HPCinsights

DoD High Performance Computing Modernization Program
Fall 2009
SC09 Edition

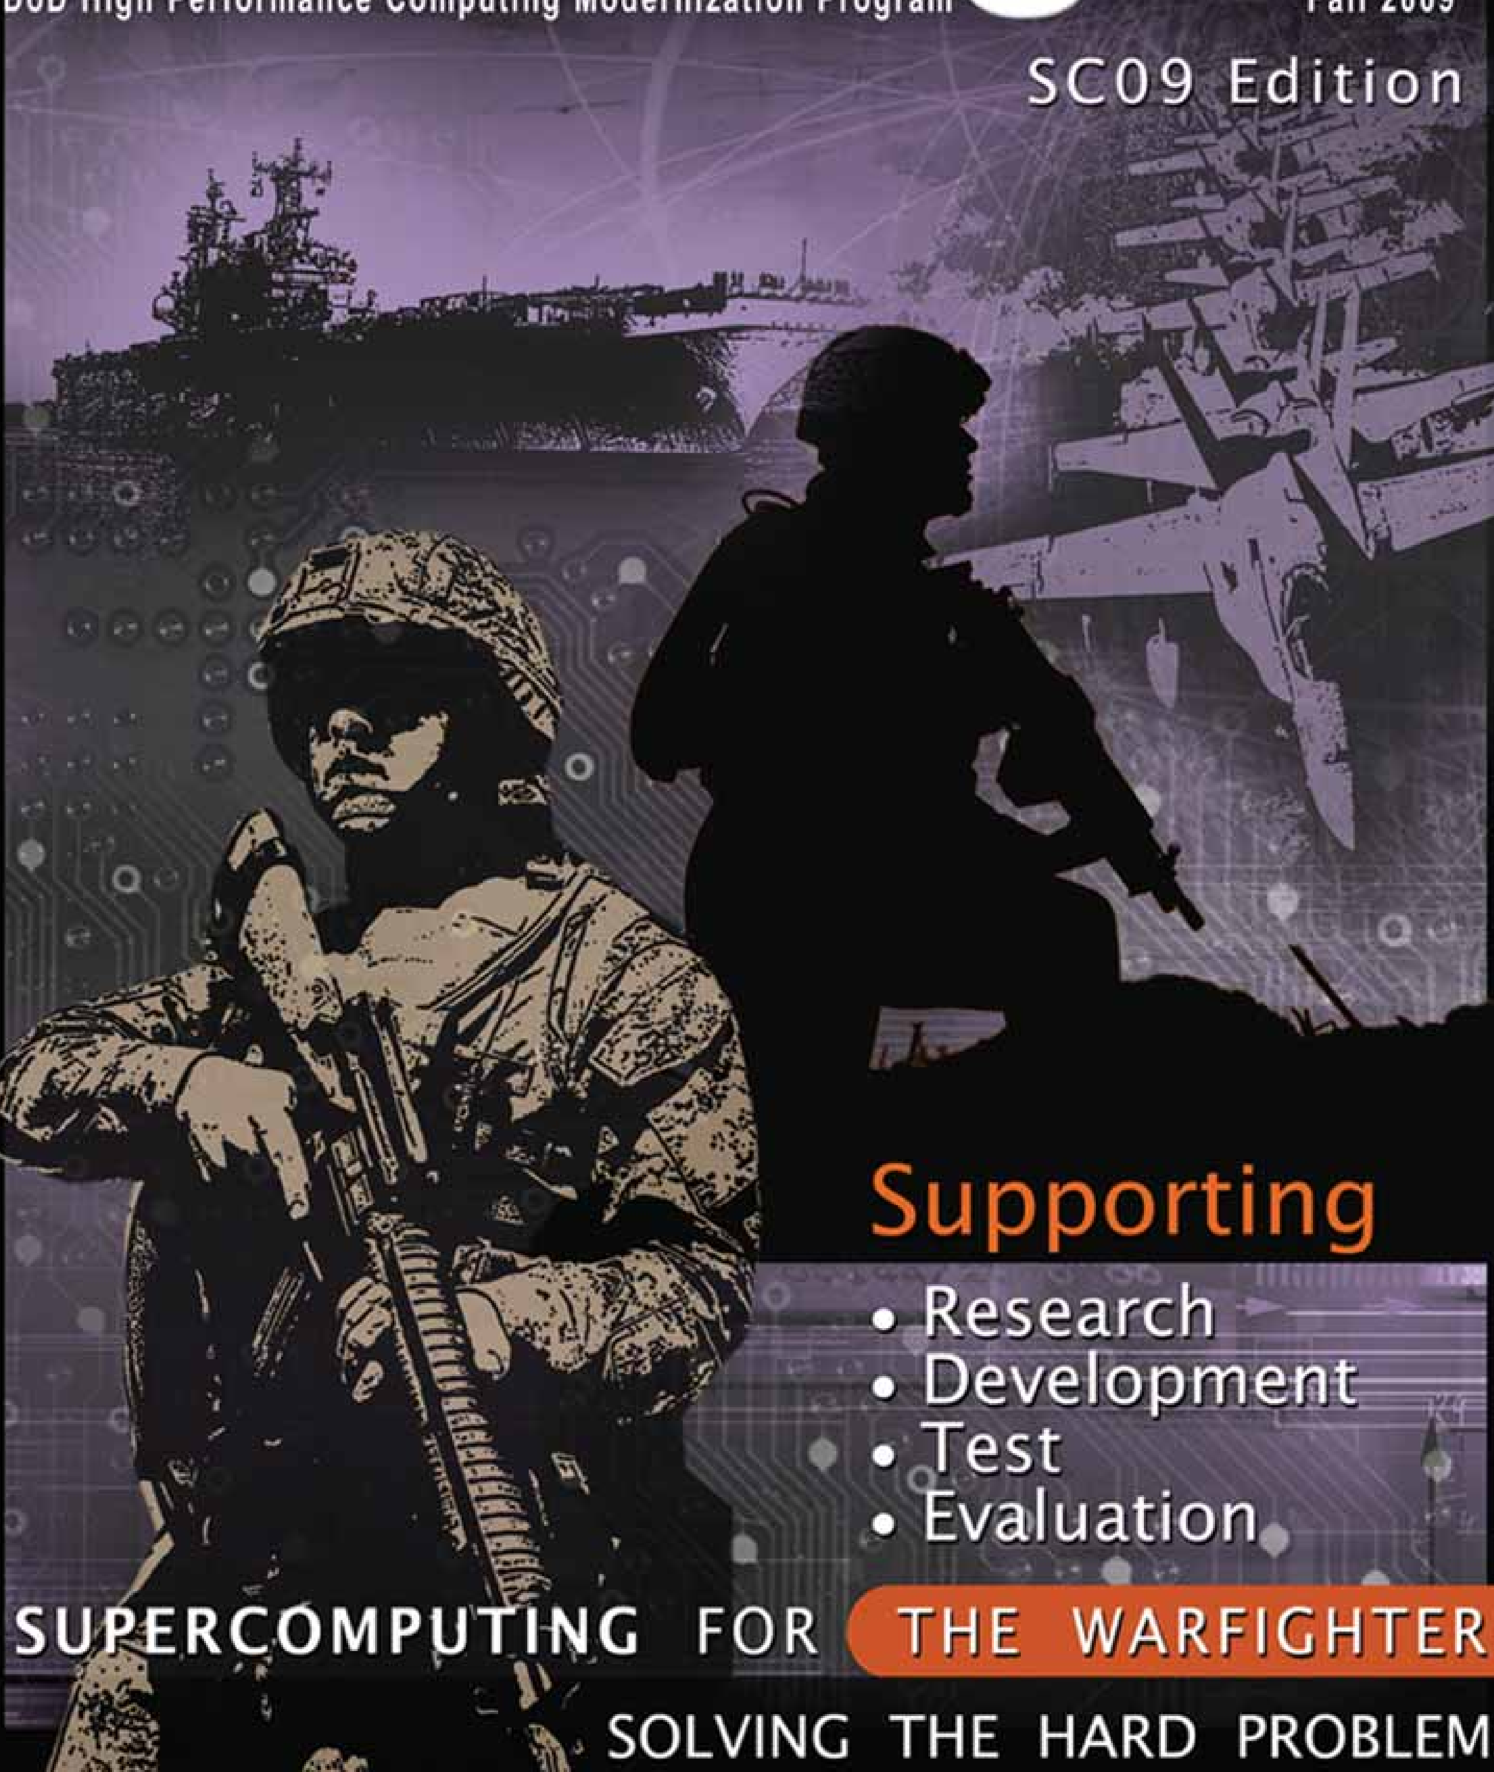

Supporting

- Research
- Development
- Test
- Evaluation

SUPERCOMPUTING FOR THE WARFIGHTER

SOLVING THE HARD PROBLEMS

HPC Insights is a semiannual publication of the Department of Defense Supercomputing Resource Centers under the auspices of the High Performance Computing Modernization Program.

Publication Team

AFRL DSRC, Wright-Patterson Air Force Base, OH
Joan Henley
Chuck Abruzzino

ARL DSRC, Aberdeen Proving Ground, MD
Debbie Thompson
Brian Simmonds

ARSC DSRC, Fairbanks, AK
Debra Damron
Mary Haley

ERDC DSRC, Vicksburg, MS
Rose J. Dykes

MHPCC DSRC, Maui, HI
Jeff Schmidt
Betty Duncan

Navy DSRC, Stennis Space Center, MS
Christine Cuicchi
Lynn Yott

HPCMPO, Lorton, VA
Deborah Schwartz

MANAGING EDITOR
Rose J. Dykes

DESIGN/LAYOUT
Betty Watson

COVER DESIGN

Contents

Department of Defense High Performance Computing Modernization Program

By Cray Henry, Director

Through research, development, test and evaluation (RDT&E), our high performance computing (HPC) community ultimately provides enhanced capability for the Warfighter and enhances our Nation's security. The Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) provides the DoD HPC community the 21st century computational tools they need to succeed.

This issue of *HPC Insights* features articles on current research and developments projects from across the country. The DoD HPC community is working on hundreds of projects focused on protecting the homeland and our forward deployed troops around the world through enhanced science-based simulations, numerical predictions, and analysis. This small sample of work is an impressive testament to the dedication and innovation this community exhibits toward our national defense.

Within the HPCMP, we are taking giant steps to meet customer requirements to support the pursuit of modernization, so apparent in this issue of *HPC Insights*. Research among the Centers requires a robust environment for managing and processing information. Engineers and scientists want easy-to-use, accessible resources. On this front, the Program is focusing on what matters most: system responsiveness, usability of our resources, and value added services aimed to enhance the customer experience at each of our DoD Supercomputing Resource Centers (DSRCs).

By building a foundation to support joint computational environments, tools, and applications, the HPCMP allows scientists and engineers to enhance processes, evaluate numerical simulations and improve computational methods. Through an enhanced storage system and common utility-enhancement server, to be deployed in FY2010, HPC users will have easier access to a wider variety of services at all DSRCs.

On behalf of the Program, I want to encourage you to take a moment to read the collective accomplishments of our HPC community. Our community really is focused on "solving the hard problems."

HPC Resource Requirements for Hurricane Storm Surge Simulations

By Mary A. Cialone, Ty V. Wamsley, Jane M. Smith, and Robert E. Jensen, U.S. Army Engineer Research and Development Center Coastal and Hydraulics Laboratory

Introduction

The vulnerability of coastal areas to storm attack, inundation, and sometimes destruction by storm surge and waves is a reality that may unfortunately increase with the ever-increasing population that seeks to reside along the coast and its accompanying development/ infrastructure. In addition, this threat might be exacerbated by an increase in frequency and severity of the hurricane hazard that some scientists predict. On a positive note, the ability to examine storm response (surge and waves) for a suite of synthetic storms provides a predictive tool for estimating the potential hazard along the coast. The simulation suite to generate statistical surge levels and waves is generally on the order of 100 to 300 storms, requiring substantial computer resources. This article introduces the models applied and the high performance computing (HPC) resources required to simulate a storm suite for development of statistical surge and wave response to hurricanes.

Modeling System

Following Hurricanes Katrina and Rita in 2005, a team of engineers and scientists specializing in coastal hydrodynamics, meteorology, statistics, and computer science collaborated to develop an integrated modeling system to estimate storm surge inundation. A schematic diagram of the system is shown in Figure 1. The modeling system was validated for Hurricanes Katrina and Rita because of the availability of high-quality storm system data and high water marks from these storms. In addition, the extent of inland inundation for these storms was extreme, allowing for a unique opportunity to validate the effectiveness of modeling the impact of topography, overland resistance, and decreases in overland wind speeds. With the system validated, a synthetic storm suite was constructed and simulated to examine the range of potential surge and wave responses in a given area.

Model Coupling/Synthetic Storms

The coupled modeling system summarized in Figure 1 includes models for simulating hurricane wind fields, wave generation and transformation, and storm surge. For a given synthetic storm parameter set (consisting of a track and time-varying wind field parameters), application of the TC96 Planetary Boundary Layer (PBL) model (Thompson and Cardone 1996) generates a time-series of wind and atmospheric pressure fields that are used to drive the offshore wave model WAM (Komen et al. 1994) and the storm surge model ADCIRC (Luettich et al. 1992, Westerink et al. 1994). The large-domain, discrete, time-dependent spectral wave model WAM is run to calculate directional wave spectra that serve as boundary conditions for the local-domain, near-coast wave generation and transformation model STWAVE (Smith, Sherlock, and Resio 2001, Smith and Sherlock 2007). In parallel with the WAM runs, the unstructured coastal ocean circulation model ADCIRC is used to simulate pressure- and wind-driven water surface elevation (storm surge) and is coupled with the nearshore wave model STWAVE. The input for each STWAVE simulation includes the bathymetry, variable bottom friction, surge, and wind (all interpolated from the ADCIRC domain). The STWAVE simulations are also forced with wave spectra interpolated on the offshore boundary from the WAM model. The wind and surge applied in STWAVE are spatially

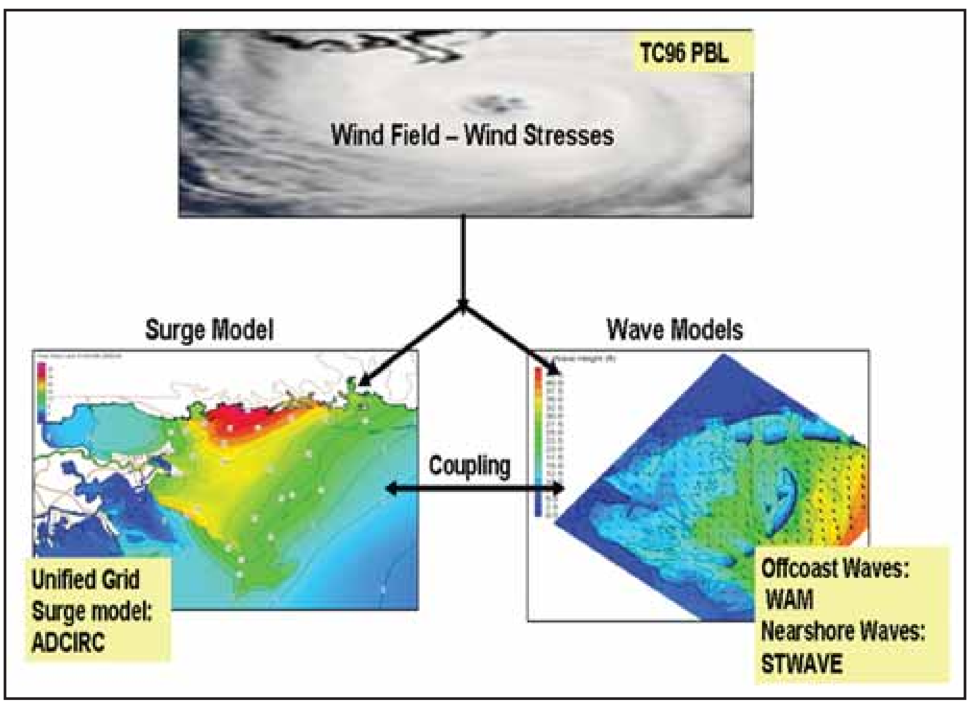

Figure 1. Diagram of modeling system for coastal inundation application

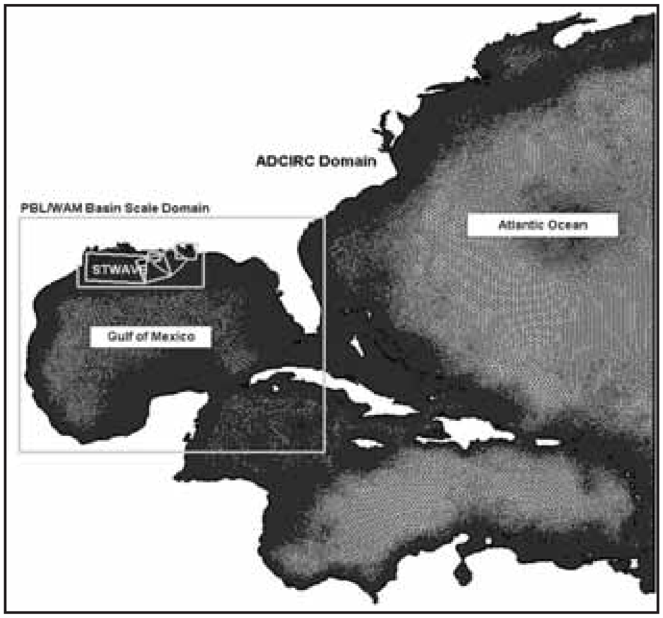

Figure 2. Model domains for system models: PBL, WAM, ADCIRC, and STWAVE

and temporally variable. STWAVE is run at 30-min intervals for approximately 1 day prior to and following storm landfall. Lastly, ADCIRC is rerun so that radiation stress gradient fields from STWAVE can be applied to ADCIRC and contribute to the computation of total water level. This procedure is repeated for each storm in the synthetic storm suite.

Model Domains

The circulation model ADCIRC covers a broad domain including the entire Gulf of Mexico and the Atlantic Ocean eastward to the 60 deg W longitude line (Figure 2). The unstructured mesh applied with ADCIRC allows for high resolution where solution gradients are large and low resolution where solution gradients are small, minimizing both local and global errors for a given computational cost (Blain et al. 1998). The ADCIRC mesh has over 2 million nodes with the majority concentrated in coastal Louisiana and Mississippi. The PBL and WAM models apply a structured grid, and their domains cover the entire Gulf of Mexico. WAM operates on approximately 61,000 grid points. The nearshore wave model STWAVE is nested within the other models and covers the nearshore zone with the seaward limit located at approximately a 30 to 50 m water depth. In the general system developed for the Louisiana and Mississippi regional applications, five STWAVE grid domains are defined to cover the overall nearshore region: the west (W) grid covers western Louisiana, the south (S) grid covers southern Louisiana, the southeast (SE) grid covers the southeastern portion of Louisiana and the western portion of Mississippi, the Pontchartrain (Pont) grid covers Lake Pontchartrain, and the Mississippi-Alabama (MS-AL) grid covers the Mississippi coast and a portion of Alabama. The five STWAVE grids total on the order of 3 million grid cells.

HPC Resources

With the model domains generated and a storm suite selected, a set of simulations are made. The PBL model is run in a stand-alone mode for all storms and does not require HPC resources. All wind and pressure fields generated from this task are transferred to the high performance computer for application by the other models. WAM and ADCIRC are executed simultaneously using the PBL wind and pressure fields as forcing conditions. Each WAM synthetic storm simulation is performed in serial mode with a simulation time requirement of 12 to 28 hr. ADCIRC is initially run for 2 days to allow the Mississippi and Atchafalaya rivers to reach a dynamic steady-state balance. This river spin-up simulation, referred to as ADCIRC0 in Figure 3, is run in parallel on approximately 256 processors, requiring approximately 2 hr to complete. ADCIRC0 is only run once and is applied to all storm simulations.

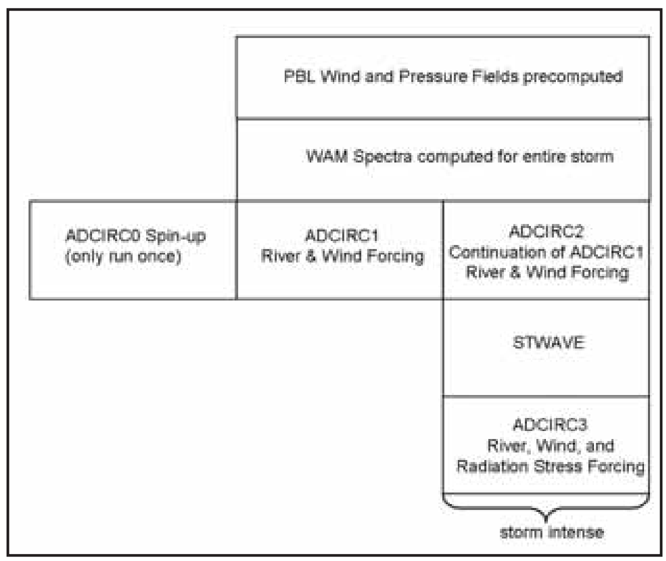

Figure 3. Coupled modeling system flow chart

With the initial ADCIRC spin-up simulation completed, the river boundary conditions are changed to radiation boundary conditions to allow flow to propagate up the river during the storm rather than reflect back into the domain. For each storm, ADCIRC is then "hot started" from the end of the river spin-up and run, depending on the forward speed of the storm, for 2 to 4 days during the initial phase of the storm (ADCIRC1). ADCIRC1 is run in parallel on 256 processors, requiring 2.5 to 4 hr to complete. ADCIRC is then hot started from ADCIRC1 and run for 2 more days during the intense portion of the storm (ADCIRC2). (The simulation is separated into these two portions to allow the intense portion of the storm (ADCIRC2) to later be hot started again from ADCIRC1, but with the inclusion of radiation stress gradients from STWAVE.) ADCIRC2 is run in parallel on 256 processors, requiring 2.5 hr to complete. With ADCIRC2 completed, water level and winds from ADCIRC2 are interpolated to the STWAVE domain, WAM wave spectra are applied to the STWAVE offshore boundary, and STWAVE is run for the intense storm period (approximately 2 days). Each of the five STWAVE grids is run in parallel on 94 processors, requiring approximately 1 hr to complete. Radiation stress gradients from STWAVE are then interpolated to the ADCIRC domain, and ADCIRC3 is hot started from ADCIRC1 and run with the additional forcing from the wave model. ADCIRC3 is run in parallel on 256 processors, requiring 2.5 hr to complete. Table 1 summarizes the HPC requirements for completing a 300 storm suite. Such an undertaking requires approximately ¾ million processor hours for a stochastic storm surge project.This system has been successfully applied to many Corps projects; however, the entire system remains cumbersome to configure and costly to compute. This is due to the loose coupling methodology described earlier, which requires multiple applications of ADCIRC in order to get feedback from other models. A new work flow that is easier to configure and incorporates tightly coupled models is being developed to allow for timely feedback into each model for improved representation of physical processes. The tightly coupled feedback does not require multiple applications of ADCIRC, which will reduce the overall computational effort and the required time to solution.

Table 1
HPC resource requirements for a 300 storm suite

Model	# proc	hr	proc-hr	# runs	Total Proc-hr
WAM	1	12-28	12-28	300	3600-8400
ADCIRC0	256	2	512	1	512
ADCIRC1	256	2.5-4	640-1024	300	192000-307200
ADCIRC2	256	2.5	640	300	192000
STWAVE	94	5	470	300	141000
ADCIRC3	256	2.5	640	300	192000
TOTAL					721000-841000

Summary

A system of models for simulating hurricane storm surge was presented in the context of the HPC resource requirements to simulate a storm suite for development of statistical surge and wave response to hurricanes. Those resource requirements are substantial. HPC requirements for completing a 300 storm suite requires approximately ¾ million processor hours. A new work flow that is easier to configure and incorporates updated models to allow for tight coupling is being developed. The new system will significantly reduce the computational effort and HPC requirement.

References

Blain, C. A., J. J. Westerink and R. A. Luettich, 1998. Grid convergence studies for the prediction of hurricane storm surge. *Int. J. Num. Meth. Fluids*, 26, 369-401.

Thompson, E. F., and V. J. Cardone. 1996. "Practical modeling of hurricane surface wind fields." *J. Waterway, Port, Coastal Engr.*, 122, 4, 195-205.

Luettich, R. A., J. J. Westerink, and N. W. Scheffner. 1992. "ADCIRC: An advanced three-dimensional circulation model for shelves, coasts, and estuaries; Report 1, Theory and methodology of ADCIRC-2DDI and ADCIRC-3DL," Technical Report DRP-92-6, U.S. Army Engineer Waterways Experiment Station, Vicksburg, MS.

Westerink, J. J., C. A. Blain, R. A. Luettich, Jr., and N. W. Scheffner. 1994. "ADCIRC: An advanced three-dimensional circulation model for shelves coasts and estuaries, Report 2: Users manual for ADCIRC-2DDI," Dredging Research Program Technical Report DRP-92-6, U.S. Army Engineer Waterways Experiment Station, Vicksburg, MS., 156 p.

Komen, G. J., L. Cavaleri, M. Donelan, K. Hasselmann, S. Hasselmann and P.A.E.M. Janssen. 1994. *Dynamics and modelling of ocean waves*. Cambridge University Press, Cambridge, UK, 560 pages.

Smith, J. M., A. R. Sherlock, and D. T. Resio, 2001. "STWAVE: Steady-State Spectral Wave Model User's Manual for STWAVE, Version 3.0," ERDC/CHL SR-01-1, U.S. Army Corps of Engineers Engineer Research and Development Center, Vicksburg, MS.

Modeling of Wideband Free-Field and Littoral Scattering of Minelike Targets Using STARS3D

By Saikat Dey, Naval Research Laboratory-Washington D.C. (NRL-DC), SFA Inc.; Charbel Farhat and Jean-François Dord, Stanford University; Felipe Bulat-Jara, Global Strategies Group (North America) Inc./NRL-DC; Angie Sarkissian, NRL-DC, and Christopher Kung, Army Research Laboratory

Problem

A recent User Productivity Enhancement and Technology Transfer (PET) Project, CEA-KY8-001 “Scalable High-Order Time-Domain Seismic/Acoustic Modeling for UGS Applications Project,” incorporated a high-order time-integration scheme into STARS3D. This allows STARS3D to maintain a high order of accuracy both temporally and spatially when combined with higher order spatial-basis functions such as the spectral-basis function. The integration scheme has added to the seismic capabilities through PET Project CEA-KY6-001 and has been used by project partners Luise Couchman, NRL-DC, now Office of Naval Research, and Steve Ketcham, U.S. Army Engineer Research and Development Center Cold Regions Research and Engineering Laboratory. STARS3D also has highly accurate modeling capabilities in the frequency domain useful for several applications of Navy interest. Project partner Brian Houston, NRL-DC, who has been researching structural acoustics in the littoral environment (i.e., the shallow-water, mine hunting problem), requested that the STARS3D code be applied to the modeling and validation of wideband free-field and littoral environment scattering from minelike targets. These scattering studies have key applications in antimine and harbor-protection areas. The goal of this study was to validate and evaluate the efficacy of the high-order approximations in STARS3D applied to the modeling of scattering in the littoral environment.

Methodology

We conducted two studies to validate STARS3D for free-field and littoral environments. These calculations were performed using MJM, the Army Research Laboratory DoD Supercomputing Resource Center (ARL DSRC) Woodcrest cluster. The first study was an *hp*-convergence study of the free-field response of an air-filled elastic shell where the monostatic response of a plane wave at various angles was examined while the elastic and fluid degree of approximation was varied. From a numerical standpoint and for the range of frequencies, the finer *h2* mesh with $p = 3$ for elastic and fluid fields converged. Figure 1 shows the elastic shell, and Figure 2 shows the convergence of the monostatic response. The second study examined the littoral response of a spherical shell buried in underwater sediment. For this study, a plane wave impinges upon the water/sediment interface at a given angle θ, and the monostatic response at two incident angles (20° and 30°) is registered in the 1 – 10 kHz range and compared with the results obtained using the T-Matrix method [1, 2]. As shown in Figure 3b, the results agree very well for all frequency ranges considered. The monostatic response at 20° shows slightly more deviation from the T-Matrix result than the response at 30°, especially at frequencies higher than 5 kHz, but the overall agreement is very satisfactory.

Importance of HPC in Mine Detection

In the past, acoustic scattering problems focused on the proper modeling of the scatterer; however, research has shown that the surrounding environment around the scatterer (sediment, water, water/sediment interface, water/air interface) contributes more significantly to the scattering results, so the focus has shifted to the

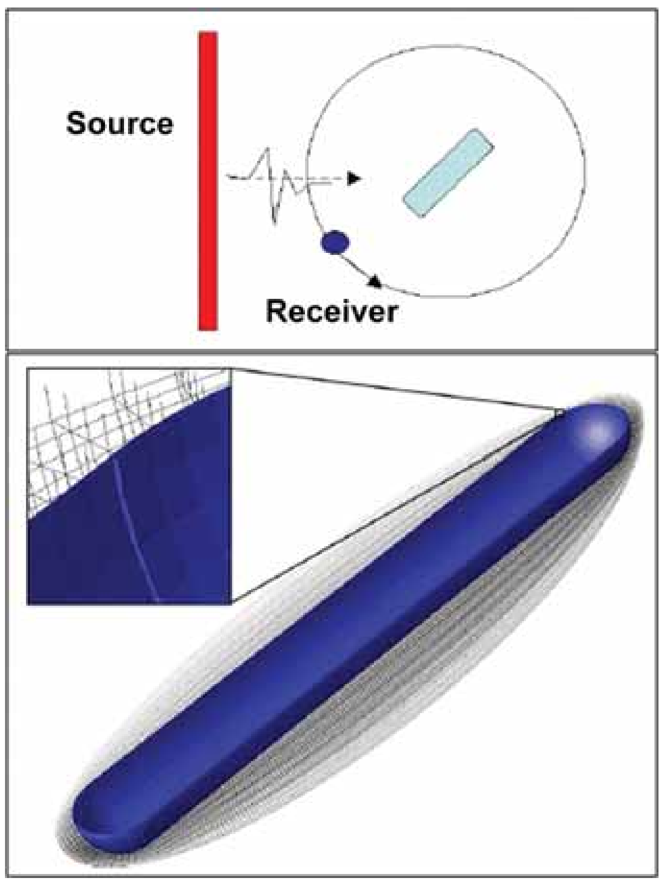

Figure 1. Scattering from an E50 shell: (a) schematic description of the problem and (b) view of 3-D finite element mesh “h1.” The refined “h2” mesh is obtained by splitting mesh edges by half

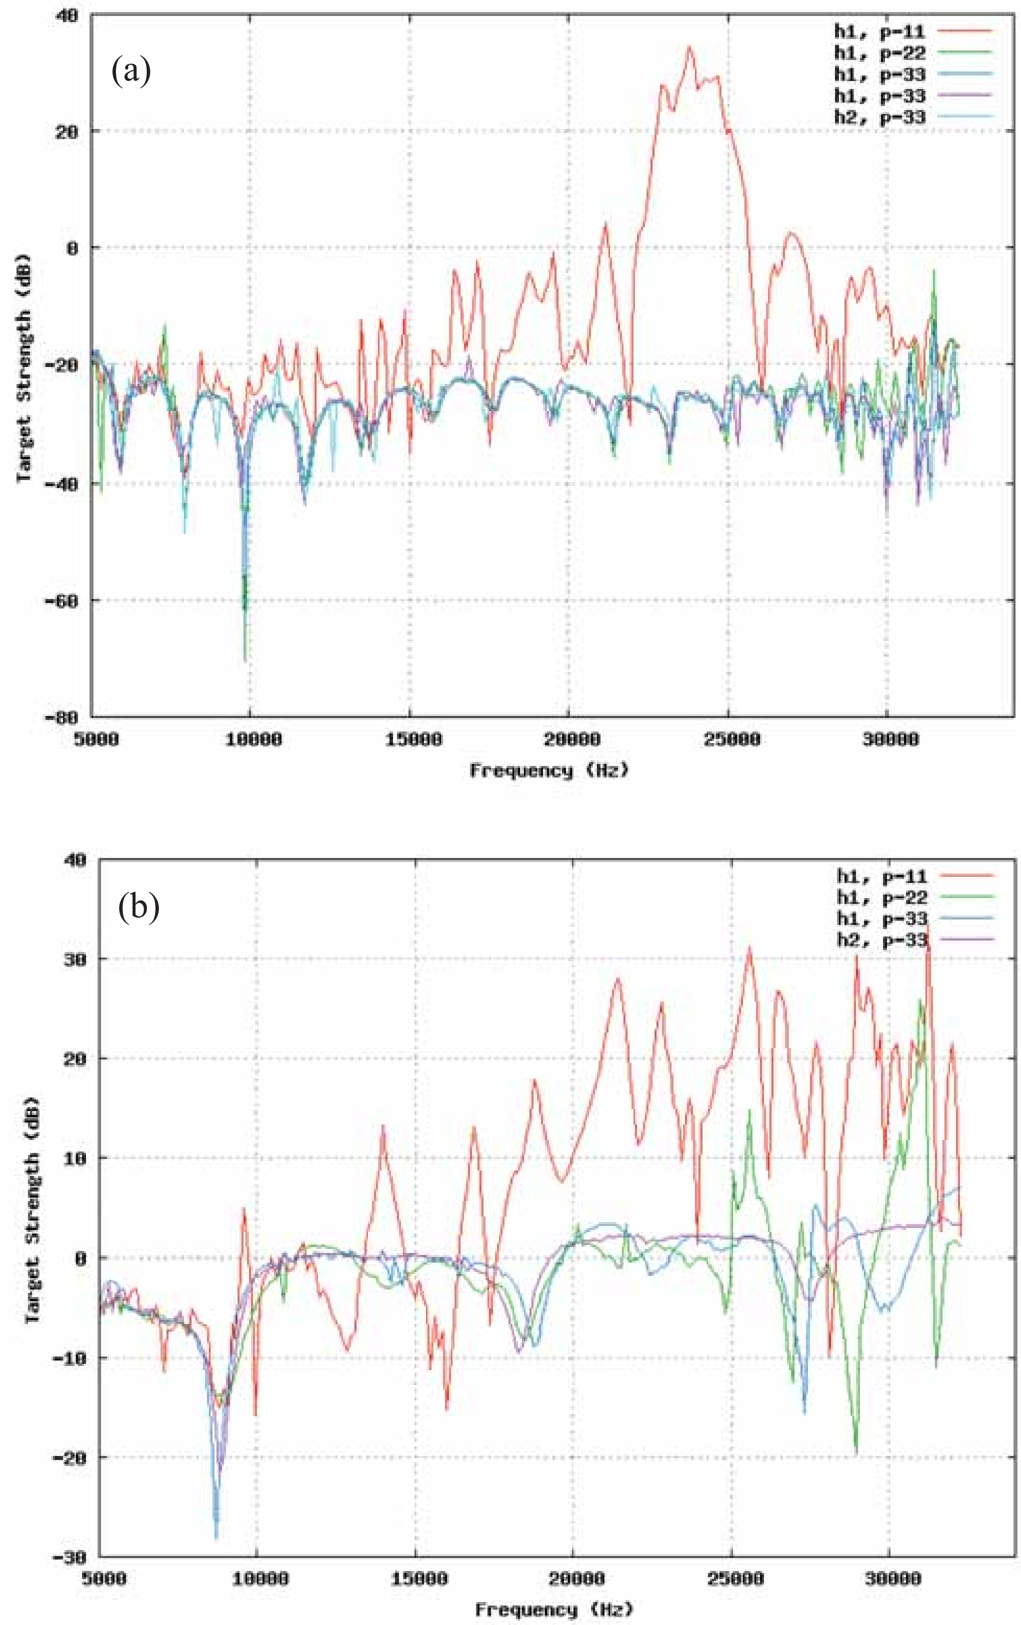

Figure 2. Effect of the hp-refinement in the scattering (monostatic response) from an E50 shell at (a) θ = 0° and (b) θ = 90°. Notice the convergence of the solution as higher order basis functions (indicated by the higher numbered p) are used

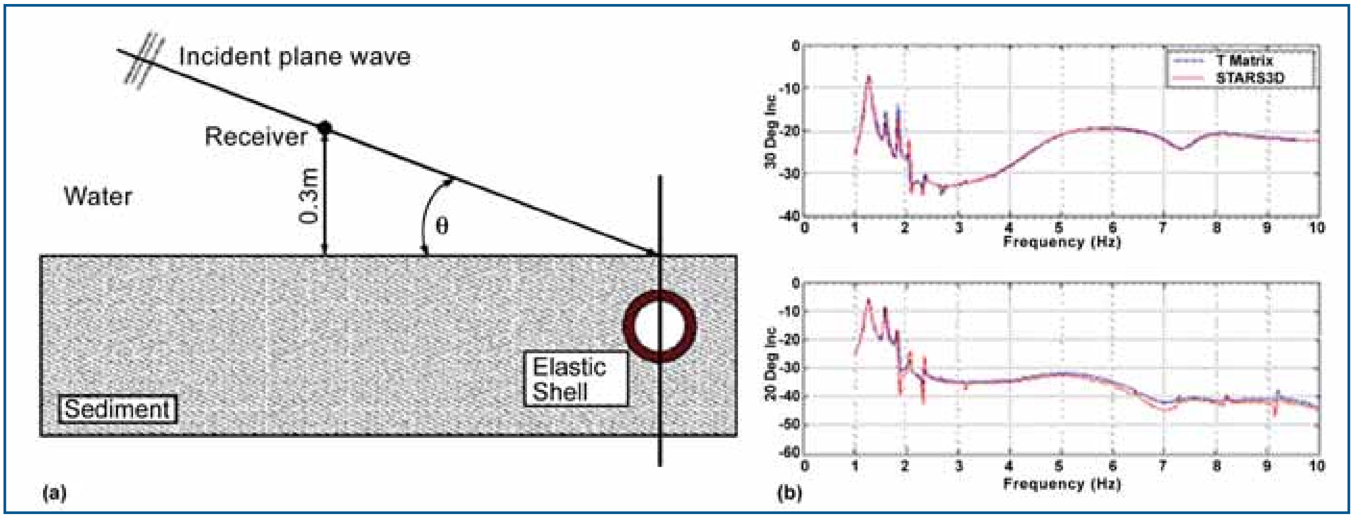

Figure 3. Scattering from an elastic spherical shell in sediment: (a) schematic description of the problem, and (b) comparison between STARS3D computations against the T-matrix method. The monostatic response is shown for 30 and 20 degree incident

proper modeling and details of the scattering environment. A critical and distinguishing feature of STARS3D is the capability to examine three-dimensional (3-D) models of mines, complex sediment layers, geometric structure of the interface between the sediment and the water, and the sea surface itself. Although this littoral example uses about 15,000 elements, evaluating large-scale environments for long-range mine hunting will ultimately require more than 107 degrees of freedom, and it is presently believed that STARS3D is the only numerical modeling technology that can be scaled to problems of this size. Such data-intensive problems are suited for the DSRC resources, and these resources will give researchers the freedom and capability to examine and evaluate realistic environments in the development of mine hunting technologies for the littoral environment. In seismic studies, STARS3D scaled well up to 64 processors during tests on MJM, and it is believed that such scaling can also be achieved with the littoral environment. STARS3D is currently available for any user on MJM.

Users Supported

This work supports Dr. Brian Houston, NRL-DC, who was interested in examining whether STARS3D modeling capabilities can be applied to the littoral environment.

DoD Impact

This work demonstrates that the STARS3D capabilities that have been developed for the seismic problem have been successfully adapted to the shallow-water mine hunting problem. The adoption of this new modeling capability is a central component of the development effort now underway with the U.S. Navy to expand this promising mine hunting approach to longer ranges and more challenging environments. NRL considers this numerical modeling technology crucial for extending high performance mine hunting to longer ranges and thus enhancing mine clearing rates.

Acknowledgments

The authors thank the DoD High Performance Computing Modernization Program for the computer time at the ARL DSRC through the PET program.

Computational Stability and Control Analysis of Fighter Aircraft

By John P. Dean, James D. Clifton, David J. Bodkin, Scott A. Morton, and David R. McDaniel, U.S. Air Force SEEK EAGLE Office

Computational Stability and Control Analysis seeks to develop a computational method for accurately determining static and dynamic stability and control characteristics of fighter aircraft with a wide range of store configurations. Studies have been focused on the USAF F-16C Falcon with and without stores. Multi-axis computational training maneuvers for use with system identification, and prescribed motion flight test maneuvers using flight test data have been investigated. Static, time-accurate computational fluid dynamics (CFD) simulations of the F-22 have also been performed.

Project Overview

Practically every fighter program since 1960 has had costly nonlinear aerodynamic or fluid-structure interaction issues discovered in flight test. The main reason for these "failures" is that the predictive methods used were not able to reveal the onset and nature of the problems early in the design phase [1]. To keep the budget overshoot under control, fixes tend to be *ad hoc* and are applied without a sound basis of fundamental understanding of the physics concerned. Using advanced CFD methods, simulations are now capable of capturing the unsteady nonlinear aerodynamic behavior that leads to the various static and dynamic instabilities associated with highly maneuverable aircraft. A new approach being investigated is system identification (SID) of CFD simulations. SID is the process of constructing a mathematical model from input and output data for a system under testing, and characterizing the system uncertainties and measurement noises [2]. The mathematical model structure can take various forms depending upon the intended use. SID is usually applied to wind-tunnel and flight-test data to obtain accurate and comprehensive mathematical models of aircraft aerodynamics, for aircraft flight simulation, control system design and evaluation, and dynamic analysis. Aircraft system identification can be used in cooperative approaches with CFD to take advantage of the strength of both approaches or having one approach fill in the gaps where the other cannot be used effectively [3]. The wide range of SID tools that have been developed for aircraft system identification can easily be used to analyze CFD data computed for aircraft in prescribed motion. An example of a SID model for lift coefficient for a clean F-16 is shown in Equation 1 and was derived from a computational training maneuver flown in CFD. Figure 1 compares the SID model prediction against flight-test data and static, time-accurate CFD simulations.

$$\begin{aligned} C_L(\alpha,\dot{\alpha},p,\dot{p},q,\dot{q}) &= C_1 + C_2\alpha + C_3\dot{\alpha} + C_4\alpha^2 + \\ & C_5 q + C_6\alpha^2\dot{\alpha} + C_7 p\dot{p} + C_8\dot{q}^2 + C_9\alpha p + C_{10}p\dot{q} + \\ & C_{11}\dot{q} + C_{12}p + C_{13}p^2\dot{p} + C_{14}p^2 + C_{15}\alpha\dot{q} + \\ & C_{16}\alpha^2 + C_{17}\alpha p^2 + C_{18}p^3 + C_{19}\dot{\alpha}q + C_{20}q^2 \end{aligned} \quad (1)$$

The ultimate desire is to generate efficient yet accurate nonlinear aerodynamic models capable of predicting force and moment coefficients for both conventional static analysis and flight-test maneuvers. Figure 5 compares wind-tunnel data against static, time-accurate CFD simulations and SID models for lift and drag for an F-16 configured with the GBU-39 Small Diameter Bomb (SDB), tanks, and missiles as shown in Figure 3. Figure 4 depicts an F-16 with the SDB performing a 3G wind-up-turn in CFD using flight test data. The image depicts isosurfaces of vorticity colored by

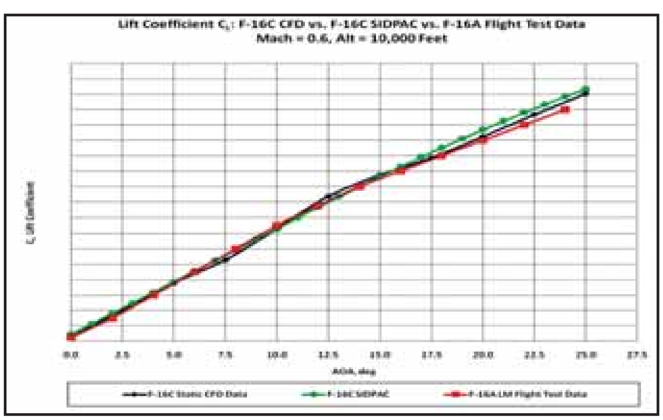

Figure 1. Comparative plots of lift coefficient for the F-16 using flight-test data and CFD simulations

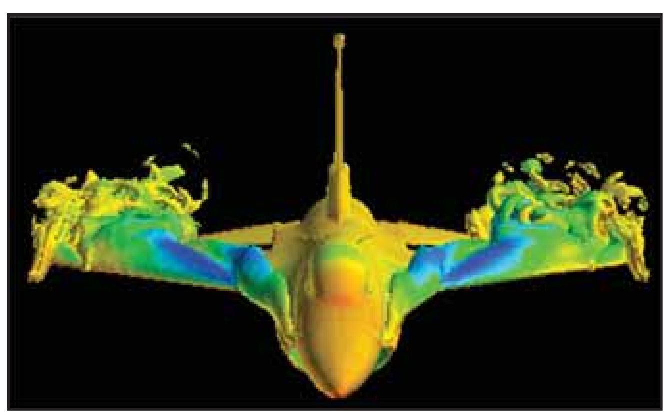

Figure 2. Image from a computational training maneuver for use with SID. Conditions are Mach 0.6, 10,000 ft, and 0 to 25 deg AOA

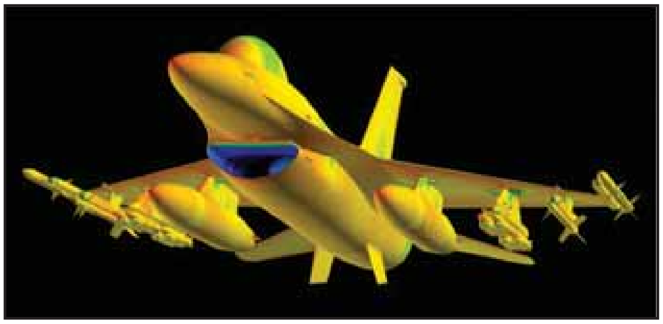

Figure 3. F-16 SDB flight-test configuration

pressure. The complete time history for lift and drag over the course of the maneuver is shown in Figure 6, comparing SID model predictions against calculated CFD values.

Conclusions and Future Work

Computational methods for accurately determining static and dynamic stability and control characteristics of fighter aircraft with various weapons configurations are being explored. These methods are being successfully applied to the F-16C and the F-22. The results of the simulations and the proposed analysis process are showing promising results and will result in improved stability and control model building times over the traditional wind-tunnel-generated-database approach as well as flexibility when encountering new weapons configurations needed by the warfighter. Favorable comparisons with flight-test and wind-tunnel data validate the outlined approach and are the subject of ongoing work. Methods for prescribed motion-flight-test-maneuver simulation have been developed and have been shown to be robust for a range of maneuvers, from wind-up-turns to high-angle-of-attack-departure testing. Integration of moving control surfaces in the flow solver with the 6DOF equations in the simulation loop is near completion. Developing "computational maneuvers" with moving controls, performing flight-test maneuvers with moving controls using flight-test data, and comparing maneuver response between CFD and flight test for the same maneuver are the subject of ongoing work by the Air Force SEEK EAGLE Office Stability & Control group.

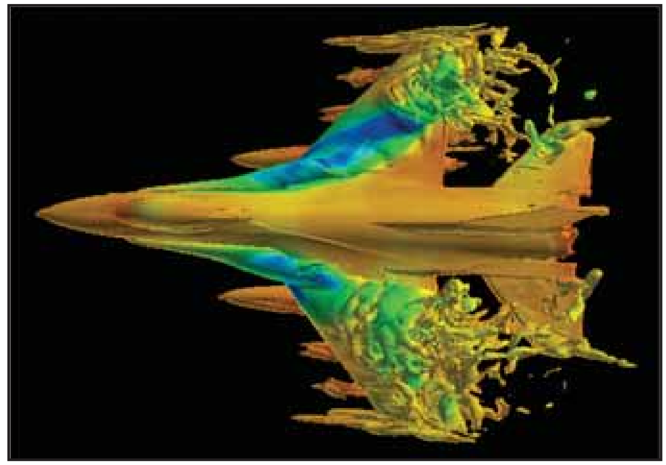

Figure 4. Image from an F-16 SDB flight test wind-up-turn-prescribed-motion simulation

References

[1] Dean, J. P., Morton, S. A., McDaniel, D. R., Clifton, J. D., Bodkin, D. J., "Aircraft Stability and Control Characteristics Determined by System Identification of CFD Simulations," AIAA Guidance, Navigation and Control Conference, Aug 2008.

[2] L. Ljung, *System Identification. Theory for the User,* Prentice Hall Information and System Sciences Series.

[3] E. A. Morelli and V. Klein, "Application of System Identification to Aircraft at NASA Langley Research Center," *Journal of Aircraft,* Vol. 42, No. 1, January-February 2005, pp. 12-25.

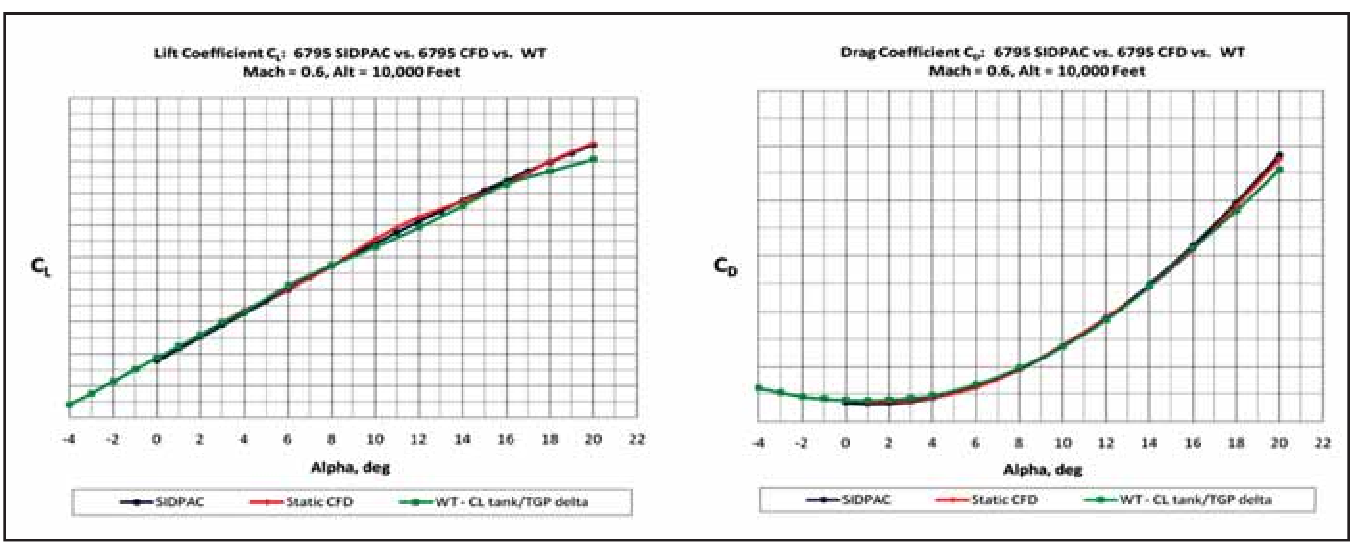

Figure 5. C_L and C_D results for the F16C SDB flight-test configuration. Plots are of a SID-model prediction, CFD-static results, and wind-tunnel results

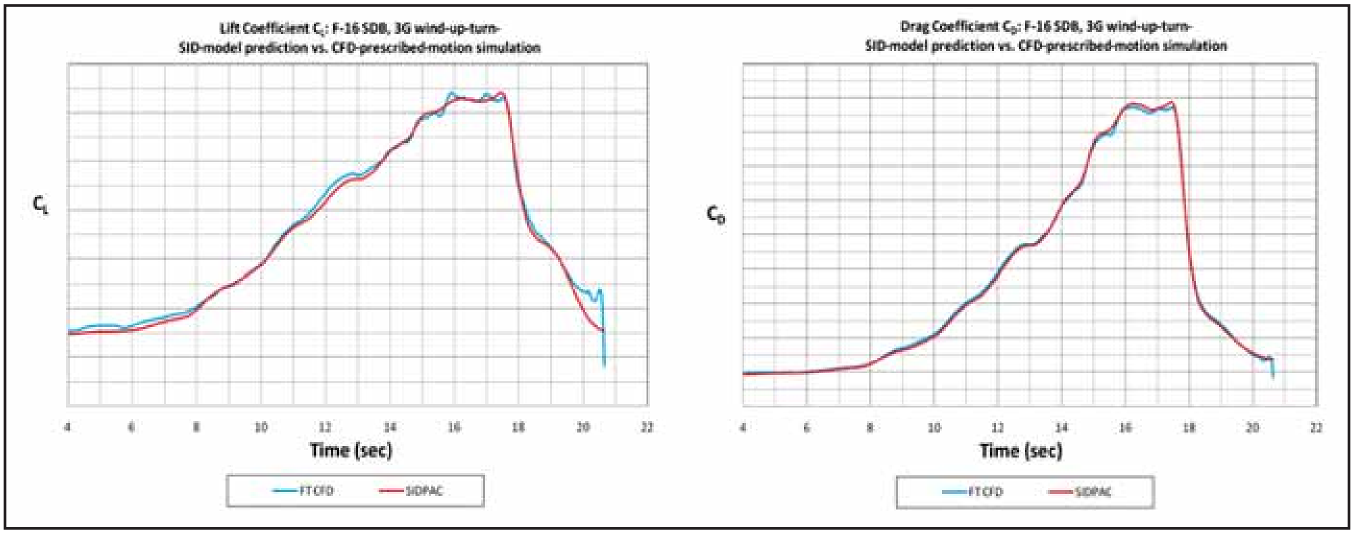

Figure 6. C_L and C_D results for a prescribed-motion 3G wind-up-turn in CFD and the corresponding SID-model prediction generated from a 10-second multi-axis computational training maneuver in CFD. Full-scale, static, time-accurate analysis of the F-22 in the clean aircraft configuration is also underway. Shown in Figure 7 is a static, time-accurate simulation depicting isosurfaces of vorticity colored by pressure at Mach 0.9

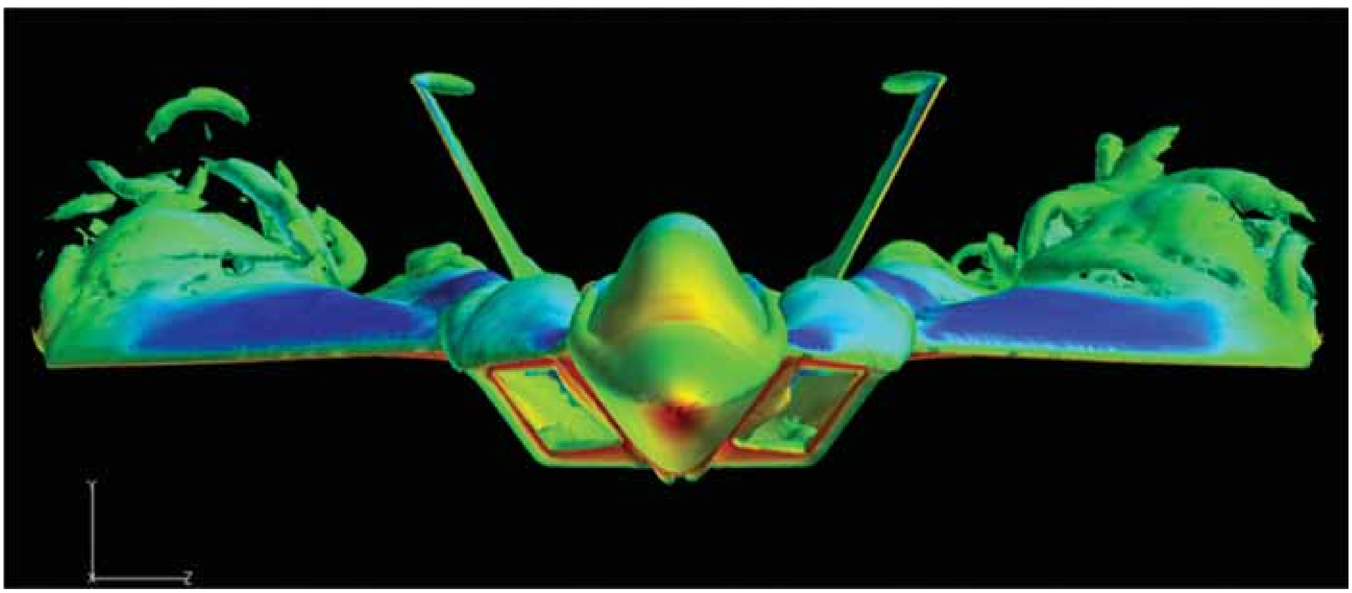

Figure 7. Image of a static, time-accurate simulation at Mach 0.9

Author and Contact: John P. Dean
Authors: John P. Dean, James D. Clifton, David J. Bodkin, Scott A. Morton, Dave McDaniel
Organization: USAF SEEK EAGLE Office, 205 West D. Avenue, Suite 348, Eglin AFB, FL, 32542
Resources: 256-1024 processors per job on *Jaws* at MHPCC
Sponsorship: DoD HPC/AF SEEK EAGLE Office Institute for High Performance Computing Applications of Air Armament (IHAAA), DoD High Performance Computing Modernization Program (HPCMP)

A 2030 Counterterrorism Vignette

*By Lee Higbie**

Even though UAVs have given him the target location, Joe Warfighter settles at the controls of his old weapon. His communication link and flack suit are new, but the steel is old—the barrel of his gun was used in the Middle East years ago. His commander sees everything he does; she can change orders or priorities in real time. His body armor is the new light-weight, super-protection unit.

Joe leans back and takes a bite of an MRG and checks the screen on his gun, looking for targets. He is startled from the lonely relaxing by his commander's voice, "You have a target at...." He punches the coordinates into his gun as she speaks. "The camo truck is Terry Terrorist's. Take it out."

He studies the screen's new view. "Got him, ma'am."

"No collateral damage. Only the truck, not the kids."

"Yes, ma'am."

"Joe. Do you see the kids there?"

"Yes, ma'am. They'll all be around the corner in a minute. I'll get him then." Joe remains riveted by his screen, waiting for the children to be safely behind the building from his target, but worrying that the target, which is creeping toward a large tree, will be under cover and out of an area of safe attack.

When a grenade explodes near him, he reflexively fires and sees his buddy at the periphery of his vision covering him as she takes out the grenade launcher. He knows his body armor provides protection only dreamt of in the last century, but is unaware of the design and testing that made it possible. The screen that rivets him shows the truck exploding and one of the children falling down. "Oh my God. Why'd I fire early? Why was my finger on the button?

Thank God, the child jumps up and runs away. The blast scared, but did not injure him. "Mission accomplished. No collateral damage," he says to his Lieutenant.

What is the connection between these imagined future capabilities? They are all being developed with modeling supported by the Arctic Region Supercomputing Center DoD Supercomputing Resource Center (ARSC DSRC).

Dr. Jerry Bernholc's modeling of ferroelectric polymers and composites is leading to high-power-density capacitors, one of the fundamental components of power conditioning systems, and high-density electronic circuits. Understanding carbon nanotubes, graphene, and molecule-based devices will lead to low-power, ultradense memories, high-gain amplifiers, and high-speed switches, all required for sensors and long-distance communication.

Dr. Emily Carter's atomic-level modeling of gun barrel materials is yielding new insights into the processes of

* Arctic Region Supercomputing Center DoD Supercomputing Resource Center, retired, and previously retired from U.S. Army Engineer Research and Development Center DoD Supercomputing Resource Center.

SUCCESS

erosion and chipping that would have shortened the life of Joe's weapon. The increased durability of artillery barrels means the accuracy and range of Joe's gun have not been eroded by time and use. Her atomic-level modeling of armor plate materials will enhance future warfighter safety.

Dr. Eric Fahrenthold's supercomputer work at the ARSC DSRC has developed the first computational models of fabrics that may be used in the next generation of body armor, such as Kevlar treated with special fluids to improve protection. Working with visualization specialists, he has produced high-resolution movies of impacts on body armor materials. These visualizations make it possible to study important details of the projectile-fabric interaction process that are difficult to observe in experiments. The goal of his research is to use high performance computing tools to complement experiments, which will accelerate the development of advanced body armor.

The two figures below are from Dr. Carter's modeling. The first figure shows a small piece of iron oxide as it undergoes shear deformation and slip. The second figure shows the atoms of an aluminum plate undergoing plastic deformation, forming partial dislocation line defects that move through the crystal lattice. Both figures are derived from approximate solutions of the Schrödinger equation, one of the most basic descriptions of matter.

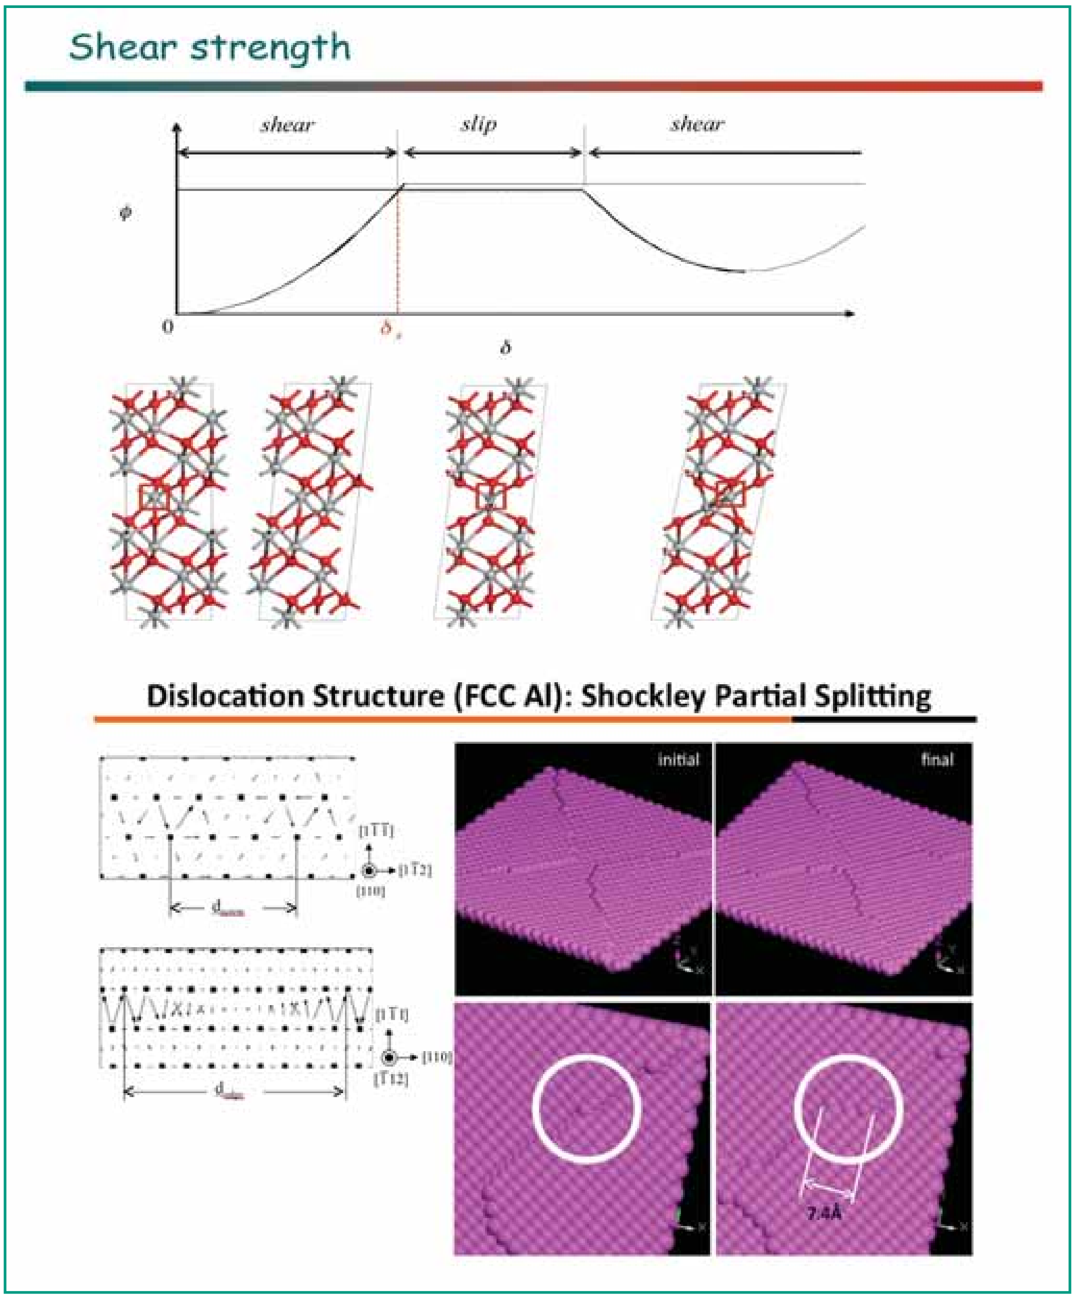

The five graphics below are from Dr. Fahrenthold's movies. All of these movies are from models of high or hypervelocity (400 m/s to 6.7 Km/s) impacts on typical armor materials. The first shows the beginning of the impact on a single layer of fabric. The second shows an intermediate stage when a projectile has penetrated the first two components of a piece of armor. The third shows the result of fabric penetration, both as modeled and as seen in the lab.

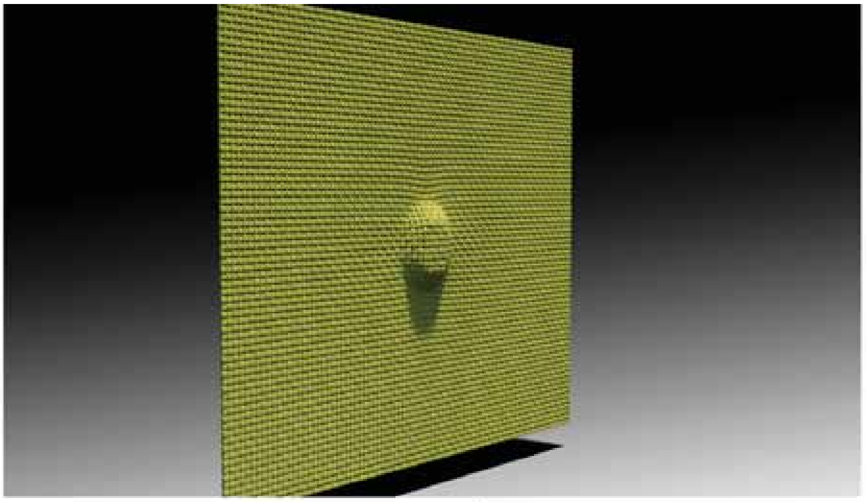

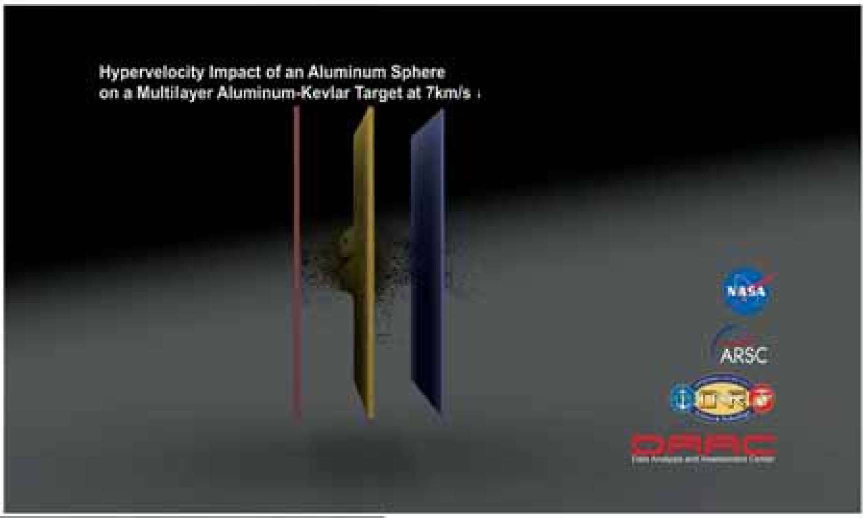

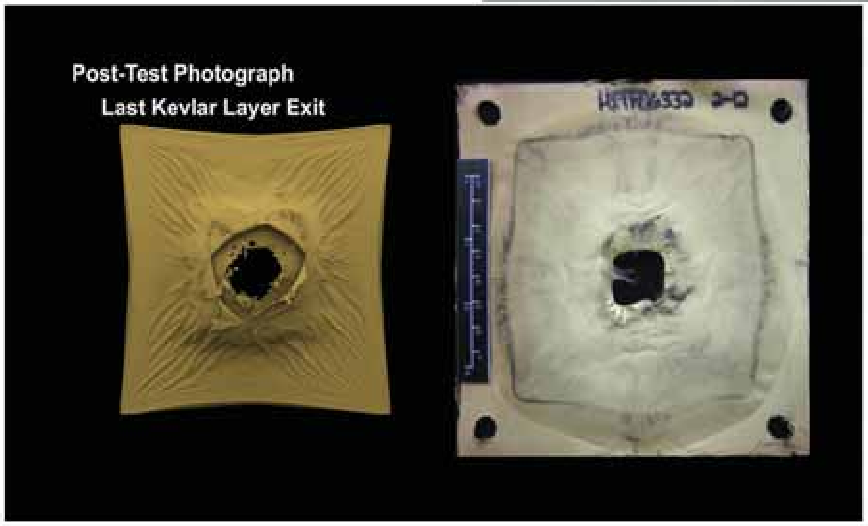

The last two frames from one of Dr. Fahrenthold's movies are of the impact of a rod on a plate at a 75° angle. In the first of the pair of frames, a small portion of the rod has fractured from the rest and is ricocheting. In the last frame, most of the rod has passed through the plate, and the fragment is nearing the edge of the plate.

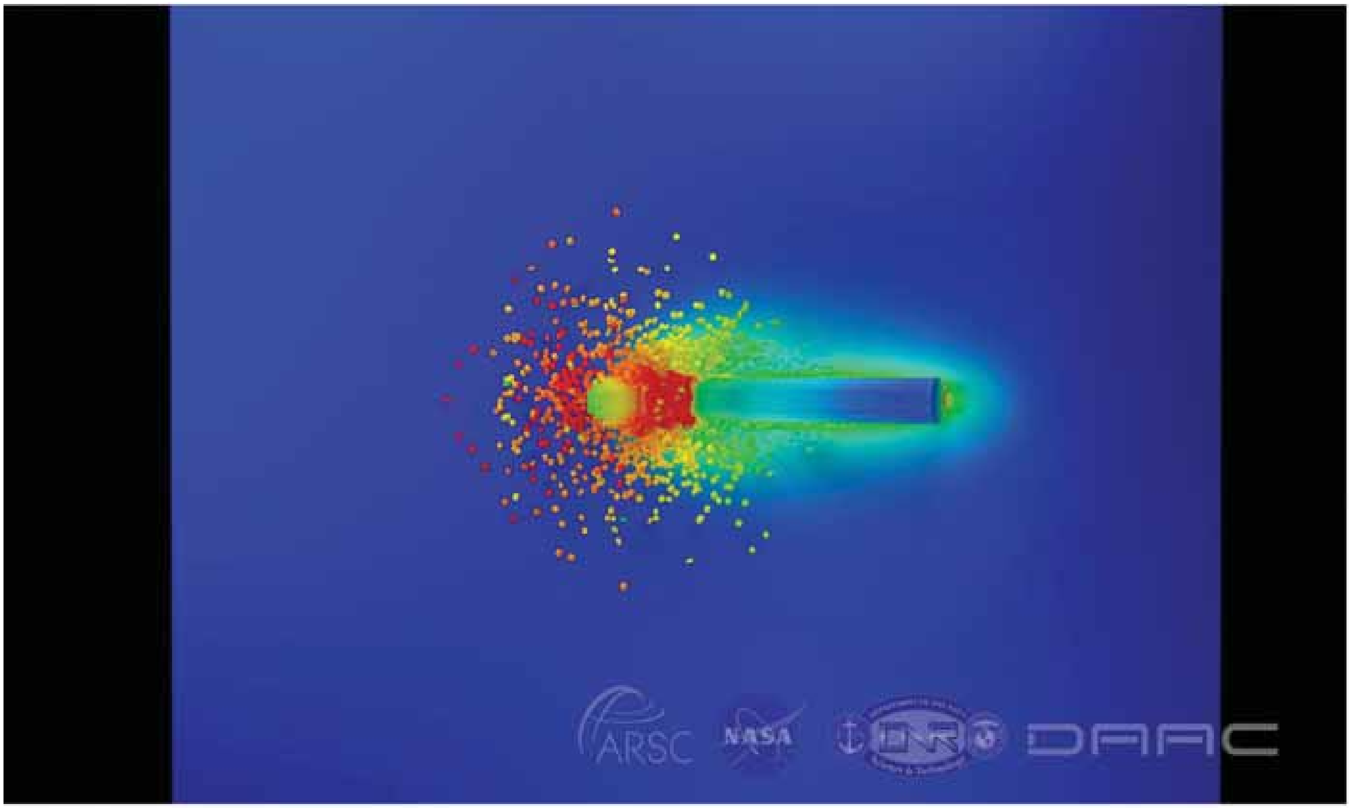

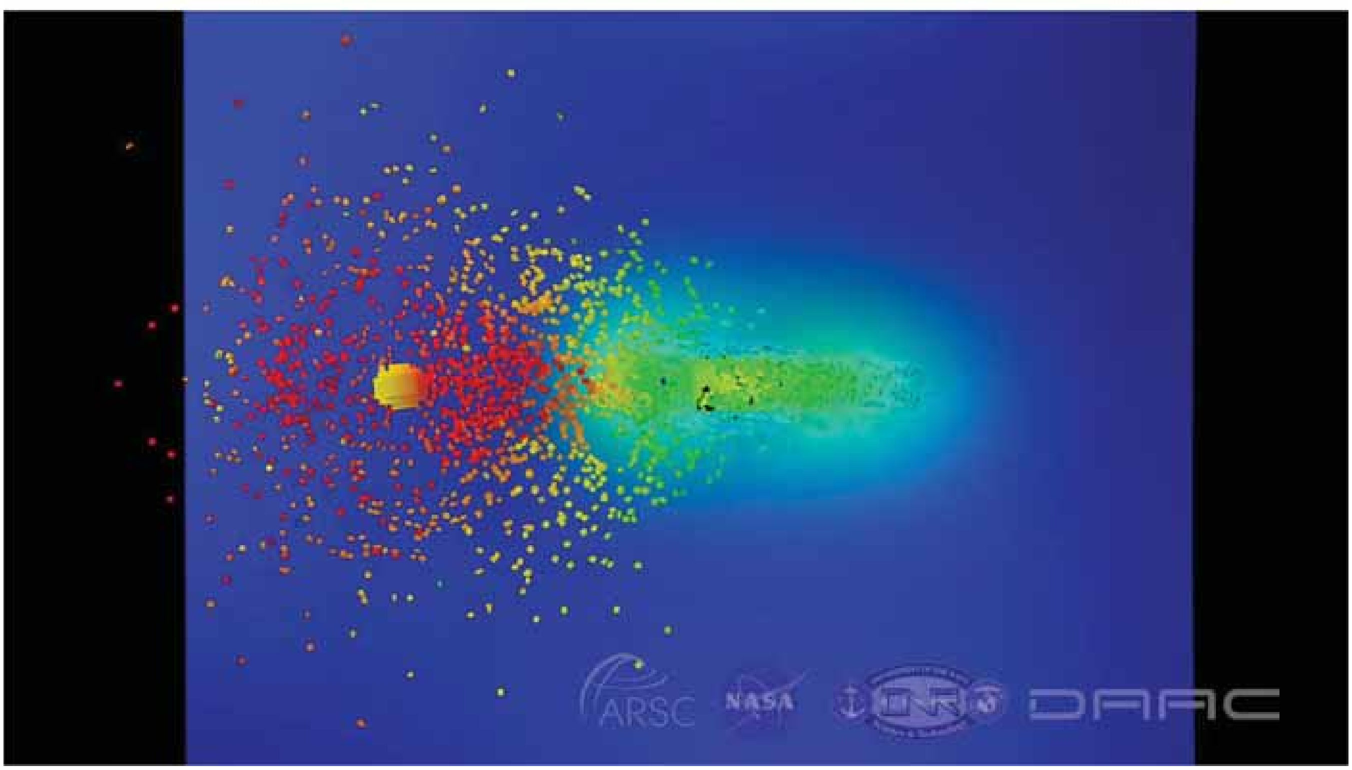

Numerical Modeling of Turbulent, Parallel, Round Jets

By Dr. Jeffrey B. Allen, U.S. Army Engineer Research and Development (ERDC) Information Technology Laboratory, and Dr. David L. Smith, ERDC Environmental Laboratory

Introduction

Because of their widespread use in a variety of engineering applications, single and offset wall jets have been extensively studied. Some of these applications include burners, boilers, film-cooling of lining walls within gas turbine combustors, fuel-injection systems, and heating and air conditioning systems [1]. Studies of twin parallel jets, however, are less common, but are still encountered in numerous engineering applications.

Some previous experimental studies involving the interaction of twin parallel jets include the inaugural work of Tanaka [2, 3], who identified the basic flow patterns and entrainment mechanisms. Specifically, he identified three primary regions along the streamwise direction consisting of the converging, merging, and combined regions (see Figure 1). Elbena et al. [4] investigated turbulent mixing of the parallel jet and found that the velocity profile downstream of the combined region matched that for a single jet. Lin and Sheu [5, 6] showed that the mean velocity is self-similar in both the merging and combined regions, but that the Reynolds shear stresses approached self-similarity only in the combined region.

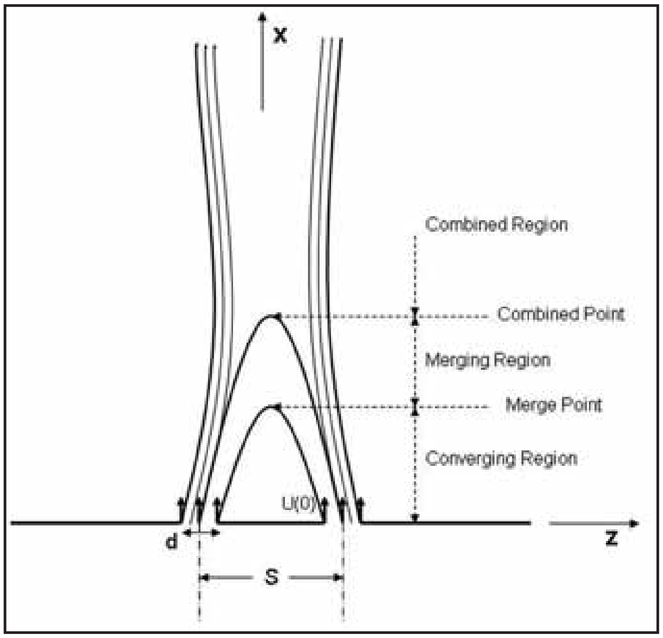

Figure 1. Parallel jet flow field regions and coordinate system

Previous numerical investigations of turbulent parallel jets are relatively few in number, but do include the work of Spall et al. [1], who investigated several planar parallel jet configurations using both the differential Reynolds stress (RSM) and standard $k - \varepsilon$ turbulence closure models. Their results showed good agreement with hot wire anemometry experiments, particularly with respect to the location of the combined and merge points.

The objective of the present work is to evaluate the effectiveness of the Finite Volume Method (FVM) using the standard $k - \varepsilon$ approach to predict the three-dimensional (3-D) evolution of twin, isothermal, turbulent, round jets at a single jet spacing ($S/d = 4$). The exit plane Reynolds number based on a nozzle exit velocity U_e (≈40 m/sec) and a jet diameter (d) of 0.01 m will remain constant at approximately 25000. Computations pertaining to stream-turbulence intensity, as well as streamwise distances to the combined merge points, are conducted and compared with the experimental data conducted by Harima et al. [7]. Qualitative comparisons between the $k - \varepsilon$ and the Large Eddy Simulation (LES) methods are also conducted. All of the numerical simulations are carried out using the finite-volume code Fluent Version 6.3 [8].

Numerical Procedure

The computational domain (see Figures 2 and 3) was defined by a right, circular frustum with lower (x/d = 0) and upper (x/d = 150d) base diameters of 50d and 100d, respectively. At the lower base, two jet openings, each of diameter d, were symmetrically positioned along the ± z axis and separated at distance s between each jet centerline. The domain was

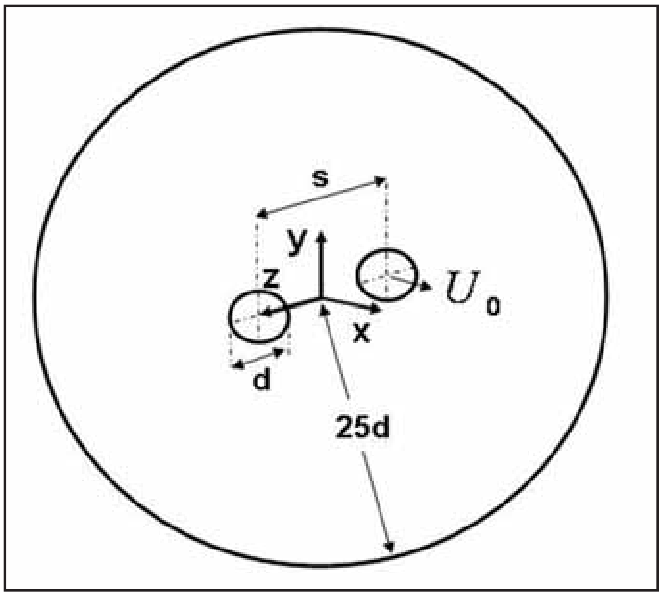

Figure 2. Lower base geometry (x/d = 0) and inlet circular jet configuration including location of reference coordinate system

bounded on the x/d = 0 plane (excluding the jet inlets) by a no-slip, solid wall and bounded along the x/d = 150 plane and lateral frustum envelope by constant (atmospheric) pressure boundaries. The inlet jets were initially assigned uniform, constant velocities of 40 m/s, contributing to a Reynolds number (based on a single jet diameter d and jet exit velocity U_e) of approximately 25000. The selected dimensions for this study were considered large enough to allow adequate evolution of the jet in the self-preserving region and provide sufficient clearance between the lateral jet envelope and the lateral domain boundaries. This was substantiated from several preliminary $k-\varepsilon$ simulations.

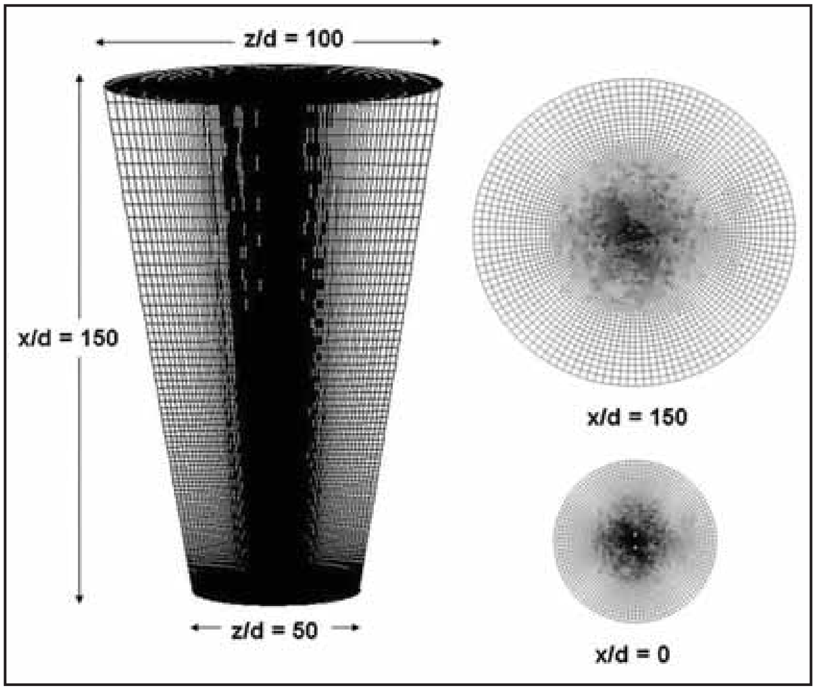

Figure 3. Representative hexagonal mesh (1.517E6 elements) corresponding to jet spacing, S/d = 4

The mesh (see Figure 3) consisted of 1.517E6 hexahedral elements consisting of a minimum cell edge distance (direction perpendicular to the wall) of 2.67E-4d and a maximum edge distance of 0.8d. In the y-direction, beginning with the first mesh node placement from the wall boundary, a grading was employed corresponding to a successive distance ratio of 0.5 percent. Although a mesh independent study was not conducted for the current LES simulation (chiefly because of the excessive time commitments needed to run a single case), several RANS simulations were conducted utilizing a variety of progressive meshes with increasing levels of fidelity. The resulting mesh comprised of the 1.517E6 elements was deemed adequate for the given domain size when using computed values for turbulence intensity for comparison purposes. The integration time-step resulting from 4 percent of the characteristic length (d) to the inlet jet velocity ratio was 1.0E-5 s. This time-step was utilized in lieu of a time-step based on the experimental sampling frequencies (10 kHz) [7], as the latter, less conservative time-step resulted in numerical instabilities and non-converging residual values.

k-epsilon Procedure

The foregoing provides the procedures involved in performing the RANS $k-\varepsilon$ turbulence closure model. Given the experimental turbulence intensity (T_i = 3.6%)) at the inlet [7], the turbulent kinetic energy (k) and dissipation rate (ε) were computed from the familiar, yet approximate, relations:

$$k = 1.5(U_i T_i)^2 \qquad (0.1)$$

$$\varepsilon = C_\mu^{0.75} k^{1.5} / L \qquad (0.2)$$

where C_μ = 0.09, and the turbulence length scale (L) was taken as 0.07d, (the factor 0.07 is based on the maximum mixing length in a turbulent pipe flow). The $k-\varepsilon$ empirical constants were taken as $C_{\tau 1}$ = 1.44, $C_{\tau 2}$ = 1.92, $C_{\tau 3}$ = 1.0, σ_k = 1.0, and σ_τ = 1.3. The pressure-based, finite-volume code Fluent Version 6.3 [8] was utilized to solve the 3-D conservation equations pertaining to mass, momentum, and transport of turbulent kinetic energy and dissipation. Interpolation to cell faces for the convection terms was performed using the second-order upwind, discretization scheme, while second-order central differences were utilized for the viscous terms. Pressure-velocity coupling was based on the Simi-Implicit Method for Pressure-Linked Equations (SIMPLE) procedure [9].

The simulations were run in the parallelized, finite volume Fluent solver [8] using an average of 16 cores (1.6 GHz, 5.68 GFLOPS), on the Air Force Research Laboratory DoD Supercomputing Resource Center (AFRL DSRC) SGI Altix 3700. Solutions at each time-step were considered converged when residuals for each of the equations (based on the L2 norm) were reduced by a minimum of four to five orders of magnitude.

LES Procedure

Initialized from the fully converged RANS $k-\varepsilon$ solution, fluctuating, turbulent velocities were superimposed at the inlet, velocity boundary (being necessarily perpendicular to the streamwise velocity) using the vortex method [8]. The vortex method is based on the Lagrangian form of the vorticity equation and Biot-Savart law, and represents each Lagrangian particle (or vortex point) by the circulation $\Gamma_i(x,y,k)$ and spatial distribution $\eta(\vec{x},k,\varepsilon)$.

$$\Gamma_i(x,y) = 4\sqrt{\frac{\pi A k(x,y)}{3N[2\ln(3) - 3\ln(2)]}} \qquad (0.3)$$

$$\eta(\vec{x}) = \frac{1}{2\pi\sigma^2}\left(2e^{-|x|^2/2\sigma^2} - 1\right)2e^{-|x|^2/2\sigma^2} \qquad (0.4)$$

Thereafter, the LES method was subsequently run utilizing the parallelized Fluent solver [8] on the AFRL SGI Altix 4700 shared-memory system over approximately 2000 core hours, utilizing an average of 16, 1.6 GHz (6.4 GFLOPS) Itanium 2 processors and culminating in approximately 43,600 time-steps.

Interpolation to cell faces for the convection terms was performed using the second order, bounded central differencing scheme. Pressure-velocity coupling was based on the PISO procedure [10]. Solutions at each time-step obtained using the segregated solver were considered converged when residuals for each of the equations (based on the L2 norm) were reduced by a minimum of five to six orders of magnitude. The subgrid-scale stresses resulting from the filtered equations were modeled in accordance with the dynamic subgrid model of Germano et al. [11] and Lilly [12]. This model obviates the need for a user-defined Smagorinsky model constant (C_s), since it is dynamically computed as a result of the information provided by the resolved scales of motion.

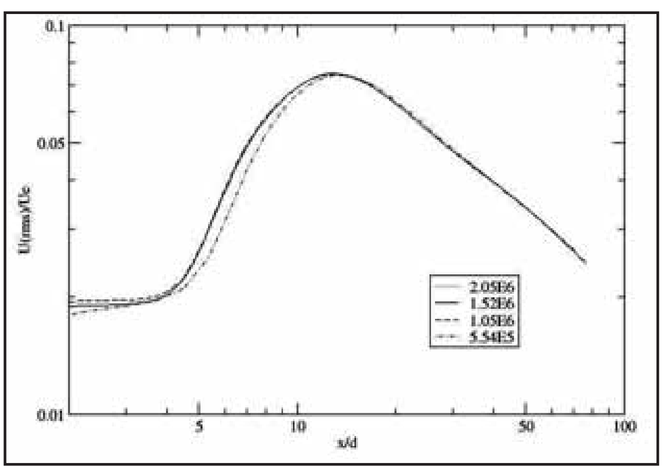

Figure 4. Sensitivity of the numerical solution to grid resolution

Results

Figure 4 shows the streamwise (y/d = z/d = 0.0) variation in turbulence intensity (U_{rms}/U_e) utilizing RANS $k-\varepsilon$ to illustrate the sensitivity of the numerical results to grid resolution. As indicated, a total of four separate grid resolutions of 5.45E5, 1.05E6, 1.52E6, and 2.05E6 cells were utilized at increasing levels of fidelity. Grid independent solutions were obtained for 1.52E6 cells.

Figure 5 shows qualitative comparisons of velocity magnitude contours between the steady-state, fully developed RANS $k-\varepsilon$ and the transient LES results. As indicated, from the RANS solution, and indeed implicit within its derivation, the effects of the Reynolds stresses on the mean flow are clearly the result of time-averaged properties. In contrast, the LES results, based on space-filtered equations of motion, allow for the computation of instantaneous flow characteristics and highly resolved turbulent flow structures.

Figure 6 shows the streamwise (y/d = z/d = 0.0) variation in turbulence intensity (U_{rms}/U_e) comparing the RANS $k-\varepsilon$ numerical results with the experimentally derived values of Harima et al. (2005). As indicated, the best agreement (within the tolerances associated with the experimental error) is found along the downstream positions aft of approximately x/d = 15. Forward of this location, the numerical predictions tend to overpredict the turbulence intensity by approximately 10 percent. The maximum turbulence intensity is also slightly overpredicted by the numerical results as compared with the experiment, with values of 14.0 versus 15.8, respectively. This maximum turbulence intensity

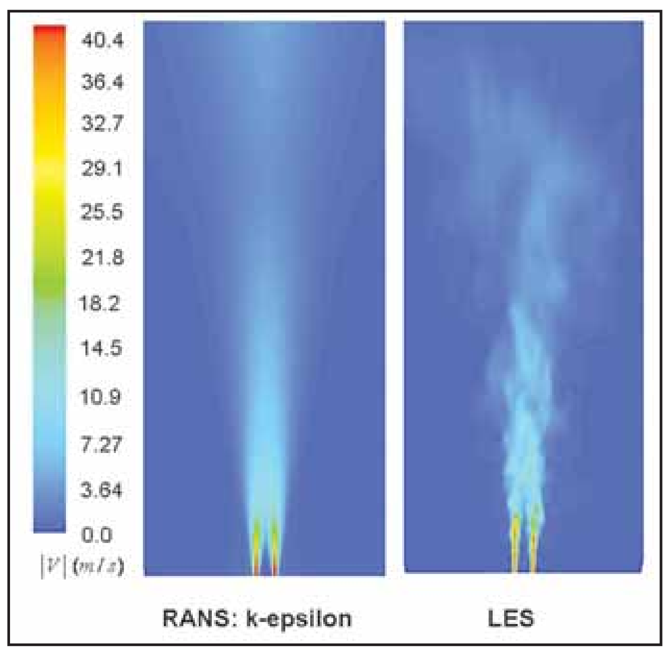

Figure 5. Contours of velocity magnitude at the y/d = 0 plane, comparing $k-\varepsilon$ and transient LES

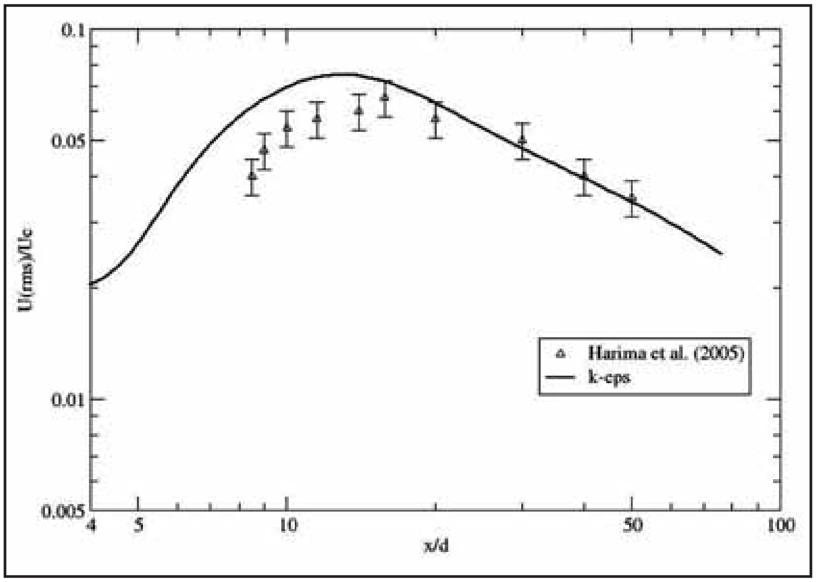

Figure 6. Experimental and numerical comparisons ($k-\varepsilon$) for the streamwise variation of turbulence intensity

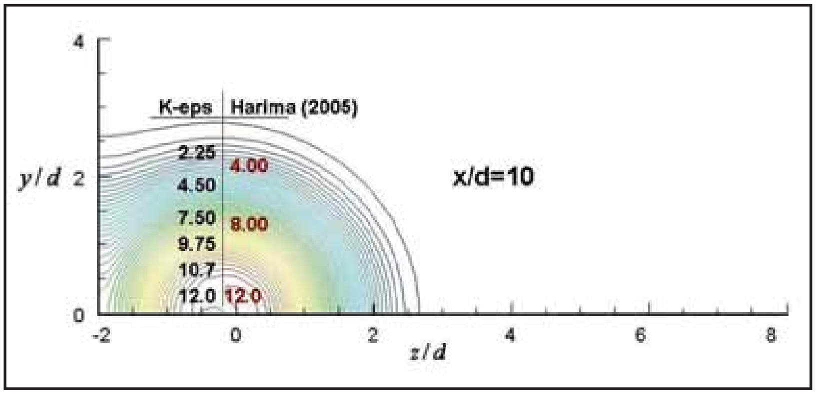

Figure 7. Turbulence intensity contour comparisons at x/d = 10

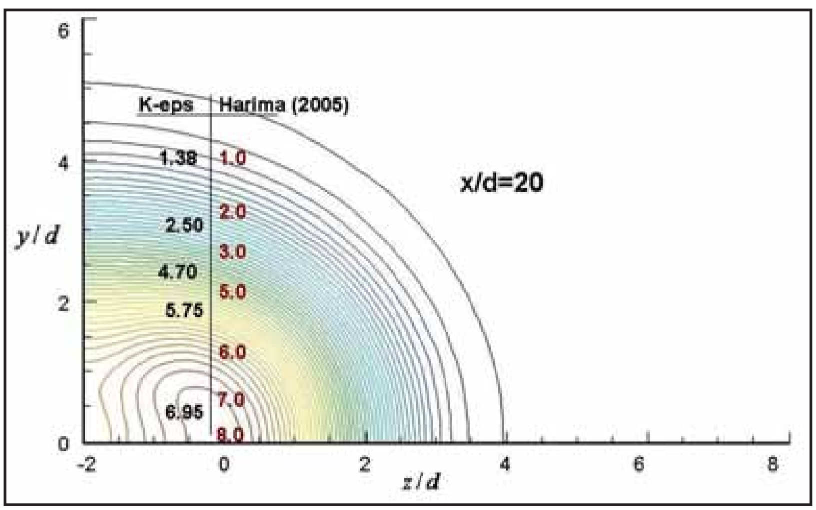

Figure 8. Turbulence intensity contour comparisons at x/d = 20

also corresponds to the maximum downstream velocity and thus defines the location of the combined point [2,3]. The area aft of the combined point corresponds to the region wherein the two jets begin to resemble a self-similar single jet.

Figures 7, 8, and 9 show contour comparisons of turbulence intensity at three streamwise planar locations (x/d = 10, x/d = 20, and x/d = 40), with $-2 \le z/d \le 6$ and $0 \le x/d \le 10$. As indicated, there is good agreement with experiments in each of these cases. Furthermore, similar to the experimental results of Harima et al. [7] and the rectangular jet experiments by Marsters et al. [13], there exists a switching tendency of the major axis between the y and z axes along the streamwise direction. The x/d = 10 planar contours, for example, reveal a major axis along the z direction, and the x/d = 20 contours swap the trend in favor of the y direction. Finally, at x/d = 40, corresponding to self-similar, single jet region, the contours reveal nearly a circular shape with no real bias for major axis distinction.

Figure 10 shows comparisons of the turbulence intensity (U_{rms}/U_e) at a total of five streamwise locations corresponding to x/d = 5, x/d = 10, x/d = 20, x/d =30, and x/d = 40 along the line bounded by $0.0 \le z/d \le 10$. This plot, like Figure 6, reemphasizes the improved agreement with experiments at locations aft of the combined point.

Conclusions

The results presented herein indicate that for a parallel jet spacing of S/d = 4, the 3-D RANS $k-\varepsilon$ turbulence model can accurately predict the evolution of the turbulence intensity downstream of the merge and combined points, but tended to overpredict upstream values from this location. Consistent with what has been identified in previous experimental findings, the numerical results also showed a tendency for the major axis (corresponding to the overall shape of the jet structure) to alternate between the two primary axes defining the plane along the steamwise direction. These latter results can be explained based on the fact that the two jets become a single, self-similar jet aft of the combined point.

References

[1] Spall, R. E., Anderson, E. A., and Allen, J. 2004. "Momentum Flux in Plane, Parallel Jets," ASME *Journal of Fluids Engineering*, Vol. 126, pp. 65-70.

[2] Tanaka, E. 1970. "The Interference of Two-Dimensional Parallel Jets (1st Report, Experiments on Dual Jet)," *Bull. JSME*, 13(56), pp. 272-280.

[3] Tanaka, E. 1974. "The Interference of Two-Dimensional Parallel Jets (2nd Report, Experiments on the Combined Flow of Dual Jet)," *Bull. JSME*, 17(109), pp. 920-927.

[4] Elbana, H., Gahin, S., and Rashed, M. I. I. 1983. "Investigation of Two Plane Parallel Jets," *AIAA J.*, 21, pp. 986-990.

[5] Lin, Y. F., and Sheu, M. J. 1990. "Investigation of Two Plane Parallel Unventilated Jets," *Exp. Fluids*, 10, pp. 17-22.

[6] Lin, Y. F., and Sheu, M. J. 1991. "Interaction of Parallel Turbulent Plane Jets," *AIAA J.*, 29, pp. 1372-1373.

[7] Harima, T., Fujita, S., and Osaka, H. 2005. "Turbulent Properties of Twin Circular Free Jets With Various Nozzle Spacing," Engineering Turbulence Modeling and Experiments 6, pp. 501-510.

[8] *Fluent 6.3 Users Guide*, Lebanon NH, 2006.

[9] Patankar, S. V. and Spalding, D. B. 1972. "A Calculation Procedure for Heat, Mass and Momentum Transfer in Three-dimensional Parabolic Flows," *Int. J. Heat Mass Transfer*, Vol. 15, p. 1787.

[10] Issa, R. I. 1986. "Solution of the Implicitly Discretized Fluid Flow Equations by Operator Splitting," *Journal of Computational Physics*, 62, pp. 40-65.

[11] Germano, M., Piomelli, U., Moin, P., and Cabot, W. H. 1991. "A dynamic Subgrid-scale eddy viscosity model," *Physics of Fluids* 3, pp.1760.

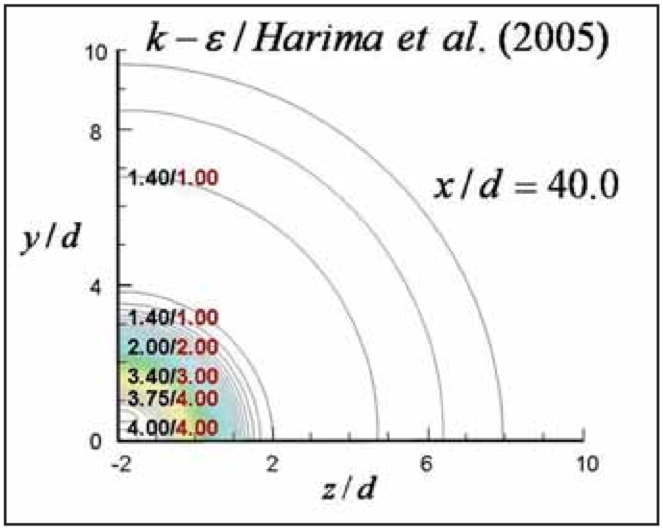

Figure 9. Turbulence intensity contour comparisons at x/d = 40

[12] Lilly, D. K. 1992. "A proposed modification of the Germano subgrid-scale closure," *Physics of Fluids* 3, 2746-2757.

[13] Marsters, G. F. and Fotheringham, J. 1980. "The influence of aspect ratio on incompressible turbulent flows from rectangular slots," *Aeronautical Quarterly*, 31:4, pp. 285-305.

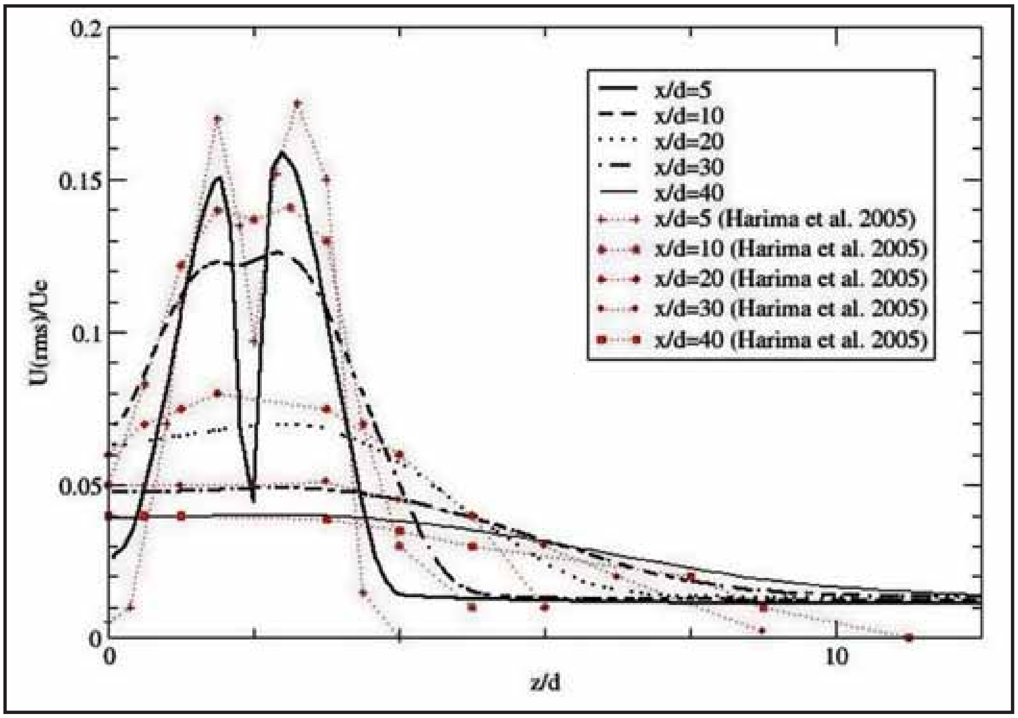

Figure 10. Streamwise variation of longitudinal turbulence intensity along the z-axis for S/d = 4

Design Studies of Turbine Blade Film Cooling with Unburned Fuel in Cross-Stream Flow

By H. Thornburg,[1] Mississippi State University; B. Sekar,[2] J. Zelina,[2] W. Anderson,[3] and M. D. Polanka,[2] Air Force Research Laboratory; C. X. Lin[4] and R .J. Holder,[5] The University of Tennessee; A. M. Briones[6] and S. D. Stouffer,[7] University of Dayton Research Institute

The demand for increased thrust in gas turbine engines, greater combustion efficiency, and reduced fuel burn has continuously driven operating temperatures and pressures higher. The increase in combustor exit temperatures is a direct result of increasing combustor fuel-air ratios, and these fuel-air ratios are approaching stoichiometric operation. In addition, combustion systems are becoming compact, resulting from requirements to reduce engine size and weight, thereby improving upon engine thrust-to-weight ratios. Reduced combustor size results in reduced residence time for combustion reactions to complete. These two factors increase the probability of unburned fuel entering the turbine and reacting with cooling or dilution air. Film cooling plays a critical role in providing effective thermal protection for components of modern gas turbine engines. However, under certain conditions, unburned fuel from the combustor can result in secondary combustion fed by the air jet employed for film cooling. Of specific concern is when this unburned fuel interacts with the cooling airflow near the metallic surfaces, resulting in secondary combustion that can cause damage of turbine blade, resulting in costly repair or engine failure. These chemical reactions of the combustor products with the film-cooling air near the vane surface can be detrimental to the turbine vane and rotor, with potential to cause catastrophic failure in a matter of milliseconds. Both experimental and computational work is necessary in order to understand the complex reacting boundary layer physics relevant to turbine vane cooling. However, knowledge of film cooling in reactive flow is limited. A literature survey indicates that nearly all of the previous studies on film cooling were conducted with pure cross-stream airflow without combustion.

Therefore, a science base is needed to understand the "burning in the turbine" phenomena and identify practical design approaches that can reduce risk of occurrence of these phenomena in future gas turbine systems. As future requirements push for compact, efficient engine designs, conventional gas turbine component design methodology may become more integrated to provide higher performance systems. Three hole configurations, namely, circular, angled circular, and fanned film-cooling holes have been tested at the combustion branch of the Air Force Research Laboratory (AFRL/RZTC) in order to understand the underlying physics of the combustion near the vane blade surface. The experimental effort for the three geometries and complementary reactive film-cooling studies of the fan-shaped film-cooling hole geometry are discussed in the papers AIAA 2009-0678[1] and AIAA 2009-0298.[2]

It is the intent of this work to compare the effect and performance of the three hole geometries on the resulting thermal boundary layer on the vane blade,

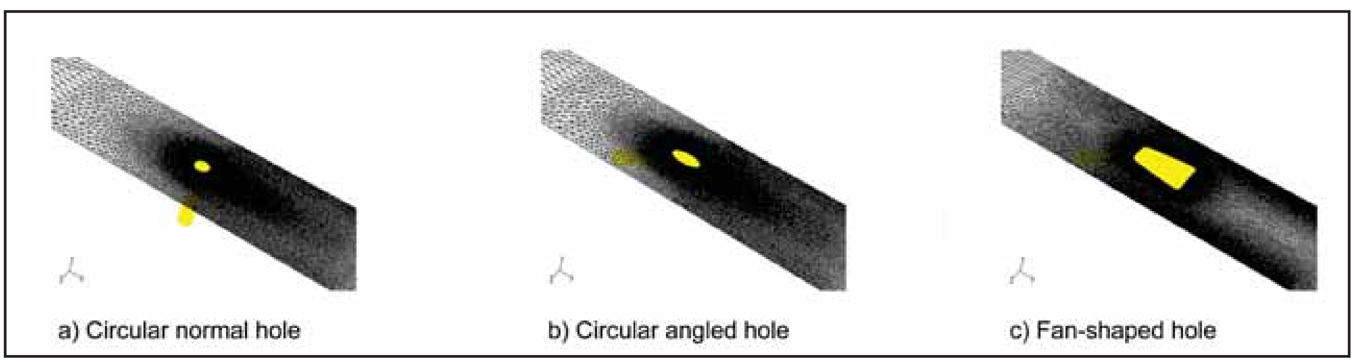

Figure 1. a-c Schematic of circular normal, angled circular, and fan-shaped hole meshes

[1] Research Professor, Mississippi State University, AIAA Senior member.
[2] Senior Research Engineer, AFRL/RZTC, 1950 Fifth St., Associate Fellow.
[3] Research Engineer, AFRL/RZTC, 1950 Fifth St., member.
[4] Associate Professor, Dept. Mech. Aerospace & Biomed. Eng., 213 Perkins Hall, The University of Tennessee.
[5] Graduate Student, Dept. Mech. Aerospace & Biomed. Eng., 213 Perkins Hall, The University of Tennessee.
[6] Research Engineer, Energy & Environmental Eng. Div., 300 College Park/0043, AIAA member, University of Dayton Research Institute.
[7] Senior Research Engineer, Energy & Environmental Eng. Div., 300 College Park/0043, AIAA Senior member, University of Dayton Research Institute.

subjected to the cross stream of combusted vitiated products, interacting with the film-cooling air. A Reynolds-Averaged Navier-Stokes (RANS) approach is employed to simulate the complex turbulent reactive flow exhibited by film-cooling flows emanating from a surface. The widely used SST k-ω turbulence model is used to model the turbulent flow. A simplified two-step, propane-air reaction scheme has been employed to model the combustion process and study the underlying physics of mixing between film cooling and cross-stream flow driving secondary combustion. The eddy-dissipation concept (EDC) approach is used to account for the turbulence-chemistry interaction. The three-dimensional geometry is modeled using a hybrid mesh. The reacting flow field and the resulting film-cooling effectiveness are predicted for circular, angled circular, and fanned-film hole geometry for two equivalence ratios, one blowing ratio, and both air and N_2 film cooling. Numerical results between air and N_2 film cooling generally agree well with experimental data in terms of relative temperature change, non-dimensionalized with respect to the N_2 film temperature. Results indicate that hole geometry plays a key role in the effectiveness of the film-cooling design. Film cooling provided by the normal circular hole is considerably lower than that provided by the angled and fanned hole for both lean and in rich conditions. Air injection feeds secondary combustion that substantially increases the wall temperature on the flat surface for a considerable distance downstream of the hole. However, the shaped hole produces a larger effective film area in the immediate vicinity of the cooling hole both axially and laterally when compared with the normal circular and angled circular configurations. For fuel-rich conditions, a distinct hot area downstream of the coolant hole generated by the secondary combustion feed by coolant air injection has been predicted. This results in negative cooling effectiveness in certain areas of the flat surface, specifically for the shaped hole. The N_2 coolant air injection provides no O_2 to feed secondary combustion for the unburned fuel exiting the combustor at high-equivalence ratios.

The cold coolant jet penetrates into the hot cross flow, where three-dimensional complex turbulent mixing occurs. Some of the coolant will be swept onto the flat surface to form a coolant layer, which is meant to protect the surface from the hot cross flow. For air injection, the cooling air will react with the chemical species in the cross flow, releasing heat that alters the thermal flow field, an addition of complex physics as compared to N_2 injection. Figure 3 shows the typical N_2 mass fraction surface contour of 0.90 value, issued from the various cooling holes.

In general, for all the cases, the coolant jet can induce a pair of counterrotating vortexes as it travels through the cross flow. For the circular hole normal injection, the film-cooling jet shoots out more vertically penetrating the cross stream, and there is little bending of the jet, compared with the angled-circular and fanned-hole cases. The fanned hole was the most effective in cooling the wall, since the jet spreads longitudinally and laterally to a greatest extent, as seen from Figure 3d, and 3e, in the immediate vicinity of the trailing edge of the film-cooling hole.

In general, with the increase of Phi, both the size of the hot area and the temperature level of the hot area increase when air is used as the coolant. The high temperature region seems to move further downstream of the film-cooling trailing edge for the angular circular hole and for the fanned hole compared with the normal circular film cooling. The increase of temperature for

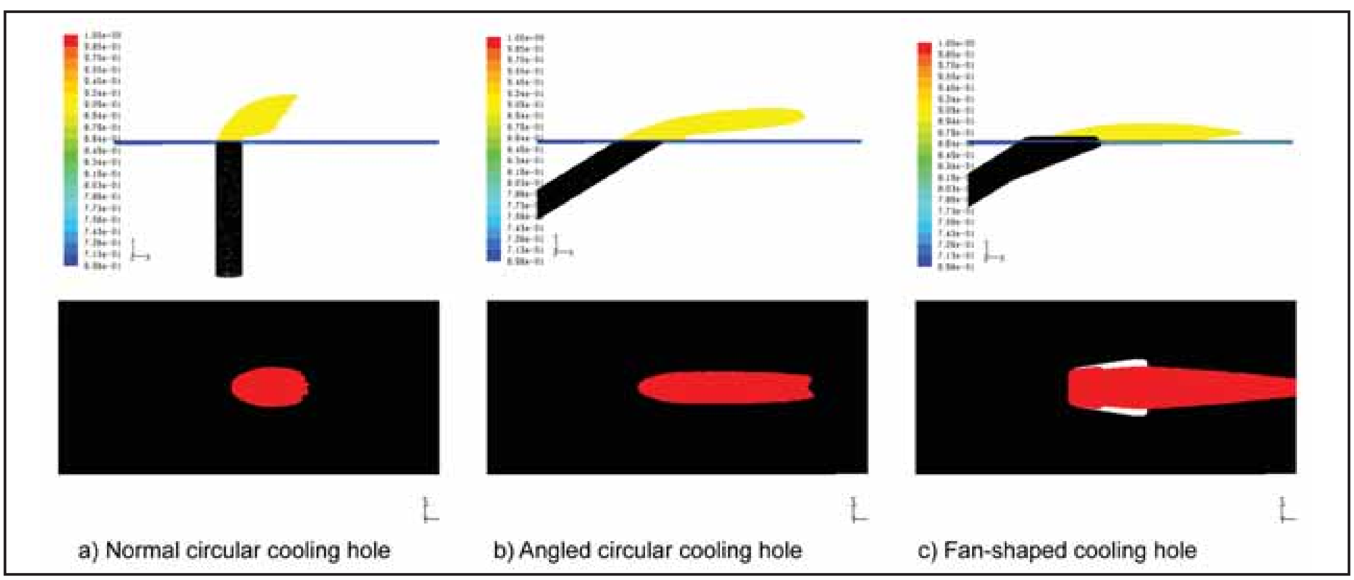

Figure 2. a-c Film-cooling jet structure for N_2 injection

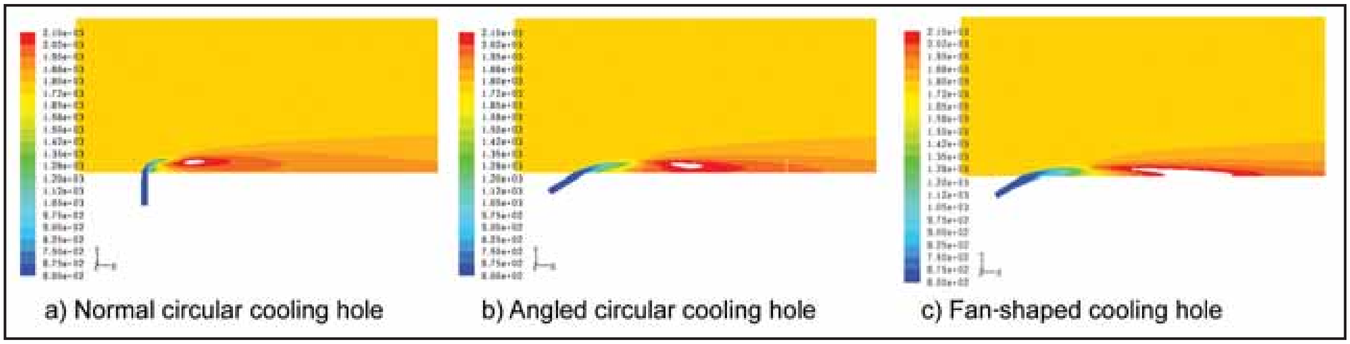

Figure 3a-c. Temperature contours on the midplane for Phi =1.5, M = 1.0, air injection

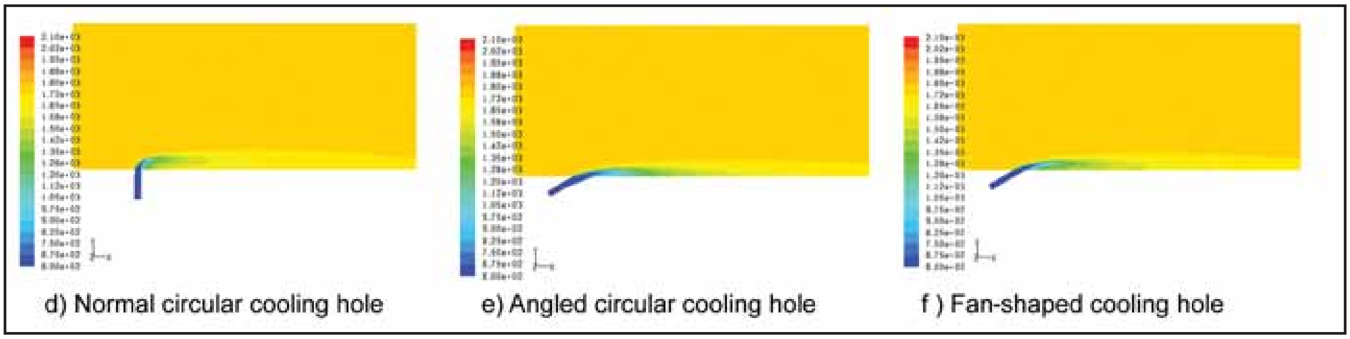

Figure 3d-f. Temperature contours on the midplane for Phi =1.5, M = 1.0, N_2 injection

rich-burning causes was almost negligible at elevated Phi when N_2 was used as the coolant.

Figure 3 shows the typical cooling effectiveness on the flat surface for air and N_2 injections at Phi = 1.5. For the cases of N_2 injection, the cooling effectiveness patterns were similar with the area of the higher level of cooling effectiveness right downstream of the coolant hole. In general, angled circular and fanned holes did better in terms of increased cooling effectiveness, compared with the normal circular hole. For cases with air injection, in addition to the area of higher level of cooling effectiveness right downstream of the coolant hole, a zone with "negative" values of cooling effectiveness appears further downstream. Since the cooling jet bends over much more for the fanned hole the cross-stream mixing occurs closer to the surface for the fanned hole causing increased chemical reactions and thus higher downstream temperatures. These temperatures were higher than the free stream and as a result, the downstream film-cooling effectiveness was negative farther downstream. The location of the negative cooling-effectiveness zone corresponds to the hot area during the secondary combustion. In classic literature on film cooling, cooling-effectiveness value is always positive because there is no reaction to add additional heat to raise the wall temperature above the free-stream temperature. The negative cooling effectiveness is one of the major unique characteristics in reactive flow film cooling, a subject deserving special attentions from researchers.

Three-dimensional CFD simulations of propane combustion during turbulent flow film cooling on a flat surface with a normal hole, an angled circular hole, and a fan-shaped hole for both air and N2 injections have been successfully carried out. Numerical simulation clearly captured the characteristics of the coolant jet, secondary combustion, and their interactions in turbulent flows. Numerical simulations compare well with experimental data in terms of relative temperature change between air and N_2 coolant injections. Film cooling from the normal circular hole performs poorly compared with the angled and fanned-hole film cooling both at the lean and in rich conditions. Under fuel lean conditions, the patterns of temperature, effectiveness, and species concentration fields for air injection were similar to that for N_2 injection for all the shapes. However, a dramatic increase in the heat release was calculated when the equivalence ratio was increased over stoichiometric. Under fuel rich conditions, a distinguished hot area downstream of the coolant hole can be generated by the secondary combustion because of air injection. Compared with N_2 injection, air injection at high Phi could result in negative cooling effectiveness in certain areas on the flat surface. This effect was most apparent for the shaped hole because of the oxygen being injected closer to the surface for this geometry. Because of this secondary combustion, air injection increases the wall temperature on the flat surface considerably for the shaped hole. Further downstream and in the immediate vicinity of the cooling hole, the film cools effectively to a greater length both axially and laterally when compared with the normal circular and fanned circular film-cooling flows.

Predicting "Ocean Weather" Using the 1/12° Global HYbrid Coordinate Ocean Model (HYCOM)[1]

By E. J. Metzger, H. E. Hurlburt, A. J. Wallcraft, Naval Research Laboratory; O. M. Smedstad, QinetiQ North America/Planning Systems, Inc.; J. A. Cummings, Naval Research Laboratory; and E. P. Chassignet, Florida State University

Introduction

The development of an advanced global ocean prediction system has been a long-term Navy interest. Such a system must provide the capability to depict (nowcast) and predict (forecast) the oceanic "weather," some components of which include the 3-D temperature (T), salinity (S) and current structure, the surface mixed layer, and the location of mesoscale features such as eddies, meandering currents, and fronts. The space scale of these eddies and meandering currents is typically ~100 km, and current speeds can easily exceed 1 ms^{-1} in the Gulf Stream (Atlantic Ocean) and Kuroshio (Pacific Ocean). Numerical ocean models with sufficiently high horizontal and vertical resolution are needed to depict the 3-D structure with accuracy superior to climatology and persistence (i.e., a forecast of no change). The accelerated development of these prediction systems would not have been possible without the computational resources provided by the DoD HPCMP. Throughout the research and development stages of numerical ocean models and data assimilation techniques, HPC has played a key role. This is especially true with regard to grand challenge projects that allowed development of high horizontal resolution global systems long before it became feasible to run them in an operational environment. In addition, nonchallenge and Capability Application Projects have also provided considerable resources toward advancement of these systems.

The existing two-model operational global ocean prediction system, run daily at the Navy DSRC, is based on the 1/8° Navy Coastal Ocean Model (NCOM) and the 1/32° Navy Layered Ocean Model (NLOM). Unlike NLOM, NCOM has high vertical resolution, but it has medium-range horizontal resolution (~15 km at mid-latitudes near 40°N), which makes it eddy-permitting. The next-generation system is based on a single model, the HYbrid Coordinate Ocean Model (HYCOM) (Bleck, 2002). It was developed as part of a multi-institutional consortium between academia, Government, and private industry. At 2.2 times the horizontal resolution of NCOM, the HYCOM-based system is eddy-resolving, a distinction associated with important dynamical implications for both ocean model dynamical interpolation skill in the assimilation of ocean data and for ocean model forecast skill (Hurlburt et al., 2008). HYCOM is also uniquely designed to allow an accurate transition between deep and shallow water, historically a challenging problem for ocean models. Its generalized hybrid vertical coordinate is a substantial advance over the vertical coordinate system

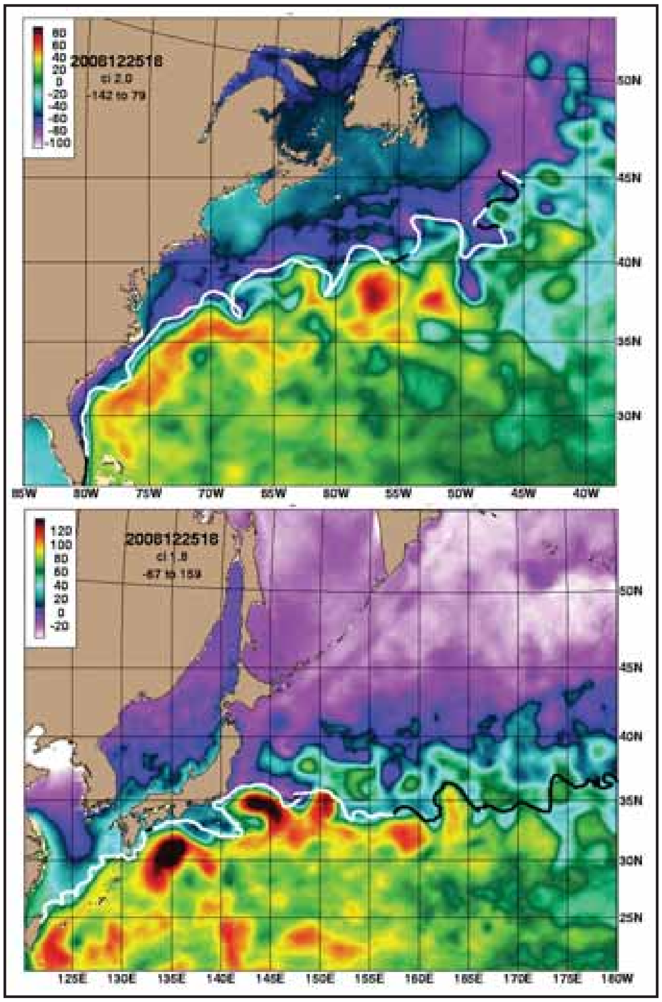

Figure 1. Sea surface height (cm) from the 1/12° global HYCOM prediction system for the Gulf Stream in the Atlantic Ocean (top) and the Kuroshio in the Pacific Ocean (bottom) on December 22, 2008. The ribbon of high gradient color shows the location of these western boundary currents; embedded within the meandering flow are warm and cold core eddies. The currents generally flow parallel to the isolines of height and are strongest where the gradients are the tightest. An independent infrared (IR) analysis of the north edge of both current systems is performed by NAVO and overlain on each image. A white (black) line means the IR analysis is based on data less (more) than 4 days old

[1] *HPC Insights* is reprinting this article, as essential information was inadvertently deleted in its last issue.

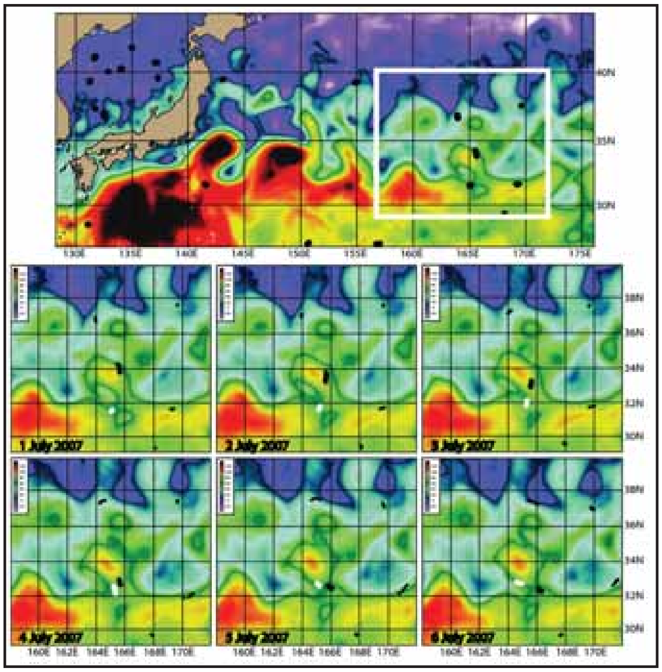

Figure 2. Sea surface height (cm) from the 1/12° global HYCOM prediction system for the Kuroshio on July 1, 2007 (top). Drifting buoy tracks over a 1-day time period are overlain on each panel. The white box defines the focused area of the bottom six panels that span the time frame July 1-6, 2007. A warm core eddy is about to detach from the Kuroshio, and two drifting buoys (highlighted in white and black) are traversing its western and eastern sides

used in NCOM. The HYCOM-based system represents the world's first eddy-resolving global ocean prediction system with both high horizontal and vertical resolution and has been validated against observational data (Metzger et al., 2008). It is scheduled for operational testing in 2009.

HYCOM Description

HYCOM is on a 1/12° global grid (mid-latitude resolution of ~7 km) with 32 hybrid vertical coordinate surfaces. The truly generalized vertical coordinate can be isopycnal (density tracking – often best in the deep stratified ocean), levels of equal pressure (nearly fixed depths – best used in the mixed layer and unstratified ocean), or terrain-following (often the best choice in shallow water). HYCOM combines all three approaches by choosing the optimal distribution at every grid point and time-step. The hybrid coordinate extends the geographic range of applicability of traditional isopycnic coordinate models toward shallow coastal seas and unstratified parts of the world ocean. It maintains the significant advantages of an isopycnal model in stratified regions while allowing more vertical resolution near the surface and in shallow coastal areas, hence providing a better representation of the upper ocean physics.

HYCOM employs the Navy Coupled Ocean Data Assimilation (NCODA) (Cummings, 2005), which is a fully 3-D multivariate optimum interpolation scheme, to assimilate observational data. The data include surface observations from satellites, including altimeter sea surface height (SSH) anomalies, sea surface temperature (SST), and sea ice concentration, plus *in situ* SST observations from ships and buoys as well as T & S profile data from XBTs, CTDs, and Argo profiling floats. The 3-D ocean environment can be more accurately nowcast and forecast by combining these diverse observational data types via data assimilation and using the dynamical interpolation skill of the model.

The 1/12° global HYCOM-based prediction system has been running daily in pre-operational mode at the Navy DSRC since December 22, 2006. Originally running on the IBM machines (*Romulus* – Power4+ and then *Babbage* – Power5+), the system was recently moved to the Cray XT5 (*Einstein*). Here it is presently configured to use 78 nodes (619 processors) to run HYCOM and produce the NCODA analyses with an additional two nodes set aside for pre- and postprocessing.

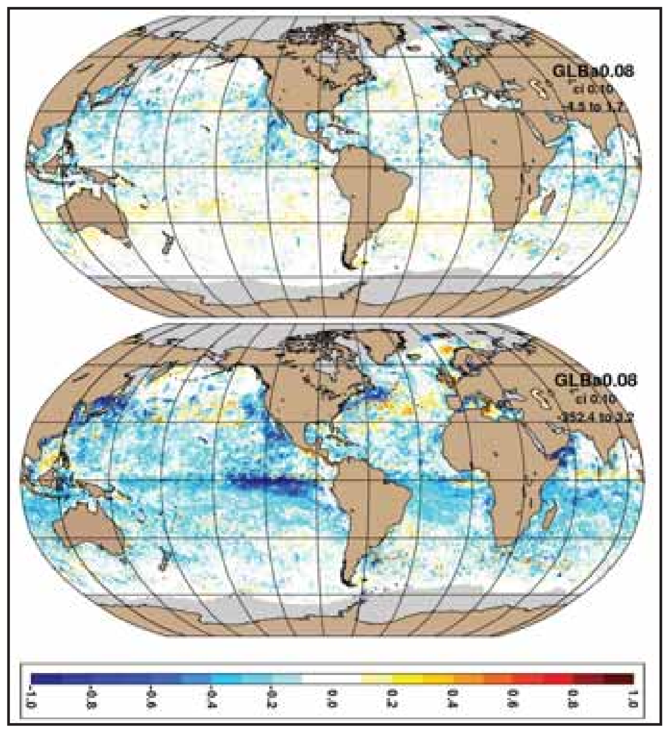

Figure 3. Sea surface temperature (SST) mean error relative to ~33,000,000 MCSST observations from the 1/12° global HYCOM prediction system at the analysis time (top) and for a 3-day forecast (bottom) over a year-long hindcast spanning June 2007-May 2008. Red (blue) colors indicate nowcast and forecast SST that is warmer (cooler) than observed. Values between ±0.1˚C are white. The gray area near the poles is an annual average ice coverage mask

HYCOM efficiently scales to large processor counts and can easily be configured to fit within the allocated resource window. Each day the system starts 5 days in arrears of the nowcast time (to assimilate all available late-arriving observational data) and then runs forward to create a 5-day forecast. It generates 3-D whole domain instantaneous archive files at 00Z each model day that are ~10 Gb.

Real-Time Results

Where possible, the HYCOM-based prediction system is evaluated using independent observations, and some examples follow. Figure 1 shows simulated SSH for the Gulf Stream and the Kuroshio systems. The assimilation of satellite altimeter SSH anomalies is essential to accurately map the circulation in these highly chaotic regions dominated by mesoscale flow instabilities. Infrared-based frontal analyses that show the northernmost edge of the currents are overlain on the panels. They provide an independent analysis of the current positions and clearly indicate the ocean nowcast/forecast system is accurately mapping these western boundary currents. Figure 2 shows an example that uses drifting buoy trajectories to validate the flow field in the Kuroshio. Drifting buoy temperature (but not velocity) is assimilated via NCODA, allowing the trajectory to be an independent validation source. The white box focuses on a warm core eddy about to detach from the Kuroshio, and a pair of drifting buoys is noted on the western and eastern sides. These two drifters pass within a half degree of each other while traveling in opposite directions. Close examination indicates the two buoys are on opposite sides of a saddle point that still connects the main current with the detaching eddy. Thus, the system is able to accurately assimilate the altimeter data and act as a dynamical interpolator. Lastly, SST forecast skill of the system is examined. Table 1 shows the mean error (bias) and root-mean-square-error (RMSE) as a function of forecast length. The bias and RMSE gradually grow with forecast length. The spatial distribution of the mean error is shown in Figure 3 for the analysis time and a 3-day forecast. In hindcast mode, the global HYCOM system has also demonstrated forecast skill on time scales up to a month for the meandering currents and eddies in some regions (not shown).

Table 1. SST error statistics as a function of forecast length from the 1/12° global HYCOM prediction system compared against ~33,000,000 satellite-based observations. The analysis is performed over a year-long hindcast spanning June 2007-May 2008 and is limited to the area between 45°S – 45°N

	Mean error	RMSE
Analysis	-.02	.36
1-day forecast	-.09	.44
2-day forecast	-.14	.52
3-day forecast	-.18	.60
4-day forecast	-.22	.67
5-day forecast	-.26	.72

Impact

A next-generation ocean nowcasting/forecasting system based on 1/12° global HYCOM is running in real-time at the Navy DSRC. It can accurately depict and forecast such features as western boundary currents and sharp ocean fronts, thus providing improved environmental awareness to the Fleet. Other naval applications include optimum track ship routing, search and rescue, anti-submarine warfare and surveillance, tactical planning, and providing boundary conditions for regional and coastal nested model. HPC resources have played a major role in making this state-of-the-art system feasible, beginning with the preliminary development of HYCOM and continuing all the way through its transition as a pre-operational product.

References

Bleck, R., 2002. An oceanic general circulation model framed in hybrid isopycnic-Cartesian coordinates. *Ocean Modelling*, 4, 55-88.

Cummings, J.A., 2005. Operational multivariate ocean data assimilation. *Quart. J. Royal Met. Soc.*, 131, 3583-3604.

Hurlburt, H.E., E.P. Chassignet, J.A. Cummings, A.B. Kara, E.J. Metzger, J.F. Shriver, O.M. Smedstad, A.J. Wallcraft, and C.N. Barron, 2008. Eddy-resolving global ocean prediction. In: Hecht, M., Hasumi, H. (Eds.), Ocean Modeling in an Eddying Regime, Geophysical Monograph 177. *American Geophysical Union*,Washington, DC.

Metzger, E.J., O.M. Smedstad, P. Thoppil, H.E. Hurlburt, A.J. Wallcraft, D.S. Franklin, J.F. Shriver and L.F. Smedstad, 2008. Validation Test Report for Global Ocean Prediction System V3.0 – 1/12° HYCOM/NCODA: Phase I, *NRL Memo. Report*, NRL/MR/7320—08-9148.

Air Force Research Laboratory DoD Supercomputing Resource Center

From the Director's Desk – Frank Witzeman


Frank Witzeman
Director, AFRL DSRC

The Air Force Research Laboratory DoD Supercomputing Resource Center (AFRL DSRC) recently participated in some exciting High Performance Computing Modernization Program (HPCMP) initiatives, and we're getting ready for the next HPC system expected under Technology Insertion 2010 (TI-10). We played key roles in the selection of two important contracts: (1) the User Productivity Enhancement, Technology Transfer and Training (PET³) effort awarded to High Performance Technologies Inc., and (2) the Storage Lifecycle Management (SLM) project awarded to General Atomics. We also completed our final design and established a schedule for necessary TI-10 facility upgrades and modifications. The decommission process was started for our 2048-processor SGI Altix 3700 system, *Eagle*, October 2009, and our 2048-processor HP XC Opteron system, *Falcon*, June 2010, both leaving just in time as we look forward to the new TI-10 system. Our remaining 9216-core SGI Altix 4700, *Hawk*, system was accepted 2 years ago and remains in high standing in the Top 500 list of the most powerful computers in the world (www.top500.org). It will remain in service for 2 more years. An interesting development for Hawk is the October switch from Platform Computing's LSF to Altair's PBS Pro for handling batch job scheduling and workload management.

Although *Eagle* and *Falcon* are at the end of their tenure with the HPCMP, we remain proud of their impact on DoD technology development and support to weapon systems. Immediately after their installation and acceptance in late 2005 and mid-2006, *Eagle* and *Falcon* were executing "Challenge" and "High Priority" projects that were resolving fluid dynamic phenomena encountered by aircraft during weapons carriage and release, investigating electromagnetic radiation characteristics of directed-energy devices, and simulating aircraft engines to refine designs for better efficiency and survivability. Today these systems are helping researchers understand the complex physics that dictate the flight of birds and insects so that micro air vehicles may be designed to mimic nature, and they are assisting in the analysis of interference between a naval hovering/landing vehicle and its host ship. In addition to such scientific discovery, significant flight-test and experiment savings of hundreds of millions of dollars, and substantial risk reductions have occurred as a result of the simulations and analyses carried out on these systems. Together *Eagle* and *Falcon* have provided over 30 million processor-hours a year for the past 4 years—a great contribution to the HPCMP.

Moving ahead to 2010 requires a significant amount of guesswork in trying to determine how much power and cooling will be needed for next-generation supercomputing systems. Our design for 2010 estimates an additional 4 MW electrical load on our facility, and an associated 1200 tons of cooling with new water-cooled chillers, effectively doubling our current capacity. We've also added a free-cooling component where we can take advantage of cold air outside and bypass the chiller plant, thereby saving operating costs in the winter months. We'll save more costs by removing our chilled water storage tank and implementing closed circuit chillers. To accommodate the new mechanical and computer systems, we'll make some major changes to the building power distribution system and install new panels and transformers.

The above facility changes become increasingly important as we expand our support to more and more unique customers who are not yet included in the HPCMP, are at the entry-level of supercomputing, or have dedicated (vs. allocated) systems located in our Center. Most, if not all, of our support to these customers results in partnerships and alliances to investigate new HPC technologies while advancing the research, development, test and/or evaluation objectives of the users' organizations. One example is the introduction of a 4536-processor SiCortex system into our environment for computational fluid dynamics algorithm development and application to air vehicle design and analysis. Another example is our hosting of a 2048-core SGI Altix ICE 8200 system (coupled with SGI Altix 450 system) purchased by the HPCMP for the AFRL Sensors Directorate under a Dedicated HPC Project Investment (DHPI) initiative. This system, named *Desch* for the Dayton, Ohio, pioneer of code-breaking computers Joseph Desch, is being used to perform radar image processing and analysis for persistent surveillance in all-weather conditions and complex terrains. There are a number of other collaborations started, and we are expecting more next year.

Beyond 2010, we are excited about the construction of a new facility at Wright-Patterson Air Force Base that will provide modern data center space in 2012. This new building is part of a multiphase Information Technology Complex (ITC) construction project supported by various Air Force organizations. It is anticipated that the AFRL DSRC will occupy 10,000 square feet that will have power and cooling capacity for up to 8 MW of computer system(s). Subsequent phases of the project, if completed, will provide more data center space and modern offices for employees. The long-term goal is to vacate our current facility, which has served as a computer center since 1971, and populate our allocated spaces in the new ITC. Space, power, and cooling become critical elements in our strategy to operate and grow the AFRL DSRC in support of the HPCMP mission and the requirements of our special laboratory customers. We remain committed to providing quality HPC service and capabilities while striving to reduce operating costs and exceed our sponsors and customers expectations.

Gotcha Supercomputer, Desch, Ribbon Cutting Ceremony

The ribbon cutting ceremony for the Gotcha supercomputer, *Desch*, was held on August 31, 2009, at the AFRL DSRC. About 51 scientists, engineers, and support staff attended the ceremony. Among them were distinguished guests Dr. Michael Kuliasha (AFRL Chief Technologist), Dr. David Jerome (Director of AFRL Sensors Directorate), Col. Cleophas Hockaday (Deputy Director of AFRL Sensors Directorate), Frank Witzeman, Jr. (Director of AFRL DSRC), Deborah Desch Anderson (daughter of Dayton's brilliant engineer and code-breaker Joseph Desch), Jim Brinker (Vice President of SGI Federal Service), Amanda Wright Lane (great-grandniece of Wright Brothers), Christopher Ristich (Chief of AFRL ATR Division), Edmund Zelnio (Inventor of Gotcha Radar Concept), Lori Westerkamp (Technical Advisor of AFRL ATR Division), Martin Justice (Chief of AFRL ATR Signature and Modeling Branch), and Patricia Ryan (Technical Advisor for AFRL ATR Signature and Modeling Branch).

The supercomputer is named after Dayton engineer Joseph Desch, who led a team that built a computing machine during World War II to decipher messages encrypted using the Nazi Enigma code. Desch died in 1987, never revealing even to his family, his highly successful, but secret, contributions to the allied war effort.

Desch is a custom-designed SGI Altix ICE 8200 supercomputer that has been optimized and will be dedicated exclusively to support real-time translation of synthetic aperture radar data into high-resolution three-dimensional video images from the Gotcha radar system being developed by the AFRL.

"The goal is to provide an extremely high-fidelity, all-weather intelligence, reconnaissance and surveillance, or ISR, capability that can observe activity over an entire city," said Dr. Dave Jerome."

Desch is a system paired with a dedicated smaller companion SGI Altix 450 system that uses shared-memory architecture, optimized to speed translation of *Desch's* high-resolution images into a virtual mosaic that users can manipulate.

Dr. Michael Minardi, Gotcha Program Manager, said that the smaller system, nicknamed *Bombe* after the World War II decoding machines designed by Joseph Desch, significantly enhances processing efficiency and real-time imagery output.

"I don't think it's possible that any honor would fit him more," Mrs. Desch said, referring to her father's intense curiosity of how complex machines work and the naming of a supercomputer dedicated to defense research in his memory.

Desch is built on 2048 Intel Nehalem processors, has 3 terabytes of random access memory, 87 terabytes of fast storage, and 16 gigabits-per-second data communication between processors, according to Tom Majumder, a Sensor's Directorate engineer.

"One of Gotcha's advantages is the flexibility of its recorded radar data coupled with the supercomputer's blazing processing speed. The same set of data can be used to look for many things just by changing the algorithms that the supercomputer processes," said 1st Lt. Curtis Casteel, an ISR engineer with AFRL's Sensors Directorate.

Desch was funded by the DoD High Performance Computing Modernization Program as one of the four FY09 DHPI awards.

Dr. Michael Minardi demonstrates Gotcha in action on a dynamic touch screen display

Eagle's "Wings" Widespread

As of October 1, 2009, the SGI Altix 3700 Supercluster (*Eagle*) will no longer run jobs at the AFRL DSRC. Its term has expired and it is time for it to move aside. Even a supercomputer becomes victim to the growth of the integrated circuit that doubles approximately every 2 years.

When *Eagle* went into production in the spring of 2005, it was heralded as a premium supercomputer and listed 23rd on the Top 500 list. With 2048 Processors (13.1 teraflops), 1 gigabyte memory, and 100 terabyte workspace, it was the belle of the ball waiting to have its dance card full.

SGI Altix 3700 vs. SGI Altix 8200 floor space comparison

The immense physical impact *Eagle* had on the AFRL DSRC facility in terms of cooling, power consumption, and its physical size has been minimized with newer technology. To put it into perspective, a new SGI Altix ICE 8200 system contains 2048 processors, the same as *Eagle*, but has a dramatically smaller footprint and, in turn, requires less cooling and power.

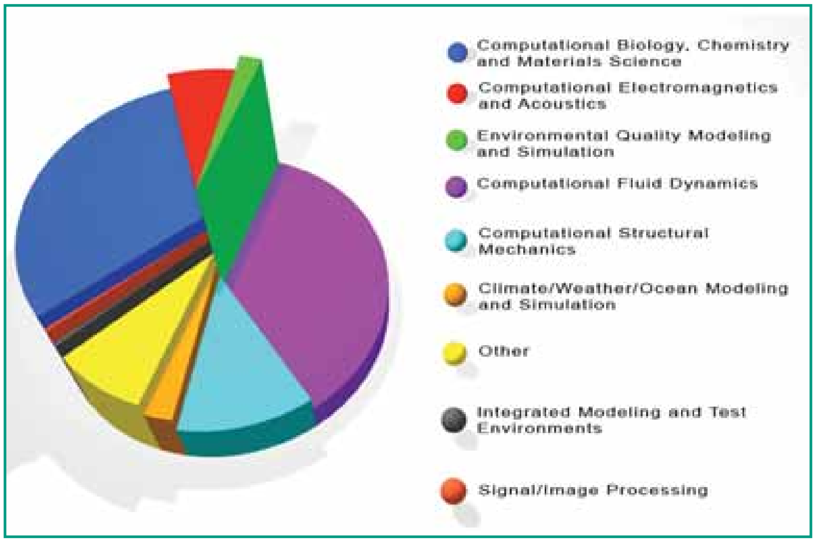

Eagle's effects are widespread and have encompassed many areas across the program. The chart below gives a graphical representation of *Eagle's* total number of runtime hours since 2005 and how the 180,108,150.413 total hours are broken down by computation technology areas (CTAs).

The knowledge and discoveries gained through the hours of research completed on *Eagle* will continue to impact the warfighter. The projects run on *Eagle* reflect the commitment of the DoD HPCMP and the AFRL DSRC to supporting projects that impact the warfighter well into the 21st century.

Projects run on the SGI Altix 3700 included the following:

- A High-Resolution (0.5mm) Simulation of a HPM Munition: Magnetron with Pulsed Power and Antennas
- Advanced Chemical Oxygen-Iodine Laser Technology Development using 3-D Navier-Stokes Simulation
- 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
- Characterization and Prediction of Stratospheric Optical Turbulence for DoD Directed Energy Platforms
- Computational Chemistry Modeling of the Atmospheric Fate of Toxic Industrial Compounds (TICs)
- Computational Simulations of Combustion Chamber Dynamics and Hypergolic Propellant Chemistry for Selectable Thrust Liquid/Gel Rocket Engines
- Computer Design and Simulation of Molecular Devices and Energy Sources for Naval Applications-Pederson
- Coupled Aircraft/Ship Performance Prediction for Dynamic Interface
- Coupled CFD/CSM/DPM Modeling of Structure Response to Blast Loading
- Decision Support for Seismic and Acoustic Sensors in Urban Terrain
- Design of Energetic Ionic Liquids
- Design of Materials for Laser Protection Applications
- Direct Simulation of Nano-scale Plasticity
- Full Annulus High Fidelity Fan and Compressor Simulations
- High-Fidelity Multidisciplinary Simulation of Biologically Inspired Micro Air Vehicles
- HYCOM Global Ocean with Tides
- Integration of Simulation and Test for Strike Aircraft Design and Development
- Modeling of Targeting Deeply Buried C4I and WMD Facilities
- Multi-scale Predictability of High-Impact Weather in the Battlespace Environment
- Polynitrogen/Nanoaluminum Surface Interactions
- Prediction Capability for High-Speed Surface Ships
- Simulation of a Dynamically Maneuvering Unmanned Combat Air Vehicle
- Statistical Fatigue and Residual Strength Analysis of New and Aging Aircraft Structure
- V-22 Roll-off in Vortex Ring State Coupled CFD/CSM/DPM Modeling of Structure Response to Blast Loading
- Vulnerability of Structures to Weapons Effects

The future of this *Eagle* is unknown at this time. It may work as a total, or in parts, for another agency (s) because although its work is finished at the AFRL DSRC, it has the capability to serve as a reliable source of computing power.

Army Research Laboratory DoD Supercomputing Resource Center

From the Director's Desk – Charles J. Nietubicz

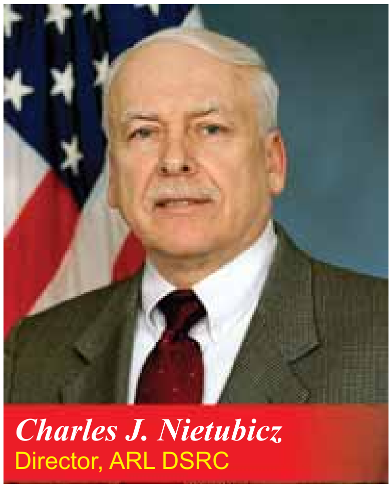

Technical Insertion 2009 in Full Swing at ARL

By Brian Simmonds, ARL DSRC Outreach Lead

Early in the morning on Monday, June 29th, the “semis” arrived with the largest of the Technology Insertion 2009 (TI-09) systems. The activity outside the building attracted the usual attention of the ARL employees arriving at work that day. There was a flurry of activity, forklifts, unpacking, and eventually the connections being made both with the hardware and the people. A new system had arrived at the ARL DSRC.

Over the next few days, the SGI technicians swarmed the newly placed cabinets and began connecting the multitude of cables and associated equipment that comprises the modern supercomputer. By Thursday, the computer was powered on, and all nodes, as well as the storage, were operational. The Linpack benchmark was then run on Friday. It showed a speed of 109.3 teraflops, making this system the most powerful computer ever at ARL.

Two months previous had seen the arrival of the test and development system *ICECUBE*. Once installed, the Customer Service and Admin teams began setting up, compiling, and tuning applications in preparation for the larger systems. The most challenging was integration of the new cueing system. This will be the first system at ARL to premiere with the PBS cueing system. PBS is the batch scheduler system that will be

the new standard across the entire Modernization Program.

The system has 10,752 cores (2688 2.8-GHz Intel Nehalem quad-core processors) with 32 TB of system memory and 600 TB of local disk storage.

The system is named *Harold* to honor Harold Breaux, an integral part of the ARL computing center landscape from the inception of the Army Supercomputing program that eventually evolved into the DoD High Performance Computing Modernization Program (HPCMP). Harold had a 21-year career as a research mathematician before transitioning to a management position in 1994, a time when the Army began its quest for acquiring supercomputers. Harold was a major player in this Army-wide effort to include efforts for acquiring systems for ARL and for the Army, where he served as Executive Secretary for the Army HPC Functional Coordinating Group chaired by the Army Director for Research and Laboratory Management. When Congress mandated that the DoD create an HPC Modernization Program, Harold was named the Army lead to the Working Group that created the Program. Wearing his ARL Management Hat, Harold led the effort to write the proposal leading to ARL being one of the first three Labs to be chosen to host a Major Shared Resource Center (MSRC). In 2005, he was given a "HERO" award by the Defense Department for his long-term contributions to the High Performance Computing Modernization Program.

The smaller of the TI-09 systems, named *TOW*, also arrived in late May. All of the TI-09 systems are expected to be online and available to users this fall.

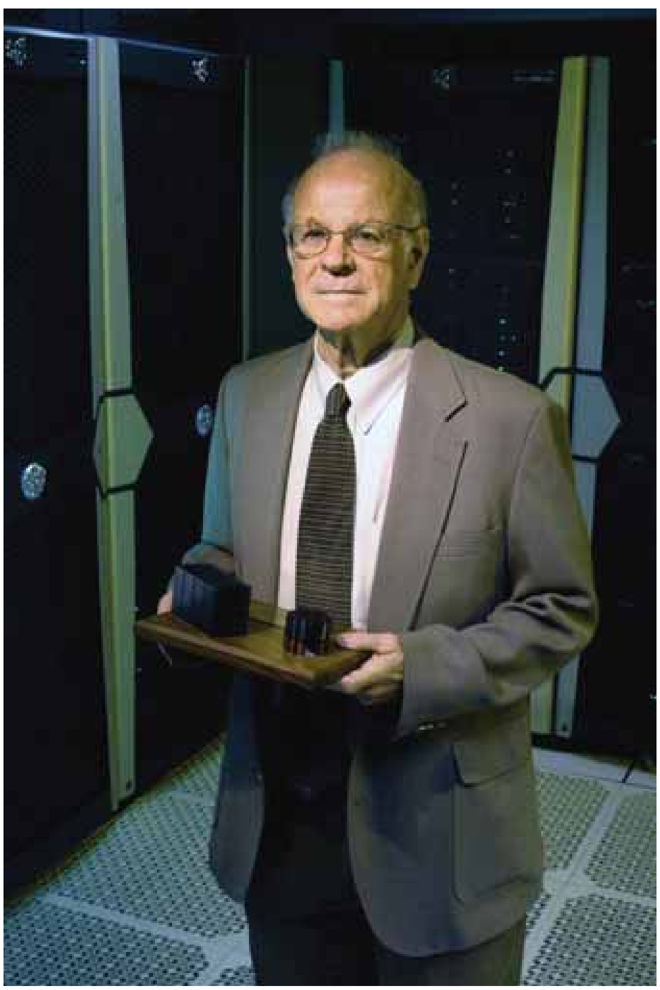

Harold Breaux with the new SGI Altix ICE and a scale model of the first Army Supercomputing Program acquisitions, a Cray2 and an SGI Power Challenge Array

Arctic Region Supercomputing Center DoD Supercomputing Resource Center

Technical Exellence

By Frank Williams, Director

Frank Williams
Director, ARSC DSRC

Each of the six DoD Supercomputing Resource Centers (DSRCs) has an important role in enhancing and sustaining the High Performance Computing Modernization Program (HPCMP). Finding ways to define and achieve an overarching goal of enhancing technical excellence for the DoD HPCMP has been an especially rewarding challenge that Charlie Nietubicz, Director, Army Research Laboratory DSRC, and I have been working on since the six DSRCs were recognized in January 2009.

Our task has focused on proposing and leading efforts designed to enhance technical excellence that extend across the six Centers and the Next Generation Technical Services contractor, Lockheed Martin. We have defined an accomplishable vision of technical excellence for the Modernization Program that uses a pathway of accomplishment leading to a state of technical excellence across the HPCMP.

Central to the pathway are a high level of customer satisfaction, unsolicited customer advocacy, and high demand for resources and services from within and beyond existing membership in the DoD HPCMP community. Achieving a state of technical excellence will be recognized when the DSRCs do the following:

- Provide a full spectrum of services and resources.
- Take holistic approaches for overall optimization of solution and service delivery.
- Through users' research and development projects, provide significant accomplishments in advancement of science and technology.
- Operate smoothly with the best technology and practices known to the community.
- Assume responsibility for programwide technical excellence through decentralized authority and appropriately centralized operational/commodity activities.

Now, while our recommendations for making steps along the path are being considered, you can contribute by voicing ways you believe the DSRCs can demonstrate excellence to you, wherever you are in the HPC community.

ARSC DSRC Open-Systems Access

By Debra Damron, ARSC DSRC, Communications Group Leader

A National Agency Check (NAC) can take months to process. In the case of Dr. Jaroslav Vacek, an ARSC DSRC user from the Czech Republic, it took 1 year and 3 months. That's a lot of supercomputing time that could have been used for investigating the properties and potential applications of molecular rotors, a project Vacek is working on with funding from the Army Research Office. Principle Investigator for the molecular rotors project is Dr. Josef Michl from the University of Colorado at Boulder.

Luckily for Vacek, he didn't have to wait more than a year to do his work. Because the ARSC DSRC does not require users to have a NAC when they apply for an account, Vacek was able to gain access to Arctic Region Supercomputing Center (ARSC) DSRC's resources in a matter of days instead of months.

The ARSC DSRC is the only Center in the DoD High Performance Computing Modernization Program (HPCMP) that provides open-research computing capabilities. The Center provides a full range of computational and data systems to foreign nationals (FNs) who are members of DoD HPCMP projects.

This includes *Pingo*, a Cray XT5 with 3456 compute nodes, and *Midnight*, a Sun Opteron cluster with 2312 compute processors. *Pingo* has 31.8 estimated TFLOPs theoretical peak performance and uses Cray SeaStar interconnect. *Midnight* has 12.02 TFLOPS and uses a SuSE Linux operating system.

The ARSC DSRC became an open-research Center Octotber 1, 2004. Since then, FNs on HPCMP project teams can apply for an account and are not required to have a NAC at the time of their application. Potential users are required to provide proof of citizenship (and current visa, if applicable) and consent to routine standard background checks in lieu of a NAC.

Rapid access to high performance environments provides principal investigators with the ability to quickly adapt or change memberships as a project group changes. Open-research systems are also good if a user only needs access for half of the year, as it's more timely than waiting for a NAC. This is especially useful in the case of interns and for offering HPC educational opportunities for short periods of time.

In most cases, having a NAC is preferable for users, because it provides more access to HPCMP machines and resources. However, being able to take full advantage of time allocated on machines in a matter of days, as opposed to a matter of months, is a huge advantage provided by the ARSC DSRC's open-system access.

Computational science interns tour the ARSC DSRC machine room. Because of the Center's open-research systems, summer interns can rapidly access HPC resources without applying or waiting for a National Agency Check

ARSC DSRC
DoD Supercomputing Resource Center

U.S. Army Engineer Research and Development Center DoD Supercomputing Resource Center

New Director Announced

Dr. Robert S. Maier
Director, ERDC DSRC

On August 21, Dr. James R. Houston, Director of the U.S. Army Engineer Research and Development Center (ERDC), announced the selection of Dr. Robert S. Maier as the Chief of the ERDC Information Technology Laboratory (ITL) Scientific Computing Research Center (SCRC) and Director of the ERDC DoD Supercomputing Resource Center (DSRC).

"My assignment is to make the SCRC an instrument of growth and scientific discovery for the seven ERDC laboratories and an exemplary operational unit of the DoD High Performance Computing Modernization Program (HPCMP)," said Maier. "Our talented staff is eager to help ERDC researchers leverage supercomputing assets to execute their R&D missions. The computational methods developed by our staff are of great importance to the entire DoD community and help us serve the HPCMP as an effective agent for cross-Center technical innovation.

"The most satisfying part of working at ERDC and ITL is the opportunity to work with scientists and engineers from other disciplines and laboratories, solving challenging problems with the help of scientific computing methods and resources," said Maier.

Dr. Maier began his career in computational science with the Control Data Corporation in 1981 as an applications analyst. He earned his Ph.D. from the University of Minnesota Department of Computer Science in 1990 in the area of numerical analysis and joined the Army High Performance Computing Research Center in 1991 as a postdoctoral fellow and then as a staff scientist in 1993. He began his ERDC career in 2003 in the ERDC DSRC.

JSU Engineering Student Interns Involved in Leading-Edge Research

By Dr. Gerald R. Morris, ERDC DSRC Computer Scientist

From May 18 – August 21 2009, the ERDC DSRC had the opportunity to involve five Jackson State University (JSU) engineering student interns in some leading-edge research. Antoinette Anderson, Ales-cia Malone, Ricky McGruder, Jarvis McWilliams, and Nikeya Peay, who are all pursuing M.S. degrees, worked with Dr. Gerald R. Morris on an HPCMP-funded project entitled "High Performance Computational Design of Novel Materials (HPCDNM)."

The HPCDNM project, which is in its fourth year, involves several basic research tasks. The students worked with Dr. Morris on HPCDNM Task 1, which is entitled "Accelerating Scientific Applications with High Performance Reconfigurable Computing." Under Task 1, these researchers are investigating a promising new alternative computational technology, wherein field programmable gate arrays

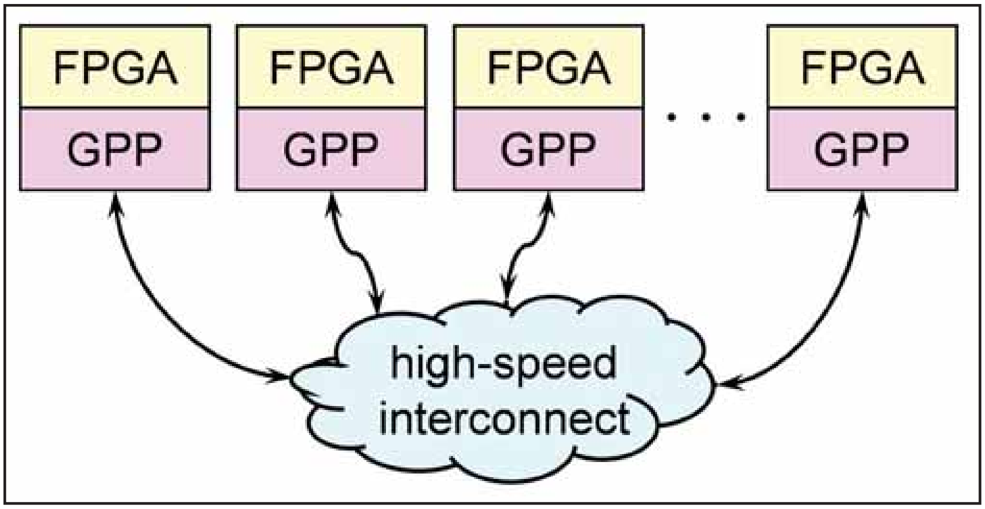

Figure 1. HPRC Cluster

(FPGAs) are used in conjunction with general purpose processors (GPPs).

As seen in Figure 1, high performance reconfigurable computers (HPRCs) are parallel computing clusters that contain multiple GPPs and FPGAs connected via a high-speed network. The GPPs perform part of the computation and usually take care of the communication between nodes. The FPGAs are configured as application-specific coprocessors such that the application runs faster while using fewer nodes. During their stint at ERDC, the JSU interns were tasked with investigating the tools, technologies, and methodologies associated with mapping floating-point scientific applications onto HPRCs. All five students, who are fully funded under a National Science Foundation grant entitled the Louis Stokes Mississippi Alliances for Minority Participation (LSMAMP), will continue their HPRC research during their matriculation at JSU. This continuation is possible, because Dr. Morris, who is an adjunct professor of computer engineering at JSU, is on their graduate committee. As part of the LSMAMP "Bridge to the Doctorate" program, all of these students will be pursuing a Ph.D. degree after completing their M.S. degrees.

U.S. Senator Roger F. Wicker Visits

By Rose J. Dykes, ERDC DSRC Editor/Writer

U.S. Senator Roger F. Wicker of Mississippi and his staff recently toured the ERDC DSRC. He expresses his strong feelings for the DoD in his recent weekly column: "As a member of the Senate Armed Services Committee, I have the opportunity to participate in many of the debates that shape policies important to our Nation's defense and the men and women of our armed forces. One of the most important bills the committee considers each year is the defense authorization measure, which sets the policy and spending priorities for the Department of Defense. Following committee approval, this bill was recently passed by the full Senate. This important measure authorizes funds our troops need to achieve their objectives in the field, as well as a number of provisions essential to defense installations in Mississippi."

The Honorable Roger Wicker (left), U.S. Senator of Mississippi, shakes hands with Dr. Reed Mosher, ERDC Information Technology Laboratory, after autographing the Cray XT4 supercomputer (Jade)

Senator Wicker replaced Senator Trent Lott, who resigned in December 2007, and joins the senior U.S. Senator from the State of Mississippi, Thad Cochran, who is the Ranking Member of the Senate Appropriations Committee and its Subcommittee on Defense. Senator Cochran serves as Vice Chairman of the U.S. Senate Committee on Appropriations.

Maui High Performance Computing Center DoD Supercomputing Resource Center

New Director – David Morton

David Morton is the new Director of the Maui High Performance Computing Center DoD Supercomputing Resource Center (MHPCC DSRC). Mr. Morton has been involved with high performance computing (HPC) since 1988. He has served in a variety of HPC leadership positions including Director of Engineering at Cray Research and SGI as well as Vice President of Engineering and Chief Technology Officer for Linux Networx. In addition, he was the Technical Director at the MHPCC and is the named inventor on three HPC patents. He received his Bachelor of Science degree in mechanical engineering from the University of Illinois, Urbana, Illinois, a Master of Science degree in mechanical engineering, University of Illinois, Urbana, Illinois, and a Master of Business Administration degree, Carlson School of Business, University of Minnesota, Minneapolis, Minnesota. His major awards include the Cray Research Leadership and Innovation Award; Defense Programs Award of Excellence for work on the LANL Blue Mountain Supercomputer; three patents issued, two of them enabling packaging technologies for new product lines; and Pi Tau Sigma (Mechanical Engineering Honor Society) member.

New Deputy Director – Marie Greene

Marie Greene joined the MHPCC as the Deputy Director in July 2009. Ms. Greene holds a Bachelor of Science degree in biochemistry and a Master of Science degree in computer science, both from the University of California at Riverside (UCR). She initially joined the MHPCC back in 2001 as the Visualization and Systems Services Lead and was responsible for establishing the MHPCC visualization facility, as well as leading the development of visualization efforts for the MHPCC DoD programs. Prior to coming to MHPCC, she founded and was director of three high performance visualization computing centers at UCR over a period of 18 years. Under her leadership, the multimillion dollar facility named the Center for Visual Computing (CVC) became self-sustaining within 3 years. Ms. Greene was the UCR Resource Control Manager for the San Diego Supercomputing Center (SDSC) and was the Acting Associate Vice Chancellor for the UCR Computing and Communications department. She was actively involved in the preparation, deployment, and analysis of over a dozen scientific experiments on eight Space Shuttle flights. She utilized image and signal processing to analyze these experiments and was a pioneer in publishing a method and new software to build three-dimensional volumetric reconstructions. Ms. Greene began her career as a software engineer for the Naval Sea Systems Command (NAVSEA), Naval Surface Warfare Center (NSWC) in Seal Beach, California. She has more than 50 periodical acknowledgments, and her visualization work has appeared in more than 200 scientific journals, reports, books, and multimedia.

Technical Director – Dr. Cliff Rhoades

Since January 1, 2008, Dr. Clifford E. Rhoades, Jr., has been serving as the MHPCC DSRC Technical Director under an Inter-Governmental Personnel Act from the Carnegie Mellon University Software Engineering Institute. Dr. Rhoades holds a Ph.D. in physics from Princeton and is one of the first five individuals honored as a Fellow of the American Physical Society for contributions to computational physics. The author of more than 100 publications and reports, Dr. Rhoades is well-known for his contributions to high performance computing and communications. He was first appointed to the Senior Executive Service (SES) in 1983, serving as chief of the Computer Systems and Research Division, National Aeronautics and Space Administration – Ames Research Center. Prior to becoming the SES Director of Mathematics and Space Sciences at the Air Force Office of Scientific Research, he was a member of the scientific staff at the University of California Lawrence Livermore National Laboratory, where, among other positions, he served as associate theoretical physics division leader and deputy computational physics division leader. He is a graduate of the Air Force Reserve Officer Training Corps program, Squadron Officer School, Air Command and Staff College and Air War College, a retired U.S. Air Force commissioned officer, and retired from the Senior Executive Service.

MHPCC DSRC AFRL Program Manager – CPT Joseph Dratz

CPT Joseph Dratz is currently the Deployed Systems Program Manager at the Air Force Research Laboratory Detachment 15 in Maui, Hawaii. In addition, he is the military Program Manager assigned to the MHPCC DSRC. CPT Dratz was commissioned in 2004 through the Reserve Officer Training Corps. His first assignment was to the Air Force Technical Applications Center (AFTAC) at Patrick AFB, Florida, where he served as a nuclear debris alert officer and as Officer in Charge of the Nuclear Test Branch. While at AFTAC, he led the Nuclear Debris Collection and Analysis Team that collected the only physical evidence of North Korea's first nuclear test. He received his Bachelor of Science degree from of the University of North Carolina. He also graduated from the Air and Space Basic Course, Maxwell AFB, Alabama, and the Squadron Officer School, Air University, Correspondence. He has been honored and decorated with the Air Force Commendation Medal and the Air Force Science and Engineering Award for Advanced Technology Development.

Much Progress at MHPCC DSRC

By David Stinson, Past Acting Director

It is truly an exciting time for the MHPCC DSRC. It became one of the six DSRCs back in January 2009, and in August 2009, it introduced a new 9216-core Dell cluster into the Program. The new system, *Mana*, went from start of install to operational in about 6 weeks. On August 21, 2009, a dedication ceremony attended by a number of dignitaries was held for the new machine with guest speakers including the Honorable Senator Daniel Inouye from Hawaii along with Dr. M.R.C. Greenwood, the new president of the University of Hawaii. Benchmarks show the new machine to be a real performer with Linpack scores putting it around number 45 on the June 2009 Top 500 List. In addition, we are upgrading power and cooling in our machine room to enable running both Jaws and Mana simultaneously. For the long haul, the MHPCC DSRC is planning for additional raised computer floor, improved office space, and modular expansion of infrastructure to accommodate future acquisitions. MHPCC hopes to offset its power costs using photovoltaics and is in the process of selecting a contractor for the $4 million Maui Energy Improvement Initiative, which will demonstrate this technology. The MHPCC DSRC hopes that these improvements will position it advantageously for success and a productive future.

On the staffing side, the MHPCC is pleased to welcome David Morton and Marie Greene as the new Director and Deputy Director, respectively. Both have prior history at the MHPCC, and David is coming with an outstanding industry background that includes experience at SGI. They join Captain Joe Dratz, special projects officer and former MHPCC Acting Director and Deputy Director, and Dr. Cliff Rhoades, Technical Director. Both Joe and Cliff are critical assets for the MHPCC and are instrumental in keeping the business processes running smoothly. Add to this an outstanding contractor team led by Executive Director Gene Bal of the University of Hawaii and a second-to-none contract and finance group including Liza Herrera, Crystal Price, and Irma Aragon at Kirtland, AFB, and you have a staff that can make anyone look good. Many thanks to each of them as the MHPCC DSRC transitions to a new management team. I have thoroughly enjoyed my tenure here on Maui, am encouraged by all the successes we have had, and am proud to have had a part in producing such a vibrant organization going into the future.

MHPCC DSRC Doubles Its Computing Power

Mana is the newest MHPCC supercomputing system, with a peak troughput rating of 103 teraflops (103×1012 floating point operations per second). *Mana* is based on Dell's new PowerEdge M610 series with half-height blade architecture, arranged in 1152 compute nodes, each with two 2.8-GHz quad-core Intel Nehalem processors and 24 GB RAM (3 GB/core) – a total of 9216 compute cores. The interconnect fabric is Dual Data Rate Infiniband, and the system is configured with nearly 400 TB of direct-attached, DataDirect disk.

MHPCC's new Dell system has been designated with the Hawaiian name Mana. Mana is a spiritual quality considered to have supernatural origin – a sacred power of the elemental forces of nature embodied in an object or person. To have mana is to have influence and authority – the power to perform

This installation is the largest in the history of the MHPCC DSRC. It advances the MHPCC DSRC to the forefront of high performance computing and places the MHPCC DSRC among the leaders in the DoD research and development community.

With the introduction of the new *Mana* system, the MHPCC overall compute portfolio doubles in size. However, the *Jaws* system will not be decommissioned. Requirements assessments and redeployment plans will likely result in the redistribution of *Jaws* assets to meet a number of niche requirements. Partitions of *Jaws* will serve as a Dedicated High Performance Computing Project Investment (DHPI), the Maui Space Surveillance System Advanced Image Reconstruction project, and are also being considered for research applications by the University of Hawaii under the Educational Partnership Agreement with the Air Force Research Laboratory.

MHPCC marked the introduction of *Mana* with a dedication event on August 21st. Dignitaries on the program included United States Senator Daniel K. Inouye and the new president of the University of Hawaii, Dr. M.R.C. Greenwood. Notable attendees included Cray Henry, Director of the DoD High Performance Computing Modernization Program, the Mayor of Maui County, state and county legislators, members of the University Board of Regents, and several members of the University executive management staff.

Dignitaries attending the Mana Dedication include (front row, left to right) Laura Ulibarri, Air Force Maui Optical & Supercomputing Site Branch Chief, Reverend Kalani Wong, Dr. M.R.C. Greenwood, UH President, Senator Daniel Inouye and his wife, Irene Hirano, Cray Henry, DoD High Performance Computing Modernization Program, Director, and his wife, Lizabeth Henry

Reverend Kalani Wong delivers the Dedication invocation for the ribbon cutting ceremony of Mana at the MHPCC DSRC Data Center

Navy DoD Supercomputing Resource Center

Preparing for the Next Decade – Navy DSRC Looks Ahead

By Tom Dunn, Director

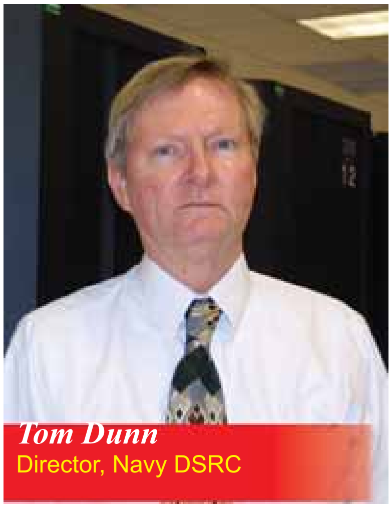

Moore's Law—the trend that processor speed doubles every 18 months—is expected to hold true for at least another decade. With that in mind, the Navy DoD Supercomputing Resource Center (DSRC) is preparing for what will likely be highly dense high performance computing (HPC) systems to be acquired in 2011 and beyond. Preparations include continuing a second phase of facility upgrades, being a strategic partner in the DoD High Performance Computing Modernization Program (HPCMP) new storage initiative, expanding disaster recovery capabilities, and continually researching for the best tools and methods of assisting our users in getting optimal performance from our existing and upcoming HPC systems.

In this publication, you'll find details of the HPCMP's Enterprise System Monitoring (ESM) effort, in which the Navy DSRC is providing leadership to and teamwork with other DSRCs. Through HPC infrastructure and outstanding outreach and support to our users, the Navy DSRC continues to provide the most robust means of cultivating the advanced research and development that are critical to the support of the warfighter.

Enterprise System Monitoring: Ensuring Wide Infrastructure Availability

By Carlos Cuevas, Lockheed Martin (LM) Network Engineering Manager, Navy DSRC; Randy Becnel, LM Technical Lead, Navy DSRC; Morris Ramsey, LM Project Manager for ESM; and Tom Brown, Associate Director and Executive Agent for ESM, Navy DSRC

The primary overall benefits from ESM are as follows:

- Provides DSRC systems status to the local DSRC and enterprise systems status to the Systems Operations Center/Local Area Network Operations Center.
- Establishes a consistent, common methodology and approach for systems monitoring of the four participating DSRCs.
- Identifies potential problems before they impact service.
- Utilizes commercial off-the-shelf products to minimize Life Cycle Cost and achieve a "Best of Breed" solution.
- Provides configuration verification and discovery in support of enterprise-wide configuration management.
- Improves data collection, analysis, and display capability.
- Reduces local effort required to maintain monitoring capability.

The ESM system is designed to provide visualization, event management, and reporting of fault and performance management information associated with the DSRC infrastructure and process operations. This will provide a single and consistent representation of data across the Centers at both the site and enterprise level. Below is a notional concept of operations for ESM that includes both classified and unclassified systems monitoring.

Currently, local DSRC staffs develop their own set of "point solution" monitoring tools, i.e., each Center uses separate Center-centric monitoring tools. The ESM project will leverage these "point solutions" and wrap them into the ESM solution, simplifying support and future enhancements by establishing a common architecture and structure. The system will also provide a single, consistent integration point for external systems requiring integration within the ESM framework. Test and production facilities will enable design and testing of the common tools in a controlled environment, prior to implementing them within the production network.

When fully deployed and integrated, Center operations staff will have the capability to monitor classified and unclassified systems and networks using physically separate management systems: Spectrum and eHealth. ESM will be used by the Consolidated Customer Assistance Center (CCAC), managers, systems administrators, and computer operators to provide more accurate information to staff and users, and will allow staff members to maintain high levels of service

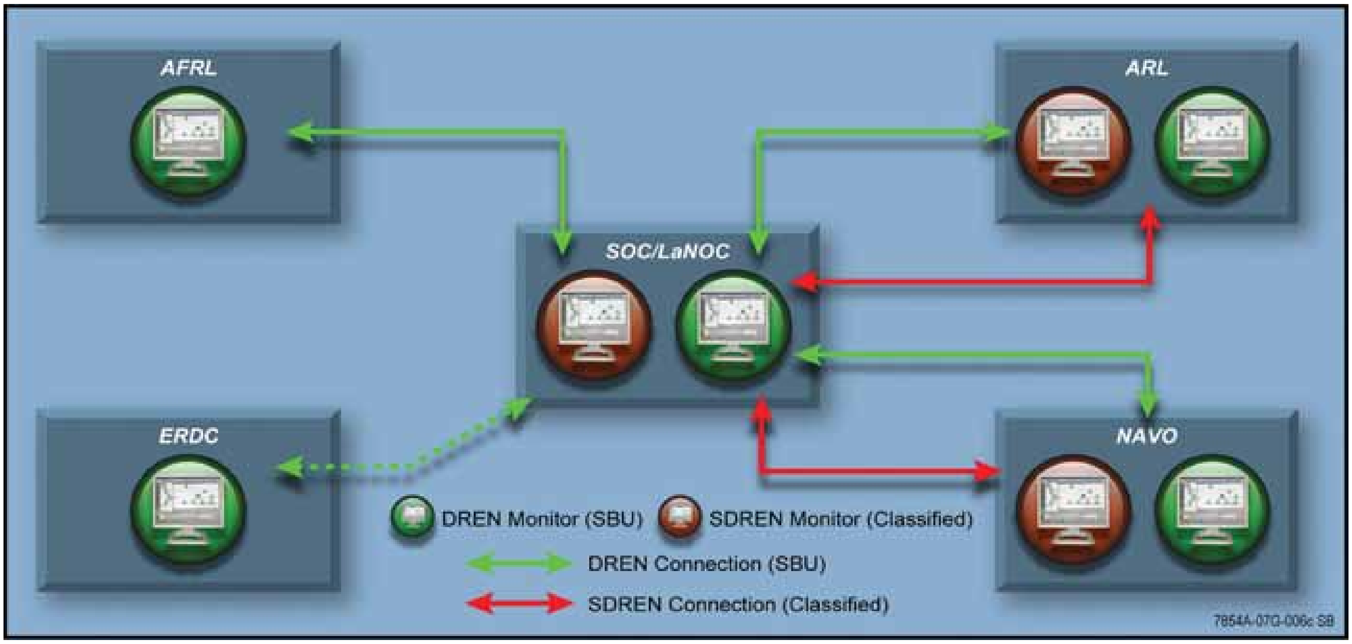

Figure 1. Notional concept of operations for ESM

availabilities by identifying potential problems quickly. ESM will also allow forecasting for systems growth and user impacts that may require manual intervention. In addition to the real-time benefits, an important and powerful capability of the ESM solution is the reporting capabilities.

ESM reporting capabilities enable the monitoring team to spot and understand negative performance trends before they noticeably affect each Center. At-A-Glance, or background, reports provide historical trends of an HPC asset's key performance indicators. Below are some report examples.

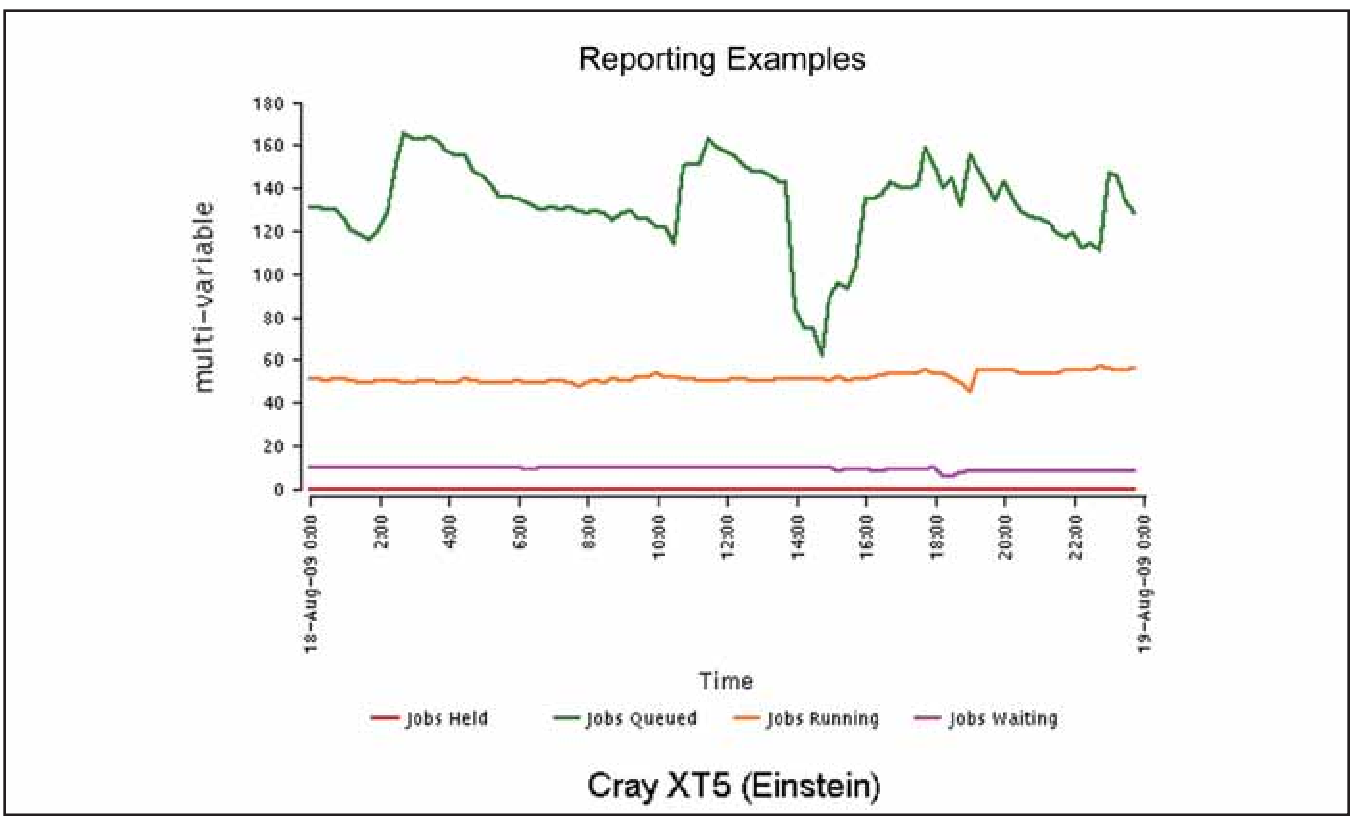

Figure 2. This report shows jobs held, jobs queued, jobs running, and jobs waiting over a 24-hr period from Einstein (CrayXT5)

The ESM project is one of the first major projects that will encompass the full HPCMP enterprise and will virtually touch all equipment within the HPC DoD Supercomputing Centers. Given the challenges facing an ESM implementation, the DoD HPC Program Office selected the Navy DSRC as Executive Agent for ESM. As such, the Navy DSRC has provided oversight and guidance on behalf of all DSRCs to Lockheed Martin in implementing the initial ESM framework they proposed. Thanks to the support provided by the Navy DSRC and the hard work of Lockheed Martin and its ESM partners, ESM is becoming a reality among the AFRL, ARL, ERDC, and Navy DSRCs.

"The ESM effort provides a unique opportunity for the DSRCs to work together toward a common approach for managing our HPC assets. It should improve our efficiency and effectiveness in delivering HPC services. As we establish common understanding and agreement on enterprise systems management approaches, I'm sure ESM will evolve over time to include new enterprise tools and services, beyond that initially proposed by Lockheed Martin," stated Tom Brown, Navy DSRC Associate Director and Executive Agent for ESM.

By providing leadership to the ESM effort, the Navy DSRC is making a major impact that will help the Centers achieve a higher degree of excellence in management of its HPC assets. This is the first cross-Center initiative of this magnitude since the Disaster Recovery project that was also led by the Navy over 10 years ago.

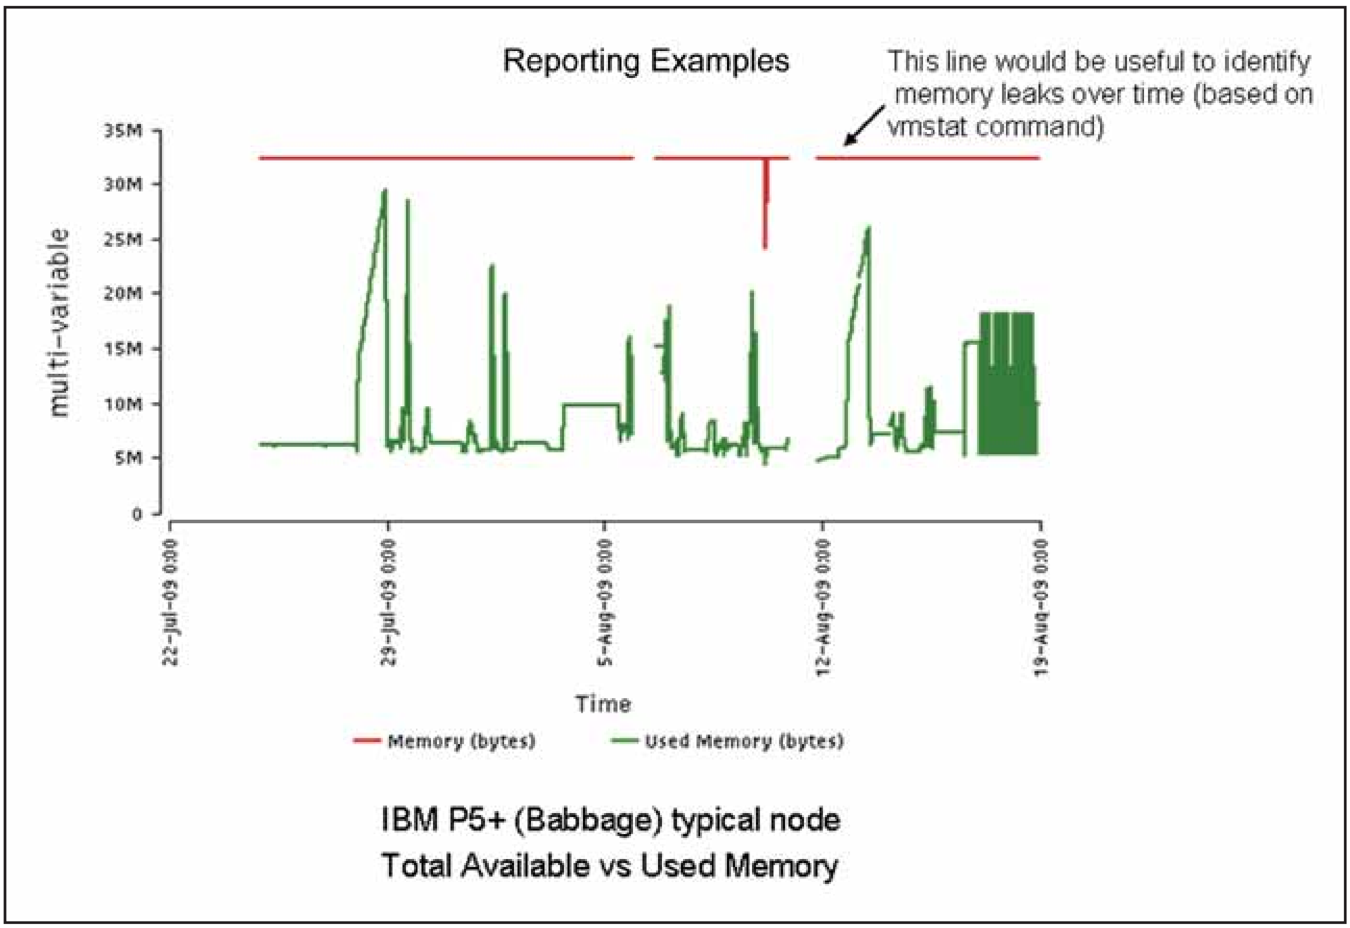

Figure 3. This report example shows Navy's HPC system Babbage's memory information. This metric identifies memory leaks over time

Data Analysis and Assessment Center

UDAAC Wins Two Department of Energy OASCR Awards

By Paul Adams, UDACC Visualization Team Lead

**"You did an incredible job isolating the vortex tubes … it looks great."
– Doug Dommermuth**

The Data Analysis and Assessment Center (DAAC) team members often meet with other members of the computational science community to discuss new technologies and discover new approaches to collaboration. One of the times for collaboration with other members of the scientific community occurred in June.

The DoD High Performance Computing Modernization Program (HPCMP) and the Department of Energy (DOE) held joint conferences at the same venue to allow participants to cross-pollinate their results and ideas. The DoD HPCMP organized their Users Group Conference, and the DOE Office of Advanced Scientific Computing Research (OASCR) organized their Scientific Discovery through Advanced Computing Program (SciDAC) Conference. As part of the SciDAC Conference, a visualization night was held in which participants could highlight their work. Two hundred peers judged a total of 29 submissions and selected 10 scientific visualizations to receive OASCR Awards. The winning submissions were recognized by receiving gold-colored statuettes. The Unclassified DAAC (UDAAC) team submitted three entries to the visualization contest and won two of the awards.

One of the award-winning submissions recounts the "Breaking Waves" simulation, which models how waves travel around ships. This is part of the Numerical Flow Analysis project led by Dr. Doug Dommermuth with visualization performed by the UDAAC over the past 5 years. The movie shows not only the improvement in the numerical simulation for that time period, but also shows the improvements in visualization techniques that help the researchers understand his data.

The second award-winning submission was for the "Multi-scale Multi-physics Simulation for Advanced Body Armor Designs" project. This simulation seeks to design better body armor, which is an essential part of force protection research. The project was a simulation of a 0.429 inch copper bullet impacting six layers of edge woven Kevlar tape at 1,070 ft/s. The simulation depicts yarn ejection at the free edges of the target, yarn pullout, and the highly flexible impact response of the fabric target, all of which are observed in impact experiments. The projectile residual velocity in the simulation shows good agreement with the corresponding experiment.

"ezVIZ performed quite well processing the large data sets" – Eric Farenthold

For more information about the DAAC, you can visit their Web site at https://visualization.hpc.mil/

For more information about the winners of the Visualization Night contest, you can visit Wired at http://www.wired.com/wiredscience/2009/08/visualizations/

DAAC
Data Analysis and Assessment Center

DoD HPCMP Hero Awards for 2009

Presented by Cray Henry at the DoD HPCMP 2009
Users Group Conference, San Diego, California,
June 17, 2009

Technical Excellence
Paul Bennett
ERDC DSRC

Technical Excellence
Danny Weddle
HPCMPO

Innovative Management
Rose Dykes
ERDC DSRC

Long Term Sustained
Margo Frommeyer
NRL-SS

DoD High Performance Computing Modernization Program Users Group Conference 2010

Chicago, Illinois, June 14-17, 2010

Renaissance Schaumburg Hotel & Convention Center

Modernization Program

DoD HPC

SOLVING THE HARD PROBLEMS

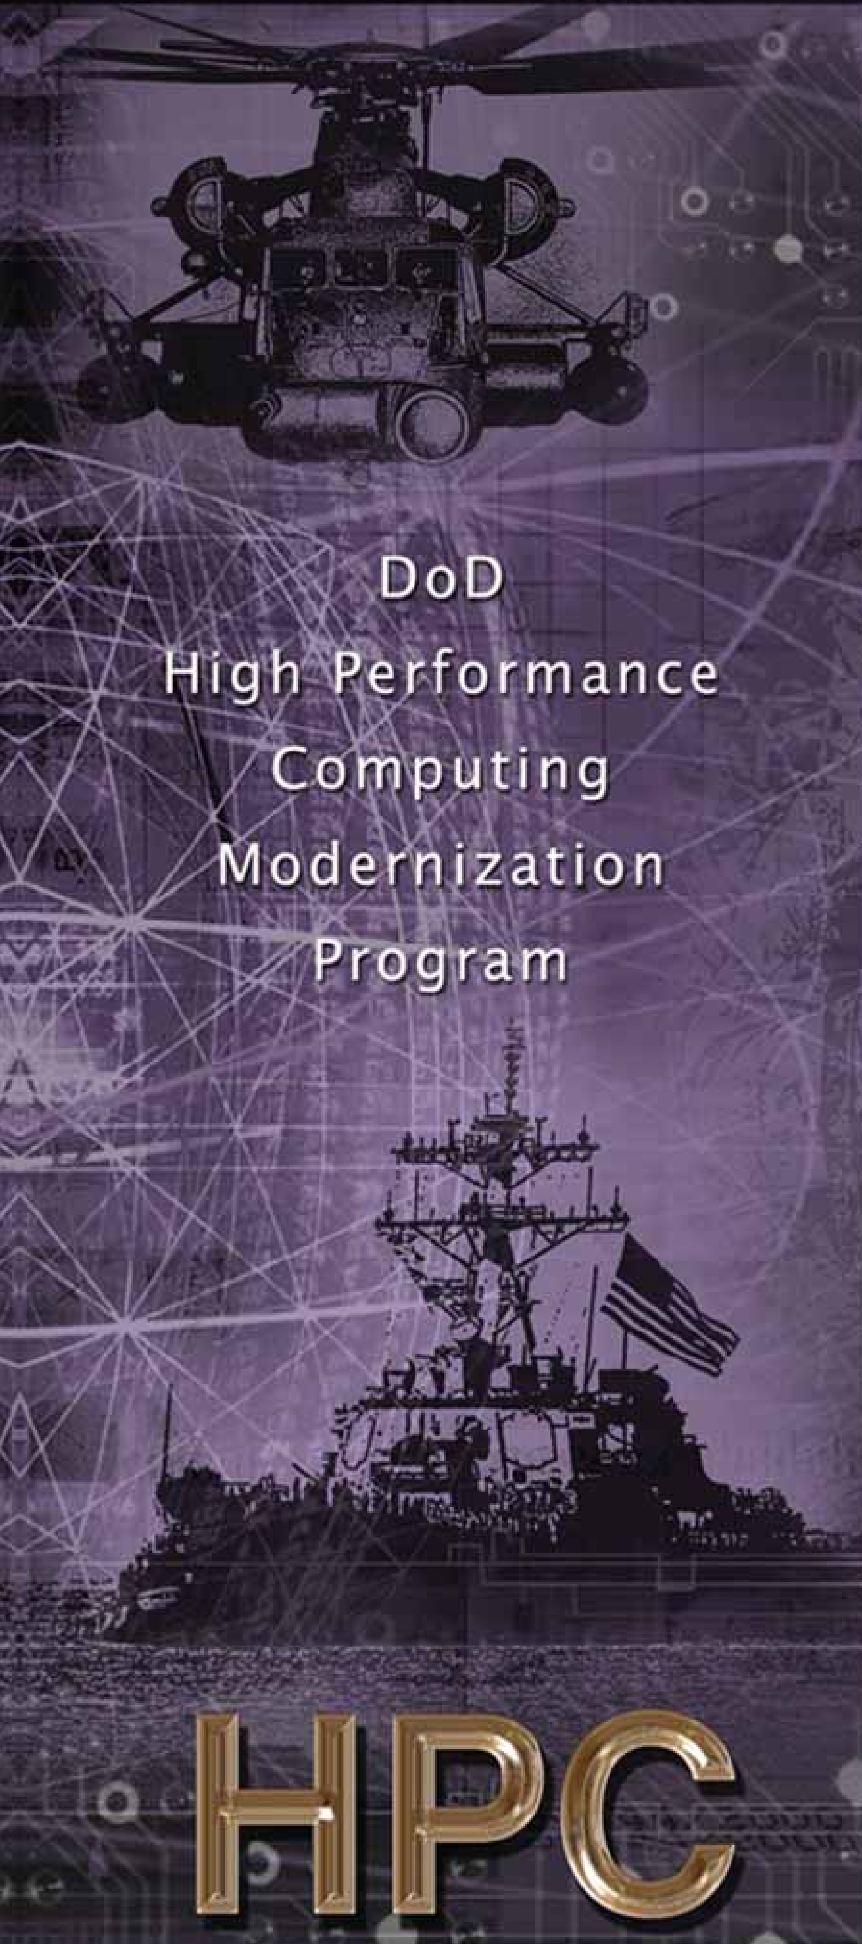

DoD High Performance Computing Modernization Program

HPC

SC09 Edition