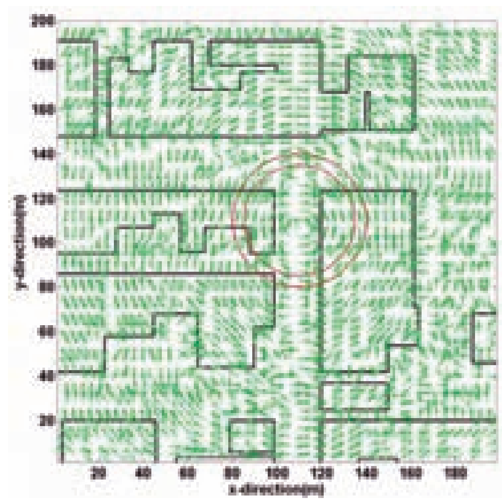

Methodology

Large-scale computing methods were combined with advanced signal processing techniques to detect, localize, and identify geo-acoustic events in a realistic urban environment. For this effort a distributed, two-sensor, time-delay array processing algorithm was used to relate acoustic signals and localized sources in a realistic urban environment. Using the output of the urban acoustic model, basic time-delay array processing was applied to a large set of dispersed two-sensor arrays where each array performs its own mutual correlation to determine source angle of arrival.

DoD Impact

This work is part of the effort of the IMTPS. The purpose of this HPCMP effort is to develop and transition C2 decision aids/tools and information management services for Current and Future Force Battle Command systems throughout all phases of operations and environments.

Harley Cudney (ERDC) recently provided the following quote: "When I described our work in acoustic propagation simulations in an urban environment, Dr. Jose Unpingco (SIP On-site - SSC-SD), enthusiastically and helpfully offered his thoughts on what we might do for the signal processing. Thoughts are one thing, but Jose has actually shown us, using some of the data we sent him, how we might approach the signal processing of the acoustic propagation domain … This would not have happened without Jose's capabilities, enthusiasm, and interest in helping. Dr. Unpingco is a resource we are glad to have found."

Rapid Generation of Unstructured Grids

Problem

Production tools for generating high quality unstructured Navier-Stokes CFD grids have not reached full maturity. Some generators handle parts of the problem well, but no single code can be effectively used to perform the entire process in a production environment.

Methodology

The commercial grid generator GridGen has mature capabilities for interacting with CAD and generating high quality unstructured surface grids. There are a large number of GridGen users in the HPC CFD user base. However, GridGen does not produce high quality unstructured Navier-Stokes volume grids. The unstructured grid generation software SolidMesh developed by Mississippi State University does generate high quality unstructured Navier-Stokes volume grids, but does not handle "bad CAD" surfaces very well. A process was developed to use GridGen surface meshes as a basis for generating volume grids with the Advancing-Front/Local-Reconnection (AFLR) grid generation routine in SolidMesh. Utility codes were developed to handle file format conversions.

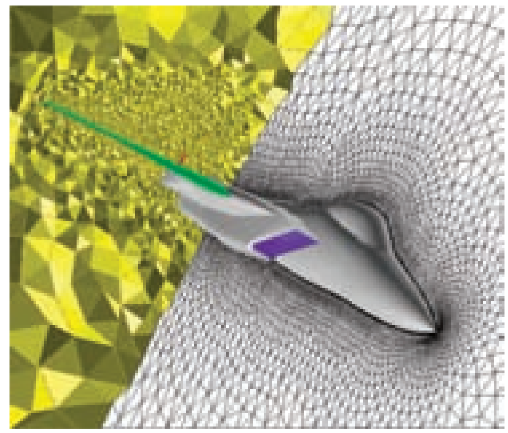

Volume grid for F-35 geometry

Users Supported

Greg Power and Jim Masters (AEDC); Darren Grove, Joe Laiosa, James Chung, and Jay Lee (NAWCAD); Jim Forsythe (COBALT Solutions); Joe Gorski and Larry Mulvihill (NSWCCD); Rick Smith (NSWCPC); Scott Morton and Bruce Jolly (AFSEO); and Balu Sekar (AFRL).

Grove commented: "I would like to extend my sincere appreciation and gratitude to all those who have assisted me in completing a recently successful high priority Joint Strike Fighter Computational Fluid Dynamics (CFD) store separation analyses project for the Navy… . Thanks to … Dr. Bobby Nichols for his file format converter that allowed exceptional viscous computational meshes to be generated around very complex geometries in a production type environment. Without these two tools, the project would not have been possible in the allotted time."

DoD Impact

This process has been used to support a number of weapon development programs including F-35, vortex combustor design, and SSN-21 submarine design. Received an accommodation letter from NAWCAD for support to F-35.

GOAL 3: ACQUIRE, DEPLOY, OPERATE AND MAINTAIN A COMMUNICATIONS NETWORK THAT ENABLES EFFECTIVE ACCESS TO SUPERCOMPUTERS AND TO DISTRIBUTED S&T/T&E COMPUTING ENVIRONMENTS

Defense Research & Engineering Network (DREN)

DREN was created to link high performance computing users and supercomputers, no matter where the person or resource is or with what Military Service they are associated (see Figure 15). From the beginning, DREN has acted as an enabler for the research, development, test and evaluation communities, the Missile Defense Agency (MDA), DoD Modeling & Simulation Office, Joint Forces Command, Defense Threat Reduction Agency, and others.

DREN enables major shared resource centers to perform secure, large-scale, remote, mass-storage for HPC disaster recovery. Although it has always been highly desirable to do in-band (live on-line) mass storage transfers, it was in the "too hard to do" category. The challenge of transferring terabits of data daily between multiple centers was out of reach. Recently, a number of advances have made these types of data exchanges a reality. Access to the DREN backbone was expanded at each of the DoD major shared resource centers to optical carrier (OC) or OC-48 (approximately 2.4 gigabits per second). These centers are the first within the DoD to have massive wide-area network (WAN) access capabilities. Anticipating rapidly rising bandwidth demands, DREN revamped its backbone nationwide using new protocol architectures (multi-protocol label switching and internet protocol security tunnels) using 9,000 byte maximum transmission units (jumbo frame enabled) which, in turn, permits high-

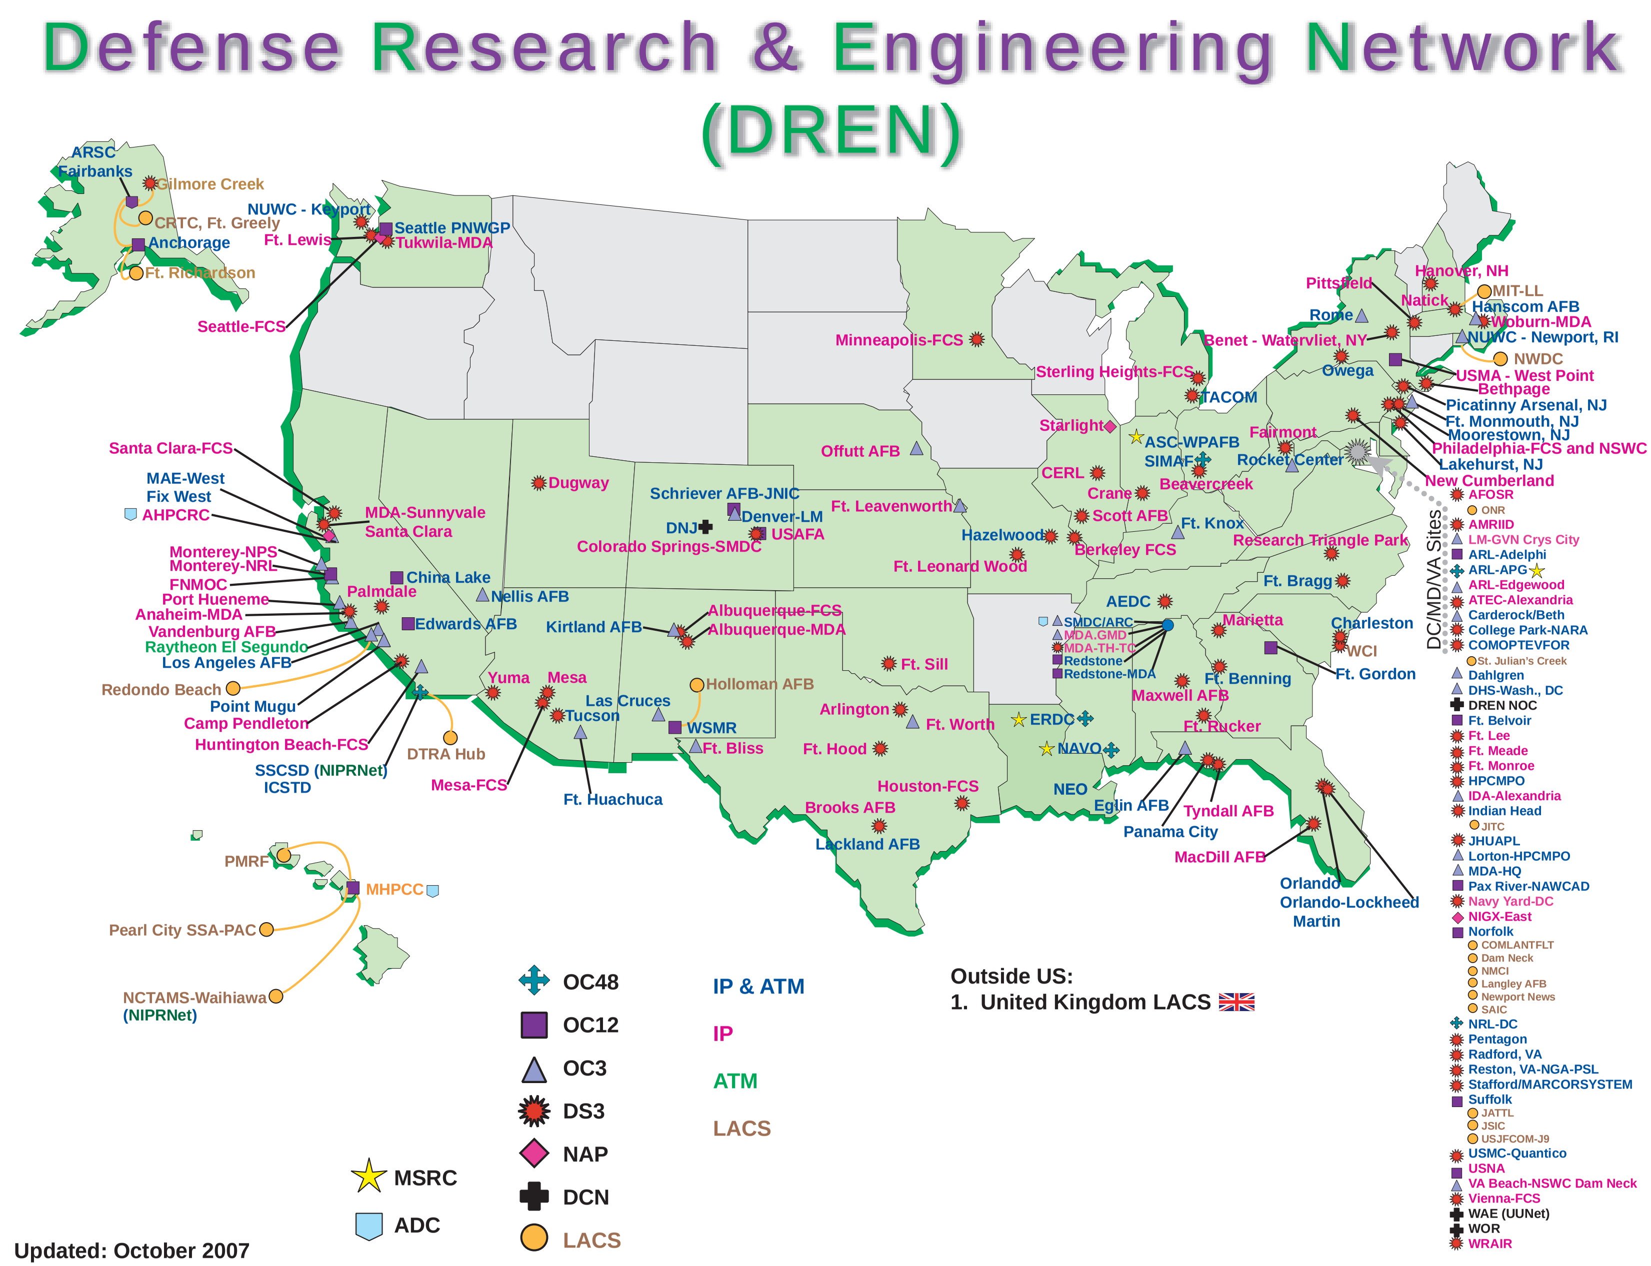

Figure 15. DREN connections, HPC centers and other network access points

GOAL 3—CONTINUED

end tuning of computational resources over thousands of miles for massive data transfers.

DREN is centrally funded for science and engineering users of DoD high performance computational resources. Other congressionally authorized groups (MDA, modeling & simulation, operational test and evaluation groups) not part of HPC line-item funding must offset service delivery point and security costs to access the DREN.

It is in the best interest of the DoD to expand continuously the pool of quality scientists and engineers working on high priority DoD problems. Potential new users often discover the availability of HPC resources through initial exposure to DREN. JFCOM in Suffolk, VA followed this pattern and eventually expanded into a joint, distributed, system-of-systems virtual communications concept for future real-time communications and network simulations. An advantage of DREN is that it makes high capacity bandwidth available to all computational resources wherever they may be. This approach makes it much easier to ensure optimal use of high performance computing assets and reduces the effective cost of these scarce resources.

Historically, we associated access to scarce and expensive resources with close proximity to major centers of civilization. Today, we have much more flexibility in the placement of new computational resources. That flexibility allows growth of new skill and job opportunities to rural (Gulf Coast, Midwest, and Southwest) or remote (Alaska and Hawaii) labor markets that otherwise would be overlooked. High bandwidth WAN access allows the HPCMP to get resources very close to specialized real-time systems while expanding the pool of potential users working on DoD problems and keeps those resources extremely busy.

The DREN will continue to expand as new computer technology, mass storage, and distributed computing environments place growing demands on bandwidth and network architecture. Core architecture will grow

Goal 3—continued

to 40 gigabits per second or more over the next few years and optics architecture will play an increasingly central role in the DREN core. Federal Agency support in both networking and security will continue to expand. We expect our role as a primary computer network defense security provider will enable more secure communications both within DoD and between other Federal Agencies in leading edge network environments.

Goal 4: Promote collaborative relationships among the DoD computational science community, the national computational science community and minority serving institutions (MSIs)

Defense Research & Engineering Network (DREN)

As one of the three major areas of DoD's High Performance Computing Modernization Program, DREN draws from the high performance computing community most familiar with Defense supercomputing for technical and security advisory group members. DREN personnel also participate in the more generalized DoD networking and security communities within the Global Information Grid, through direct participation on DoD control boards and technical advisory councils, and by participating as a Tier 2 DoD Computer Emergency Response Team for hostile acts of intrusion and compromise.

DREN contributes to overall federal agency networking and security through the Large Scale Network (LSN) and Joint Engineering Team (JET). These groups maintain and extend US technological leadership in leading-edge network technologies and coordinate federal agency networking activities, operations, and plans represented by DoD DREN, Department of Energy, National Aeronautics and Space Administration, the National Science Foundation, Next Generation Internet, and Internet 2. The JET and LSN are part of the White House's Office of Science and Technology Policy Interagency Working Group.

DREN peers (exchanges network traffic) at international exchange points including Starlight in Chicago, and the Pacific Northwest Gigapop in Seattle, and at advanced exchanges including Next Generation Internet Exchanges East and West in Maryland and California. DREN actively participates in international science exchanges such as the Australian Meteorological and Oceanographic Society and Asian Pacific Advanced Networks projects.

User Group Conference

The annual Users Group Conference (UGC) brings together the DoD high performance computing community to advance research techniques, report performance improvements, and accelerate software development efforts to improve the capabilities and effectiveness of the warfighter.

The HPCMP hosted the 2007 UGC in Pittsburgh, Pennsylvania on 18–21 June 2007. The conference featured five parallel tracks with more than 120 technical sessions and a poster session featuring 20 posters covering the extensive computational work being pursued in the DoD using HPCMP resources. The conference was also an opportunity for communities of interest to meet in Birds of a Feather meetings throughout the conference.

Attendees and interns at the UGC 2007 Users Group Conference

Goal 4—continued

Keynote speakers included:

- Brigadier General Nick Justice, Deputy Program Executive Officer, Army Command, Control, Communications-Tactical;
- Dr. André van Tilborg, DUSD(S&T); and
- Dr. Kathleen Carley, Professor of Computer Science, Institute for Software Research International, Carnegie Mellon University.

More than 360 people from 39 states attended the 2007 UGC.

Education and Outreach Program

The HPCMP Education and Outreach program offers opportunities to work with seasoned and well-respected scientists and engineers located at Department of Defense laboratories and test centers across the country. The Education and Outreach program consists of four components:

- Joint Educational Opportunities for Minorities (JEOM) summer internship,
- summer colloquium and summer institutes,
- workshops to train science, technology, engineering, and mathematics (STEM) faculty at Historically Black Colleges and Universities (HBCUs) and Minority Serving Institutions (MSIs), and
- travel funds and scholarships for STEM undergraduates.

In 2007, HPCMP established the JEOM internship to ensure a competitive and diverse workforce within the DoD by offering training, mentorship, and career development opportunities for graduate and undergraduate students in the STEM disciplines. JEOM interns work for 10 weeks during the summer and gain experience in key areas of defense-related research, such as biotechnology, maneuverability and terrain physics simulation, battlespace environments simulation, molecular modeling, materiel impact estimation, weapons design and testing, computer science, information technology, computational biology

Goal 4—continued

and chemistry, structural mechanics, fluid dynamics, and scientific visualization. In FY 2007, the Education and Outreach Program funded 21 JEOM interns from 13 universities at six locations:

Group of student interns that participated in the 2007 HPCMP JEOM Program

- High Performance Computing Modernization Program Office – Arlington, VA
- Naval Air Warfare Center, Aircraft Division – Patuxent River, MD
- Air Force Research Laboratory – Rome, NY
- Air Force Research Laboratory – Wright-Patterson Air Force Base, OH
- Army Research Laboratory – Aberdeen and Adelphi, MD
- Maui High Performance Computing Center – Kihei, HI

HPCMP funded a faculty workshop at North Carolina Agricultural and Technical University (NCA&T) to train STEM faculty from HBCUs and MSIs and to develop a collaborative environment for supporting computational STEM disciplines at MSIs. Attendees became candidates for visiting faculty appointments at DoD laboratories and major universities working on DoD-relevant projects. In July 2007, HPCMP sponsored a summer institute at the Computational Center for Molecular Structures and Interactions at Jackson State University, at which JEOM interns and students from Puerto Rico and several Historically Black Colleges and Universities presented the results of their research. The day-long event concluded with awards for outstanding research presentations.

The Education and Outreach Program committed funding to support eight additional Supercomputing 07 (SC07) broader engagement (BE) participation grants.

GOAL 4—CONTINUED

As a result, the SC07 BE committee accepted all applicants that the Review Team identified as excellent candidates. This commitment from the JEOM program made a difference for those participants who otherwise might not have been able to participate at SC07.

For JEOM FY 2008, HPCMP's recruitment efforts began by participating at the Grace Hopper Celebration of Women in Computing Conference, the Richard Tapia Celebration of Diversity in Computing Conference, SC07, and the Black Engineer of the Year Conference. Several talks were presented during the SC07 Education Program and at the BE booth at SC07. Representatives from DoD organizations participating in the JEOM program attended the talks and interviewed many potential JEOM candidates.

As a result, the JEOM internship program received five times the number of applicants as in 2007. JEOM will fund 50 interns at 10 DoD sites during the summer of 2008 in addition to two summer institutes at NCA&T and the University of Southern California. HPCMP will host the summer colloquium prior to the 2008 UGC in Seattle, WA covering four topic areas:

- career opportunities within the DoD science and technology community;
- professional development as scientists, engineers, and researchers;
- introduction to high performance computing; and,
- HPC software applications.

The Colloquium will feature presentations by principal investigators from the 10 DoD host locations, visiting faculty members, HPCMPO staff, and JEOM summer interns.

The HPCMP Education and Outreach program will continue to grow the JEOM internships, the summer colloquium, and the summer institutes to accommodate the rising interest for talented STEM students and faculty members at DoD HPC Centers and laboratories nationwide.

GOAL 5: CONTINUOUSLY EDUCATE THE RDT&E WORKFORCE WITH THE KNOWLEDGE NEEDED TO EMPLOY COMPUTATIONAL MODELING EFFECTIVELY AND EFFICIENTLY

User Productivity Enhancement and Technology Transfer (PET)

The PET contracts offered 45 training events this past year, attended by 561 students, covering subjects ranging from code profiling and error estimators to user training on codes such as FLUENT, ABAQUS®, EnSight, and Xpatch®. See Table 5 for a sampling of courses given in FY 2007. Many PET courses are

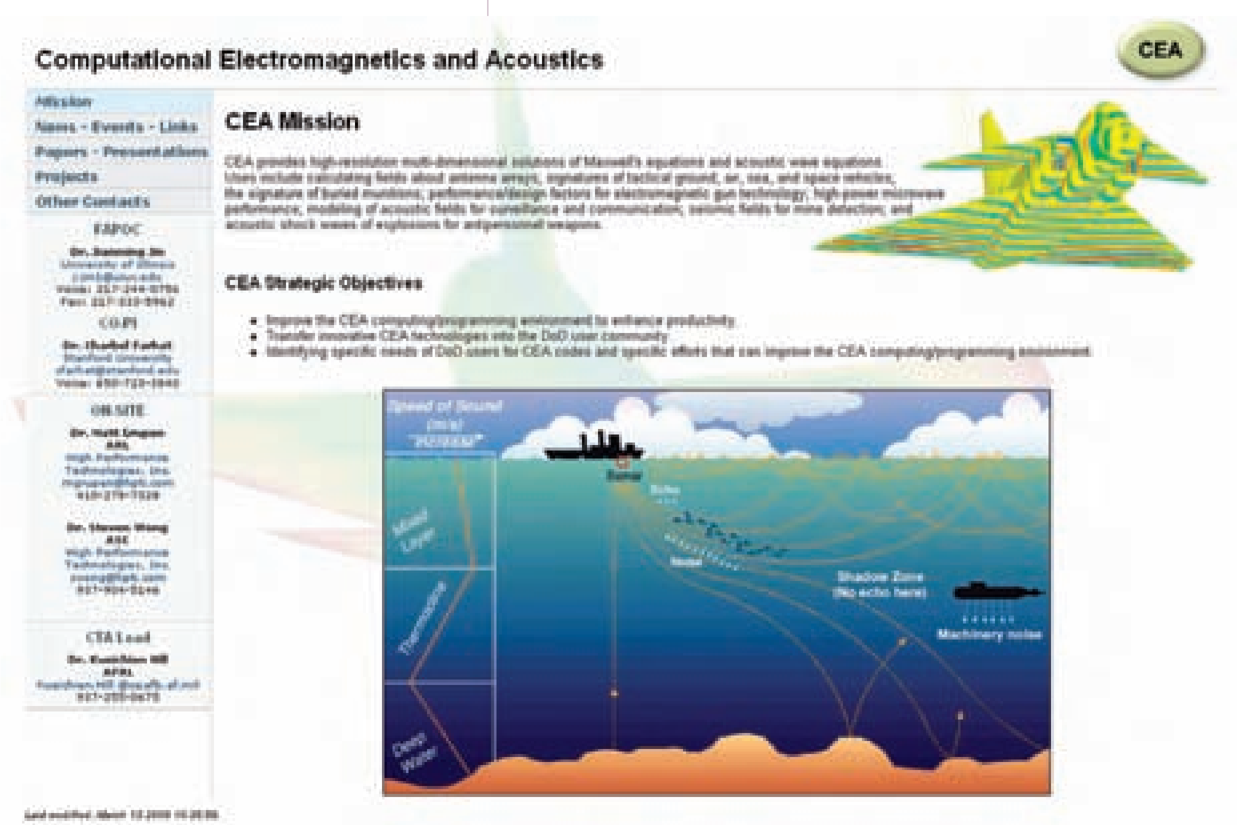

Figure 16. Sample of OKC webpage

captured on video and can be downloaded from the PET OKC (https://okc.erdc.hpc.mil) onto the users' desktops and viewed at their leisure.

While the OKC contains online training opportunities, it also is a repository of PET-developed technical reports, presentations, and points-of-contact. The computational electromagnetics and acoustics OKC front page (Figure 16) is typical of those for all the PET functional areas.

Table 5. A Sampling of Training Courses Given in FY 2007

Course Title	Number of Attendees
ABAQUS	17
Advanced GEMACS and CEM Framework	8
A Workshop on Crosslight's APSYS Package	6
SolidMesh	9
Parallel Performance Evaluation Tools Workshop	8
Numerical Libraries Workshop	11
CFD Tools and Technology Workshop	13
Programming with Parallel MATLAB MPI	18
Python for Signal Processing	17
Contact in LS-DYNA	7
Intermediate-Advanced Ensight	10
Paraview	12
Perl Programming and Applications	11
MARS and MARS/GEMINI	12
FORTRAN 90 Overview and Intro to MPI	12
Monte Carlo Techniques	7
Finite Element Tearing and Interconnecting (FETI) for Electromagnetics	12
ENS Seminar on Photonic Crystal Simulation	12
Intel Advanced Optimization Techniques for EM64T	8
CARLOS	15
Fundamentals of Xpatch	20
DOE Atomistic & Quantum Modeling and Molecular Dynamics Codes	18
Workshop on HPC Network Simulation for the DoD	25
Study of the Impacts of the Levels of Detail and Resolution for Virtual Target Models	24
Biomolecular Modeling Workshop	26

Section 3

Financial Statements

FINANCIAL STATEMENTS

FY 2007 BUDGET RESOURCES

Financial Analysis

HPCMP funds are used for: 1) capitalization, sustainment, and operations at the MSRCs; 2) annual capitalization for selected ADCs and DHPIs; 3) wide area network services for the DoD HPC community; 4) investments in human capital and key HPC software applications; and 5) expert HPC services from leading academic institutions. Figure 17 displays FY 2007 spending by component and Figure 18 shows FY 2008 planned spending by component.

We use multiple contracting officers in support of different efforts. Contracting officers at the General Services Administration support HPC equipment and services purchases, and contracting officers at various DoD installations support our service contracts. This structure is necessary because the program requires multiple contracts and contract types to ensure that state-of-the-art technical capabilities are made available to DoD scientists and engineers in a timely manner. Contracts are a combination of firm fixed price, cost, and/or indefinite delivery/indefinite quantity. All procurement awards are made for commercially available systems. Acquisitions are accomplished competitively to the fullest extent possible and encourage the inclusion of small, disadvantaged businesses and MSIs.

We evaluate the effectiveness of each program component by measuring actual cost and schedule performance versus planned cost and schedule performance and through the measurement of actual

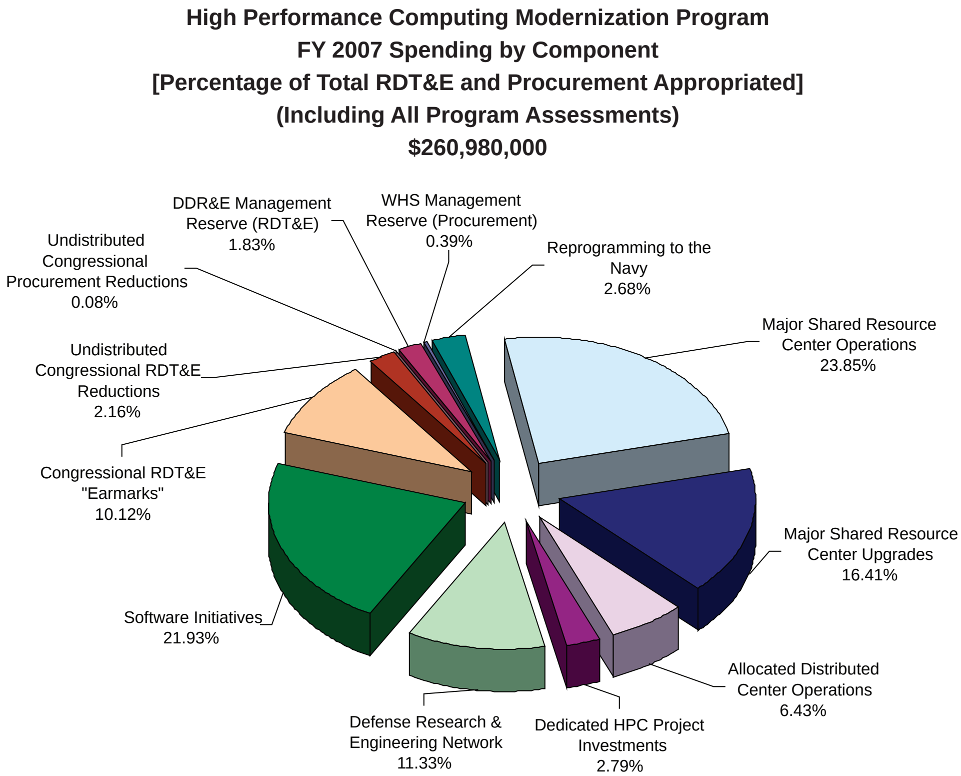

Figure 17. HPCMP FY 2007 spending by component

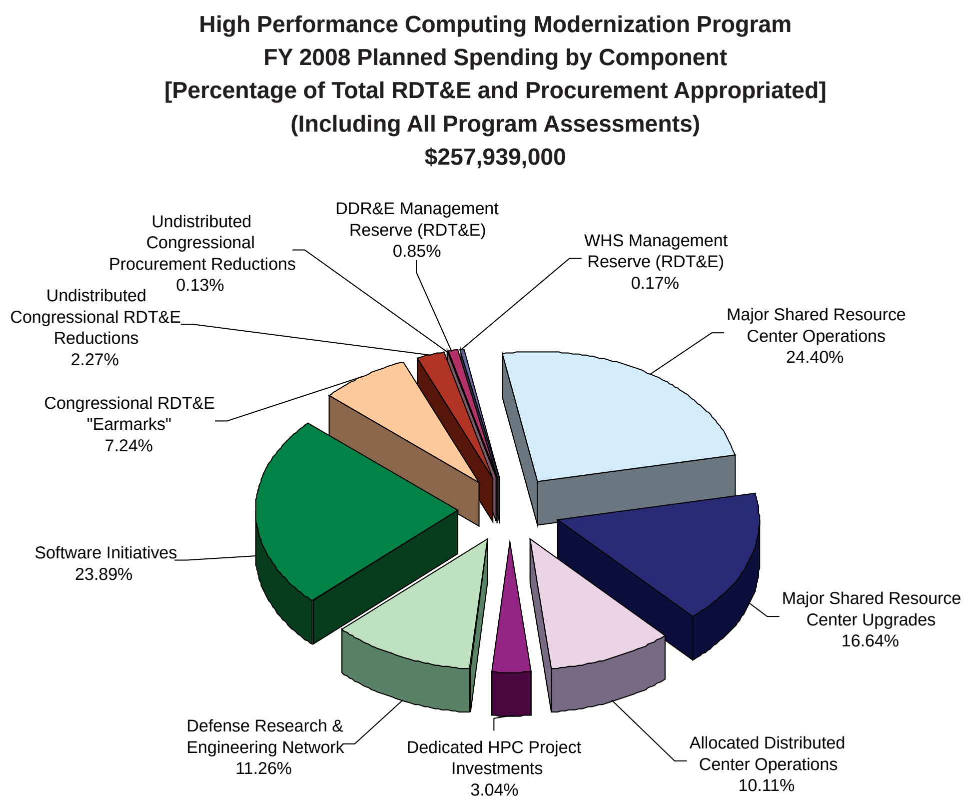

Figure 18. HPCMP FY 2008 planned spending by component

Budget Resources—continued

outcomes versus planned outcomes. The MSRC contractors submit several reports regularly including a monthly and quarterly cost performance report and quarterly contract funds status report. Each contract specifies, as a deliverable, a work breakdown structure to facilitate the review of smaller task components. Cost/schedule status reports are one of the primary tools used for oversight management of the MSRCs.

The balance sheet on page 94 shows the cumulative value of the program.

Obligations and Costs

Our Financial Manager conducts semi-annual reviews with each major component manager and major field activity to review actual cost performance against budgeted cost goals in a tailored work breakdown structure format with special attention on variance analysis. Significant variances are reported to our Program Director and corrective actions taken. We receive approximately $250,000,000 each year in funding appropriated for the DoD. Cash flow during 2007 is illustrated by the Cash Flow Statement on page 95.

While the program has leveraged major cost performance improvements in computer technology since 1994, validated requirements have always exceeded the computing capability available to address those requirements. This occurs: 1) because the use of science-based models and simulations to answer research questions and solve engineering problems has increased dramatically; and 2) because fully funding the HPC requirement is unaffordable given the entire scope of activities the DoD budget must address. While fiscal resources do not fully meet the computational requirements of the science and technology and test and evaluation communities, the returns provided are substantial with resources allocated to the highest priority projects. The FY 2007 Income Statement on page 96 shows these shortfalls.

High Performance Computing Modernization Program
Balance Sheet
As of February 29, 2008

Assets and Equity

Hardware (Total investments since program start)		$1,087,175,741	
Less: Cummulative Depreciation			
	Fiscal Year 1994–2004:	$876,719,117	
	Fiscal Year 2005:	$37,436,429	
	Fiscal Year 2006:	$22,614,429	
	Fiscal Year 2007:	$7,155,571	
	Fiscal Year 2008:	$0	$143,250,195
Software[1] (Total investments since program start)		$275,197,895	
Less: Cummulative Depreciation			
	Fiscal Year 1994–2002:	$144,482,111	
	Fiscal Year 2003:	$16,843,239	
	Fiscal Year 2004:	$12,049,047	
	Fiscal Year 2005:	$10,812,484	
	Fiscal Year 2006:	$6,473,330	
	Fiscal Year 2007:	$2,455,179	
	Fiscal Year 2008:	$0	$82,082,505
Manpower Contracts[2, 3] Software Development			
Exercised Contract Value		$31,773,301	
Less: Value Consumed Remaining		$25,042,620	
Exercised Value			$6,730,681
Maintenance Contracts[2, 3] Hardware Maintenance			
	Fiscal Year 2008:	$14,648,877	
Software Maintenance			
	Fiscal Year 2008:	$1,286,486	
Less: Value Consumed Hardware Maintenance			
	Fiscal Year 2008:	$10,016,077	
Software Maintenance			
	Fiscal Year 2008:	$889,562	$5,029,724
Intellectual/Operations Government Labor			
	Fiscal Year 2008:	$21,190,152	
Contract Labor			
	Fiscal Year 2008:	$90,507,498	
Less: Value Consumed Government Labor			
	Fiscal Year 2008:	$15,194,695	
Contract Labor			
	Fiscal Year 2008:	$68,754,330	$27,748,625
Facilities			
	Fiscal Year 2008:	$38,897,687	
Less: Value Consumed			
	Fiscal Year 2008:	$27,863,298	$11,034,389
		Total Assets	**$275,876,119**

Liabilities

Uncompleted Software Development		$6,730,681	
Maintenance Contract Liabilities Hardware			
	Fiscal Year 2008:	$4,632,800	
Software			
	Fiscal Year 2008:	$396,923	
Intellectual/Facilities Expense Government Labor			
	Fiscal Year 2008:	$5,995,458	
Contract Labor			
	Fiscal Year 2008:	$21,753,168	
Facilities			
	Fiscal Year 2008:	$11,034,388	
	Total Liabilities	**$50,543,418**	
		Program Equity	**$225,332,701**
		Total Liability and Program Equity	**$275,876,119**

(1) Research, Development and Engineering funding used to develop inventory software.

(2) Office of Management and Budget Circular A-11, Section 300 - Planning, Budgeting, Acquisition, and Management of Capital Assets (Paragraph 300.4), defines capital assets as land, structures, equipment, intellectual property (e.g., software), and information technology (including IT service contracts) that are used by the Federal government and have an estimated useful life of two years or more. Therefore, manpower is treated as a capital asset.

(3) Small consumable items such as computer tapes and supplies are considered as expense items and not carried as inventory items.

High Performance Computing Modernization Program
Cash Flow Statement
October 1, 2006 — September 30, 2007

	Fiscal Year 2007
Revenue	
Research, Development and Engineering Funding	
President's Budget	$175,313,000
Congressional Funding	$34,350,000
Department of Defense Reprogramming - In	$0
(Less Congressional Undistributed Reductions)	**($12,632,000)**
(Less Unreleased Obligation Authority)	**($4,677,134)**
Net Research, Development and Engineering Funding	**$192,353,866**
Procurement Funding	
President's Budget	$51,317,000
Congressional Funding	$0
Department of Defense Reprogramming - In	$0
(Less Congressional Undistributed Reductions)	**($208,000)**
(Less Unreleased Obligation Authority)	**($1,020,000)**
Net Procurement Funding	**$50,089,000**
Net Revenue	**$242,442,866**
Investments	
Major Shared Resource Center Upgrades	$43,815,000
Allocated Distributed Center Upgrades/Dedicated HPC Project Investments	$6,274,000
Software Development	$24,551,790
Expense	
Major Shared Resource Center Operations	$69,614,173
Allocated Distributed Center Operations	$33,989,391
Defense Research & Engineering Network	$32,577,690
Software Initiatives	$31,620,822
Net Expense	**$242,442,866**
Balance (As of September 30, 2007)	**$0**

High Performance Computing Modernization Program
Income Statement
October 1, 2006 — September 30, 2007

	Fiscal Year 2007
Income	
Research, Development and Engineering Funding	
Major Shared Resource Center Operations	$69,614,173
Allocated Distributed Center Operations	$33,989,391
Defense Research & Engineering Network	$32,577,690
Software Initiatives	$56,172,612
Procurement Funding	
Major Shared Resource Center Upgrades	$43,815,000
Allocated Distributed Center Upgrades	$6,274,000
Defense Research & Engineering Network	$0
Software Initiatives	$0
Total Income	**$242,442,866**
Expense[1]	
Research, Development and Engineering Funding	
Major Shared Resource Center Operations	$69,614,173
Allocated Distributed Center Operations	$33,989,391
Defense Research & Engineering Network	$32,577,690
Software Initiatives[2]	$31,620,822
Depreciation of Capital Assets	
Hardware (Depreciated based upon a 48-month life-cycle)[3]	$49,857,000
Software (Depreciated based upon a 60-month life-cycle)[4]	$19,896,726
Total Expense	**$237,555,802**
Balance (As of September 30, 2007)	**$4,887,064****

** This positive number reflects an increase in the value of the Program's capital assets. The Program has invested more this year than has been depreciated. Thus, the value of the inventory has gone up by the amount shown.

Note 1: Expenses include travel; supplies; government and contractor salaries and training; maintenance of hardware and software; studies and analysis; annual operations investments; communications, utilities, facilities lease, and facilities maintenance.

Note 2: Software initiatives are separated into 2 distinct categories—expenses associated with research and development, management, education/training, and expert services; and capitol assets resulting from developed software.

Note 3: Depreciation for HPC hardware is calculated using a 48-month straight-line depreciation method. Current HPC technology development results in predictable obsolescence. Generally, after 48 months of use, HPC systems are retired with little or no residual value. Fiscal year 2007 depreciation includes the 12-month value calculated for all systems in the inventory between October 1, 2006 through September 30, 2007.

Note 4: Depreciation for HPC software is calculated using a 60-month straight-line depreciation method. A period of 60 months is used because it is the typical life cycle of HPC software before significant modifications are required. Fiscal year 2007 depreciation includes the 12-month value calculated for all software in the inventory between October 1, 2006 through September 30, 2007.

Budget Resources—continued

Financial Trends

Except for minimal inflation adjustments, HPC budgets are essentially flat. We address urgent new requirements by adjusting priorities within the existing funding profile. We increased the overall capability of our HPC systems by about 80%, and added or upgraded systems at the ADCs. However, even with these increased capabilities, we are unable to meet validated DoD requirements. Development of the portfolios and institutes will continue. The DoD HPC user community will continue to be supported by the PET efforts. Our Software Protection Initiative will continue to mature. Figure 18 on page 92, breaks out program-wide and planned spending for 2008.

The Income Statement on page 96 shows that currently we have a continuing deficit. The dollars we spend are not keeping up with the rapidly growing needs of the scientific community. Figure 19 displays spending by vendor in FY 2007 and Figure 20 shows planned spending by vendor for FY 2008.

Summary

We deploy, sustain, and upgrade commercially available high performance computing environments and networking services in support of DoD laboratories and test facilities. We have substantially improved the Department's computational capabilities with the objective of providing the DoD the technology to ensure dominance on the battlefield by the early fielding of the most advanced computing capability available.

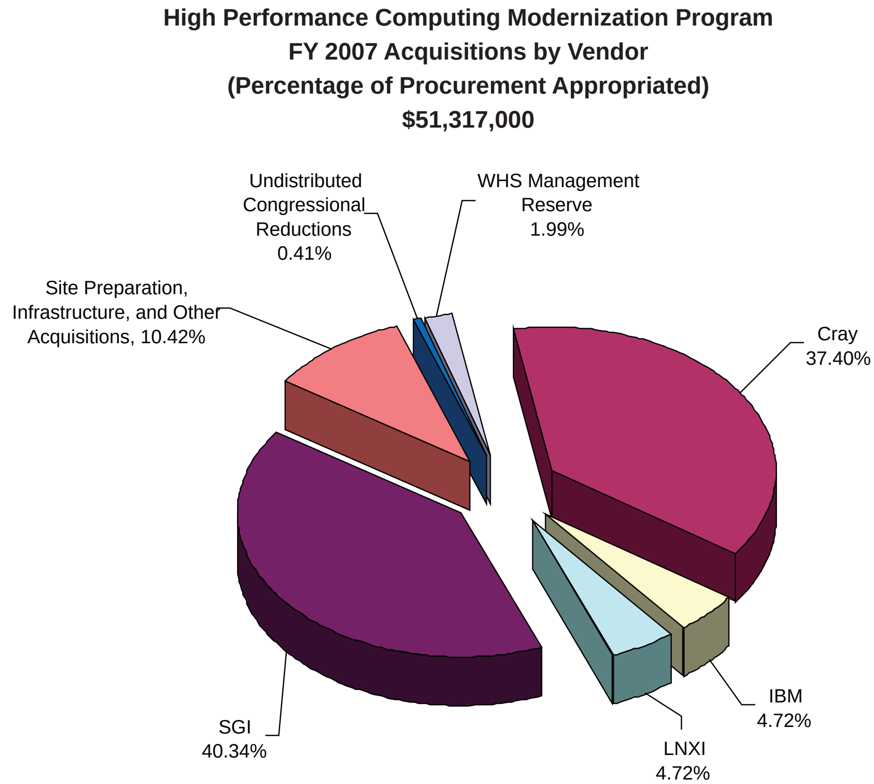

Figure 19. HPCMP FY 2007 acquisitions by vendor

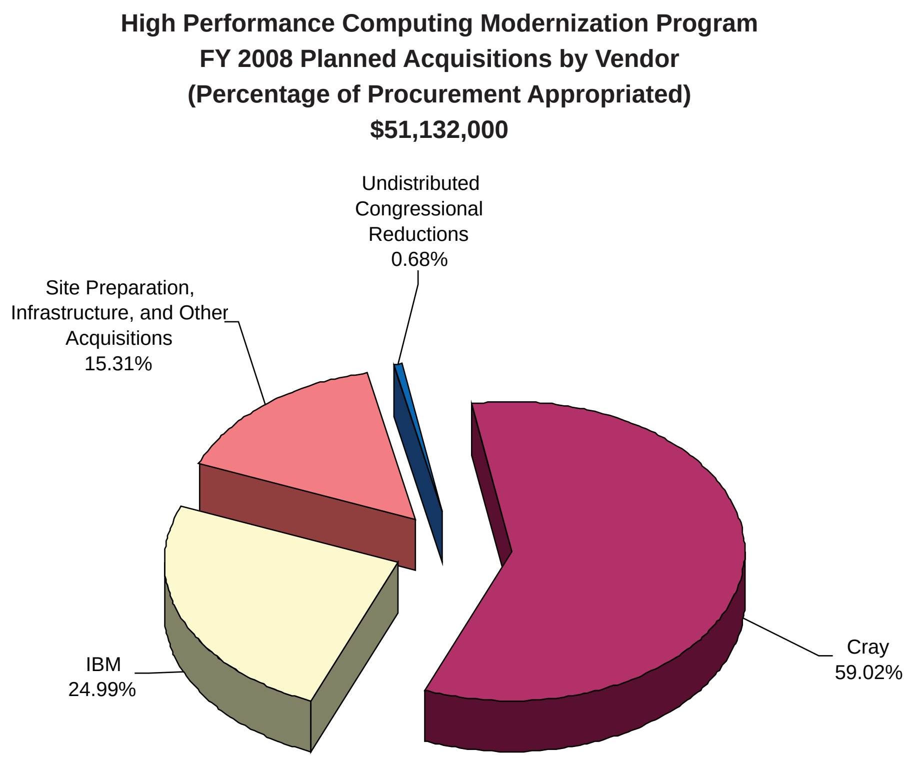

Figure 20. HPCMP FY 2008 planned acquisitions by vendor

Acronyms

Acronyms

2D	two-dimensional
3D	three-dimensional
AAS&C	aircraft and armament stability & control
ADCs	allocated distributed centers
AFB	Air Force Base
AFRL/RHDR	Air Force Research Laboratory, Human Effectiveness Directorate, Directed Energy Bioeffects Division, Radiofrequency Radiation Branch
AHPCRC	Army High Performance Computing Research Center
APPSPack	asynchronous parallel pattern search; APPSPack is a derivative-free optimization software for solving non-linearunconstrained, bound-constrained, and linearly-constrained optimization problems.
ARL	Army Research Laboratory
ARSC	Arctic Region Supercomputing Center
ASC	Aeronautical Systems Center
BE	broader engagement
BEI	Battlespace Environments Institute
BFI	bullet and fragment impacts
BHSAI	Biotechnology HPC Software Applications Institute for Force Health Protection
C4ISR	command, control, communications, computers, intelligence, surveillance and reconnaissance
CAP	capability applications projects
CBoD	Centers Board of Directors
CCM	computational chemistry, biology, and materials science
CDLT	collaborative and distance learning technologies
CE	computational environment
CEA	computation electromagnetics and acoustics
CFD	computational fluid dynamics
CICE	sea ice model
CNMOC	Commander, Naval Meteorology and Oceanography Center

CPU	central processing unit
CREATE	Computational Research and Engineering Acquisition Tools and Environments
CSM	computational structural mechanics
CST	collaborative simulation and testing
CTAs	computational technology areas
CWO	climate/weather/ocean modeling and simulation
DHPIs	Dedicated HPC Project Investments
DIOT	Distributed Implementation and Operations Team
DoD	Department of Defense
DOE	Department of Energy
DREN	Defense Research and Engineering Network
DUSD(S&T)	Deputy Under Secretary of Defense (Science and Technology)
ED&PMB	Engineering Design and Process Management Board
EM	electromagnetic
ENS	electronics, networking, and systems/C4I
EOTC	education, outreach, and training coordination
ERDC	Engineer Research and Development Center (USACE)
EQM	environmental quality modeling and simulation
ESMF	Earth Systems Modeling Framework
ET	enabling technologies
FLOPS	FLoating-point OPerations per Second
FMS	forces modeling and simulation
FDTD	finite difference time domain
FY	fiscal year
GFs	gigaflops
GPS	global positioning system
HBCUs	Historically Black Colleges and Universities
HI-ARMS	HPC Institute for Advanced Rotorcraft Modeling and Simulation
HPC	high performance computing or high performance computer
HPCMP	High Performance Computing Modernization Program
HPCMPO	High Performance Computing Modernization Program Office
HPM	high power microwave
HYCOM	Hybrid Coordinate Ocean Model

IHAAA	Institute for HPC Applications to Air Armament
ICD	initial capabilities document
ICEPIC	Improved Concurrent Electromagnetic Particle-In-Cell
IM	insensitive munitions
IMT	integrated modeling and test environments
IMTPS	Institute for Maneuverability and Terrain Physics Simulation
I-SSA	Institute for Space Situational Awareness
JEOM	Joint Educational Opportunities for Minorities
JET	Joint Engineering Team
JHL	Joint Heavy Lift Rotorcraft program
JFCOM	Joint Forces Command
LSN	large scale network
MDA	Missile Defense Agency
MFT	multiphase flow target response
MHPCC	Maui High Performance Computing Center
MSIs	Minority Serving Institutions
MSRCs	major shared resource centers
NASA	National Aeronautics and Space Administration
NAVAIR	Naval Air Systems Command
NAVO	Naval Oceanographic Office
NCA&T	North Carolina Agricultural and Technical University
NOGAPS	Navy Operational Global Atmospheric Prediction System
OC	optical carrier
OKC	Online Knowledge Center
OneSAF	One Semi-Automated Forces
OOS	One Semi-Automated Forces (OneSAF) Objective System (OOS)
OSD	Office of the Secretary of Defense
OSD (PA&E)	Secretary of Defense (Program Analysis and Evaluation)
PCID	physically-constrained iterative deconvolution
PET	User Productivity Enhancement and Technology Transfer
PEUO	physics-based environment for urban operations
RDT&E	research, development, test, and evaluation
RF	radio frequency
ROI	return-on-investment

S&T	science and technology
SAS	Software Applications Support
SC	supercomputing
SIP	signal/image processing
SMDC	Army Space and Missile Defense Command
SOF	Special Operations Force
SSA	space situational awareness
SSN	Space Surveillance Network
STAP	space-time adaptive processing
STEM	science, technology, engineering, and mathematics
T&E	test and evaluation
TI	technology insertion
TPD	technical planning document
UGC	Users Group Conference
US	United States
USACE	US Army Corps of Engineers
USAMRIID	US Army Medical Research Institute of Infectious Diseases
UTC	user training coordination
VED	virtual electromagnetics design
WAN	wide area network

HIGH PERFORMANCE COMPUTING MODERNIZATION PROGRAM OFFICE
10501 FURNACE ROAD, SUITE 101
LORTON, VA 22079
703.812.8205

www.hpcmo.hpc.mil