HPCinsights

DoD High Performance Computing Modernization Program

Fall 2011

SC11 Edition

DoD Supercomputing Resource Centers
Networking/Security • Software Applications Support

SOLVING THE HARD PROBLEMS

AFRL • ARL • ERDC • MHPCC • NAVY

SUPERCOMPUTING FOR THE WARFIGHTER

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
Maria Zimmer
Chuck Abruzzino

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

ERDC DSRC, Vicksburg, MS
Rose J. Dykes

MHPCC DSRC, Maui, HI
Betty Duncan

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

HPCMPO, Lorton, VA
Deborah Schwartz
Denise O'Donnell
Leah Glick

MANAGING EDITOR
Rose J. Dykes, ERDC DSRC

DESIGN/LAYOUT
Betty Watson, ACE-IT

COVER DESIGN
Chandra "Pat" Caldwell, ACE-IT

Contents

DoD High Performance Computing Modernization Program – by Director – John E. West 1

HPC at Work

Tropical Cyclone Prediction Using COAMPS-TC 2
Characteristics of Bluff Body Stabilized Turbulent Premixed Flames 6
Faster Tracker: Simulating Satellite Tracking Using Parallel Computing 11
Optimization of Nanocrystalline Silicon Carbide Ceramics Through Atomistic Simulations 13

DoD Supercomputing Resource Centers

AFRL DSRC

From the Director's Desk – Frank Witzeman 16

ARL DSRC

From the Director's Desk – Dr. Raju Namburu 17
Energy Aware Scheduler Saves Resources 18
ARL Making Room for More Horsepower 19

ERDC DSRC

From the Director's Desk – Dr. Robert S. Maier 20

MHPCC DSRC

From the Director's Desk – David Morton 22
GNET: Visualizing Real-Time sFlow Data on the DREN at the MHPCC DSRC 23
HPCMP Portal Initiative 24
MHPCC DSRC Energy Efficient Computing Initiative 25

Navy DSRC

From the Director's Desk – Tom Dunn 27

DoD HPCMP Users Group Conference 2011 28

Announcements

DoD HPCMP Hero Awards for 2011 29
DoD HPCMP Users Group Conference 2012 29

About the Cover: U.S. Marines move to a firing line while conducting marksmanship training on the flight deck of USS New Orleans off the coast of southern California, September 4, 2011. The Marines are assigned to the 11th Marine Expeditionary Unit's maritime raid force.
U.S. Marine Corps photo by Cpl. Chad J. Pulliam

DoD High Performance Computing Modernization Program

By John E. West, Director

1 October was the start of the Federal Government's fiscal year and, in a break with recent history, it started off feeling like a time of significant change for our Program. I hope by the time you read this you've begun to find that this isn't a time of change so much as it is a time of renewal.

From a program management perspective, the most significant development this fiscal year is that management of the High Performance Computing Modernization Program (HPCMP) has moved out of the Office of the Secretary of Defense to the Department of the Army. From a user's perspective, however, this change should be of little immediate consequence. HPC services and expertise will continue to be delivered out of the various hosting organizations in the Army, Navy, and Air Force – just as they always have been. The Program will remain a joint activity of the Department of Defense, committed to maintaining its focus on the needs of the entire DoD Research, Development, Test, and Evaluation community.

With the transition to the Army also comes a change in leadership for the Program, and I am excited to have been selected as the next director of the HPCMP. I am also humbled by the remarkable legacy of the successes that Cray Henry and his predecessors in this role have delivered to our country.

HPC increases and focuses human insight and expertise with unmatched intensity on our most pressing challenges. I remain committed to the idea that supercomputing is unique among the instruments of human technical endeavor, serving (as Dan Reed put it) as the universal intellectual amplifier.

Put simply, supercomputing makes the world a better place.

Within the Department of Defense, supercomputing plays a key role in accelerating the transition of new capabilities into the hands of our fighting forces. Supercomputing expertise and technologies also enable the DoD to reduce and manage risk:

- in research, where HPC enables DoD to explore new theories and evaluate them well beyond what is financially possible using experiment alone;
- in acquisition, through the use of validated applications in design and testing;
- and in operations, where real-time calculations produce just-in-time information for decision makers on the battlefield.

From vehicles that more completely protect their occupants from the effects of improvised explosive devices to aircraft that fly farther and faster, the DoD HPCMP has an incredible legacy of success. It is our job to build upon that success in order to continue to solve the DoD's most demanding problems.

DOD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM

John E. West
HPCMP Directorr

Tropical Cyclone Prediction Using COAMPS-TC

By J. D. Doyle and R. Hodur, SAIC, Monterey, California; S. Chen, J. Moskaitis, H. Jin, Y. Jin, and A. Reinecke, Naval Research Laboratory, Monterey, California; P. Black, SAIC; J. Cummings, Naval Research Laboratory, Stennis Space Center, Mississippi; E. Hendricks, T. Holt, C.-S. Liou, M. Peng, C. Reynolds, K. Sashegyi, J. Schmidt, and S. Wang, Naval Research Laboratory, Monterey California

A dramatic scenario played out recently during August 2011 as Hurricane Irene threatened many communities along the U.S. Eastern Seaboard, from Florida to New England. Basic questions such as where Irene would track and how strong it would become had profound implications for the millions of people in its path and billions of dollars that were vulnerable. The potential impact of tropical cyclones on military operations can also be enormous. An extreme example is the infamous Typhoon Cobra, also known as Halsey's Typhoon after Admiral William Halsey, which struck the Navy's Pacific Fleet in December 1944 during World War II. Three destroyers were lost, and a total of 790 sailors perished. More recently during Irene, the decision to "sortie" Navy assets from Norfolk, Virginia, and other ports along the Eastern Seaboard many days in advance of the storm was critically dependent on forecasts of Irene's track, intensity (maximum sustained wind speed at the surface), and storm structure (such as the size of the storm or radius of key wind speed thresholds). Similarly, multiple sorties of the Navy Pacific Fleet in the Philippine Sea were needed to avoid Typhoon Nanmadol, which exhibited erratic movement and was poorly forecasted.

The demand for more accurate hurricane forecasts with longer lead times is greater than eve due to the enormous economic and societal impact. There has been spectacular improvement of tropical cyclone[1] (TC) track prediction; a 3-day hurricane track forecast today is as skillful as a 1-day forecast was just 30 years ago. However, there has been almost no progress in improving TC intensity and structure forecasts due to a variety of reasons ranging from a lack of critical observations under high wind conditions and in the TC environment to inaccurate representations of TC physical processes in numerical weather prediction (NWP) models.

A new version of the Coupled Ocean/Atmosphere Mesoscale Prediction System (COAMPS®) has been developed by the Naval Research Laboratory (NRL) in Monterey, California, and is designed specifically for forecasting tropical cyclones. This COAMPS-TC system is comprised of data quality control, analysis, initialization, and forecast model subcomponents. The Navy Variational Data Assimilation System (NAVDAS) is used to blend observations of winds, temperature, moisture, and pressure from a plethora of sources such as radiosondes, pilot balloons, satellites, surface measurements, ships, buoys, and aircraft. Enhancements to the NAVDAS system for COAMPS-TC include the addition of synthetic observations that define the TC structure and intensity (based on the TC reports in real-time from the National Hurricane Center (NHC) and the Joint Typhoon Warning Center (JTWC)). Also, as part of the TC analysis procedure, the pre-existing circulation in the COAMPS-TC first guess fields is relocated to allow for an accurate representation of the TC position during the analysis. Following this step, the analyzed fields are initialized to reduce the generation of spurious, high-frequency atmospheric gravity waves. The sea surface temperature is analyzed directly on the model computational grid using the Navy Coastal Ocean Data Assimilation (NCODA) system, which makes use of all available satellite, ship, float, and buoy observations. Both the NCODA and NAVDAS systems are applied using a data assimilation cycle in which the first guess from the analysis is derived from the previous short-term forecast.

The COAMPS-TC atmospheric model uses the nonhydrostatic and compressible form of the dynamics and has prognostic variables for the three components of the wind (two horizontal wind components and the vertical wind), the perturbation pressure, potential temperature, water vapor, cloud droplets, raindrops, ice crystals, snowflakes, graupel, and turbulent kinetic energy. Physical parameterizations include representations of cloud microphysical processes, convection, radiation, boundary layer processes, and surface layer fluxes. The COAMPS-TC model contains a representation of dissipative heating near the ocean surface, which has been found to be important for tropical cyclone intensity forecasts. The model also contains an optional hybrid time-differencing scheme that can be selected at run-time that allows the scalars to be computed on a forward time-step. The COAMPS-TC system also contains a flexible nesting design that has proven useful when more than one storm is present in a basin at a given time as well as special options for moving nested grid families that independently follow individual tropical cyclone centers of interest.

The COAMPS-TC system has the capability to operate in a fully coupled air-sea interaction mode. The atmospheric module within COAMPS-TC is coupled to the NRL-developed Navy Coastal Ocean Model (NCOM) to represent critical air-sea interaction processes. The COAMPS-TC system has an option to predict ocean surface waves and the interactions between the atmosphere, ocean circulation, and waves using the Simulating WAves Nearshore (SWAN) model. A sea spray parameterization can be used to represent the injection of droplets into the atmospheric boundary layer due to ocean surface wave breaking and shearing.

[1] Strong tropical cyclones are known as hurricanes in the Atlantic and eastern Pacific, and typhoons in the Western Pacific.

COAMPS-TC has been tested in real-time in both coupled and uncoupled modes over the past several tropical cyclone seasons in the Pacific and Atlantic basins. These real-time tests have been conducted in conjunction with the National Oceanic and Atmospheric Administration (NOAA)-sponsored Hurricane Forecast Improvement Project (HFIP), which is focused on accelerating the improvement in hurricane intensity forecasts. In these real-time applications, the atmospheric portion of the COAMPS-TC system makes use of horizontally nested grids corresponding to resolutions of 45, 15, and 5 km. The 15- and 5-km resolution meshes track the TC center, which enables the TC convection to be explicitly resolved and more realistically represented on the finest mesh in an efficient manner. The forecasts make use of Navy and NOAA global models for lateral boundary conditions. The model is typically run four times daily for the W. Atlantic, E. Pacific, and W. Pacific regions and is triggered by the NHC and JTWC warning message (which contains observational estimates of the storm position and intensity) when a storm reaches a 30-knot intensity. The forecasts are routinely disseminated in real-time to NHC, JTWC, and HFIP researchers. The forecast graphics are also available in real-time at http://www.nrlmry.navy.mil/coamps-web/web/tc.

The development of COAMPS-TC as well as new capabilities for the operational global and mesoscale systems has been made possible through our HPC Challenge Project entitled *Tropical Cyclone Track and Intensity Predictability*. The computations have been performed primarily on the DoD High Performance Computing Modernization Program (HPCMP) Navy DoD Supercomputing Resource Center (DSRC) at Stennis Space Center, Mississippi, on the Cray-XT5 (*Einstein*) and the IBM Power 6 (*Davinci*) systems. We have also made use of the Air Force Research Laboratory DSRC Altix (*Hawk*) and the Engineer Research and Development Center (ERDC) DSRC SGI Altix ICE (*Diamond*) and Cray XE6 (*Garnet*). The real-time COAMPS-TC system is run on 160 cores on a Cray-XT5 at the Navy DSRC. The COAMPS-TC forecasts, which extend out to 5 days, take approximately 2 hours of wall clock time to complete. Experiments to evaluate the skill of numerical models such as COAMPS-TC often require several hundred cases to achieve statistically meaningful results, which underscores the necessity of extensive DoD HPC resources.

Real-time COAMPS-TC forecasts have been conducted using the HPC platforms in the past 3 years. An example of the intensity forecast skill of COAMPS-TC for a large number of cases (more than 200 cases at the 24-hr forecast time) in the W. Atlantic region is shown in Figure 1. The COAMPS-TC model was the best numerical prediction model for hurricane intensity during the 2010 season in the Atlantic basin for the 36- to 60-hr forecasts, which is a critical time period for forecasters and DoD decision makers. Other numerical models included in this analysis are operational models run by NOAA (HWRF, GFDL), an experimental NOAA model (HWRF-X), and the Navy's current operational limited area model (GFDN). This promising performance is a result of a large effort devoted to developing and improving COAMPS-TC over the past 3 years. Without the support of the HPCMP supercomputers, this excellent level of skill could not have been achieved.

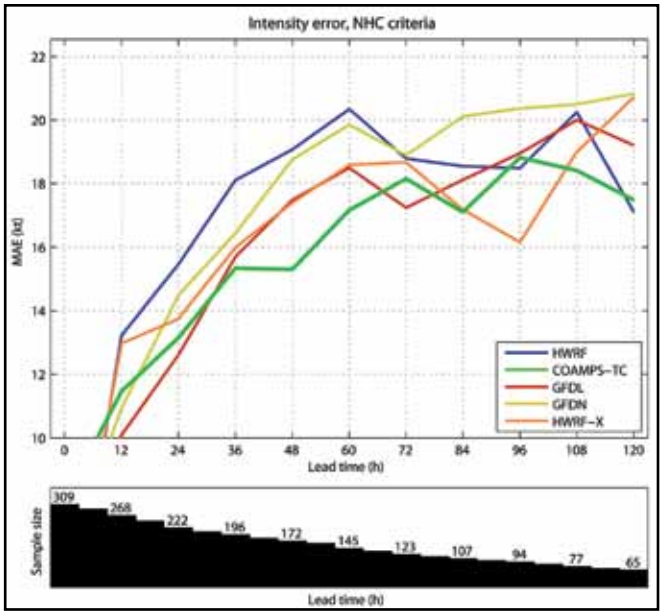

Figure 1. Wind speed mean absolute error (MAE) (knots) as a function of forecast time for the 2010 season in the Atlantic basin for a homogeneous statistical sample. Numerical models included in this analysis are Navy COAMPS-TC, operational models run by NOAA (HWRF, GFDL), experimental NOAA model (HWRF-X), and Navy current operational limited area model (GFDN). Number of cases are shown at bottom

An example of a real-time COAMPS-TC forecast for the recent Hurricane Irene is shown in Figure 2. The composite National Weather Service radar reflectivity is shown in the top panel near the time of landfall in North Carolina at 1148 UTC 27 August 2011, and the COAMPS-TC predicted radar reflectivity at 36 hr valid at 1200 UTC is shown in the bottom panel. The COAMPS-TC forecast shown in Figure 2 is for the model second grid mesh (15-km horizontal resolution). The model prediction was remarkably accurate in not only the track and eventual landfall location, but also quite accurate with regard to the storm intensity, structure, and size, an especially important characteristic of this particular storm in such close proximity to the U.S. East Coast. One noteworthy aspect of Irene was its large size, with tropical storm force winds (34 knot) radially extending outward from the eye for nearly 200 miles. The large size of Irene is also apparent in the observed radar reflectivity in Figure 2. The COAMPS-TC prediction captures the large size of the storm, as well as the precipitation shield that is present to the north and northeast of the storm. This large shield of heavy precipitation caused severe river flooding as it slowly moved north through the mid-Atlantic and Northeast U.S. The simulated radar reflectivity for the COAMPS-TC grid mesh 3 (5-km horizontal resolution), shown in Figure 3, illustrates the capability of the model to capture the finer scale features, such as the eye wall and rain bands, in

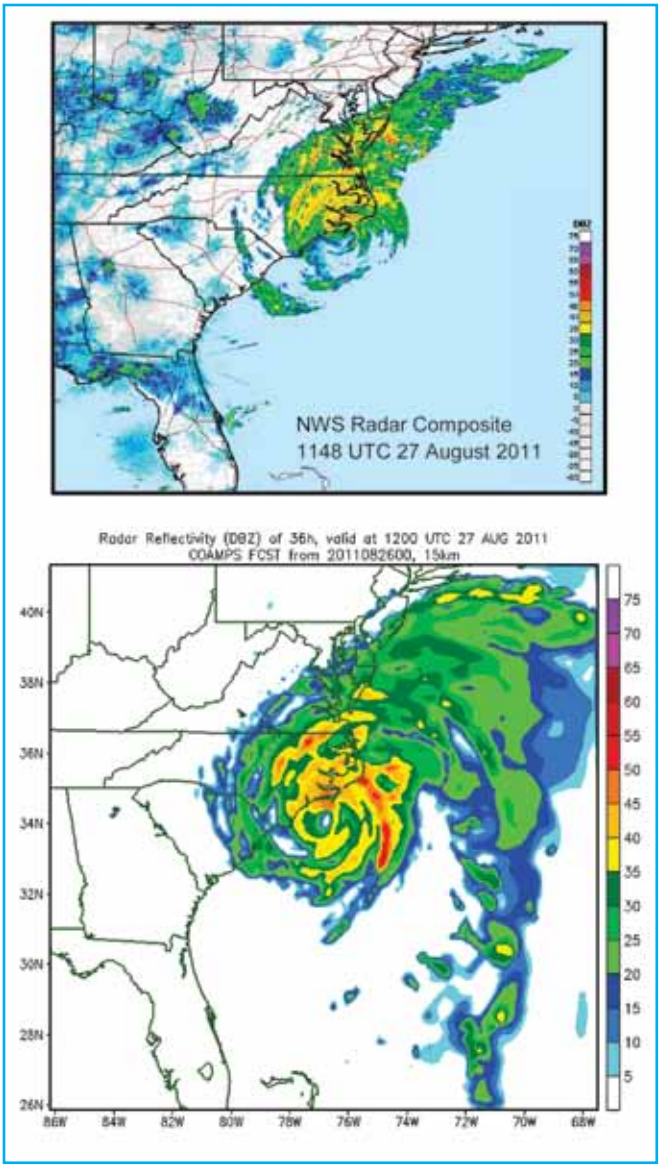

Figure 2. NWS composite radar reflectivity valid at 1148 UTC 27 August 2011 (top panel) and COAMPS-TC 36-hr forecast radar reflectivity performed in real-time and valid at 1200 UTC 27 August (bottom panel) for Hurricane Irene. COAMPS-TC reflectivity shown for second grid mesh, which has horizontal resolution of 15 km

generally good agreement with the observed reflectivity. It should be noted that the asymmetric structure of the observed precipitation with greater coverage and intensity to the north is captured by the model forecast (Figure 2).

Overall, the Navy's COAMPS-TC real-time intensity predictions of Hurricane Irene outperformed other leading operational governmental forecast models, as shown in Figure 4. All of the available models except for COAMPS-TC had a tendency to overintensify Irene often by a full storm category or more. These real-time COAMPS-TC forecasts were used by forecasters at the National Hurricane Center as part of an experimental HFIP multimodel ensemble. The COAMPS-TC consistently provided remarkably accurate real-time intensity forecasts during the period 23-28 August 2011, when critical

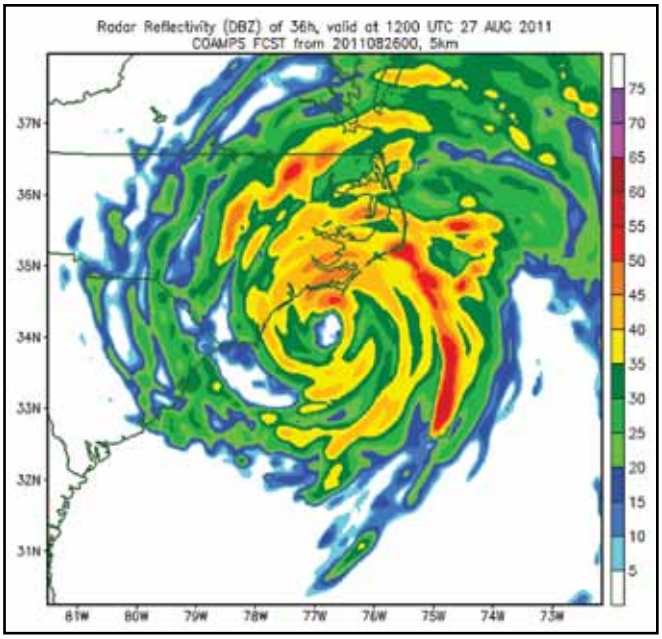

Figure 3. COAMPS-TC 36-hr forecast radar reflectivity valid at 1200 UTC 27 August for third grid mesh, which has horizontal resolution of 5 km

decisions were made by forecasters and emergency managers including evacuations.

While research is ongoing to improve deterministic atmospheric forecasts through advancements to the forecast model and more accurate estimates of the initial state, simultaneously there has been interest in obtaining probabilistic information derived from ensemble forecasts. An ensemble of forecasts from equally plausible initial

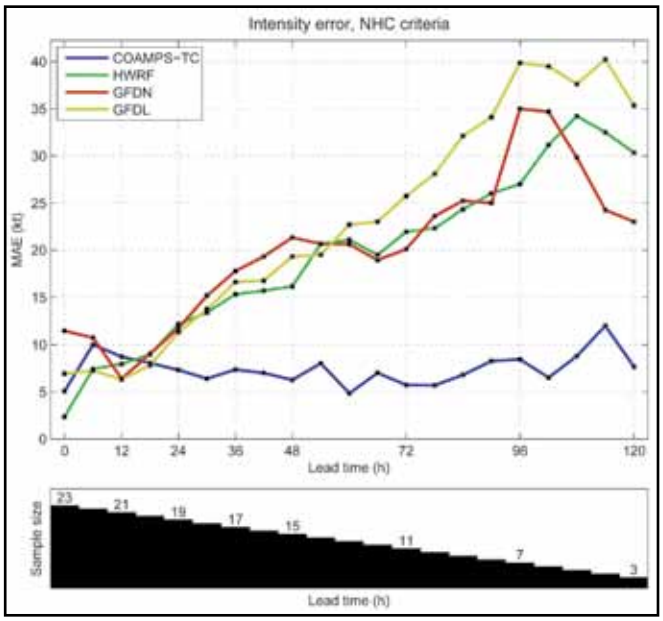

Figure 4. Wind speed mean absolute error (MAE) (knots) as function of forecast time for Hurricane Irene for homogeneous statistical sample. Numerical models included in this analysis are Navy COAMPS-TC, operational models run by NOAA (HWRF, GFDL), and Navy current operational limited area model (GFDN). Number of cases are shown at bottom

states and model formulations offers a computationally feasible way of addressing inevitable forecast uncertainties, offering improved forecasts through ensemble statistics such as mean quantities, as well as quantitative estimates of forecast error and variance. Although the concept of ensemble modeling is relatively simple, the performance of an ensemble forecast system is sensitive to the basic ensemble architecture. At the Naval Research Laboratory, we are designing new ensemble methods for both the global and mesoscale atmospheric forecast systems. Because of the high computational demands associated with ensemble development and verification (especially when one is interested in severe or rare events), exceptional computational resources are necessary to perform this research.

A new COAMPS-TC ensemble system that is capable of providing probabilistic forecasts of TC track, intensity, and structure has been developed by scientists at NRL in Monterey, California. This system makes use of a community-based Data Assimilation Research Testbed (DART) capability developed at the National Center for Atmospheric Research, which includes various options for Ensemble Kalman Filter (EnKF) data assimilation. The COAMPS-TC DART system constitutes a next-generation data assimilation system for tropical cyclones that uses flow-dependent statistics from the ensemble to assimilate observational information on the mesoscale. A real-time COAMPS-TC ensemble system was run in a demonstration mode in 2011 for the W. Atlantic and W. Pacific regions. Ten-member forecasts were performed twice daily to 5 days using three nested grids with horizontal resolutions of 45, 15, and 5 km. The data assimilation cycle, which was run every 6 hr, used 80 members. Examples of probabilistic products for Hurricane Irene are shown in Figure 5 for both track (top panel) and intensity (bottom panel). This is a real-time forecast initialized at 1200 UTC 23 August, which is 4 days prior to landfall. The probabilistic track product shows the TC position from the individual ensemble members every 24 hr and ellipses that encompass the 1/3 and 2/3 of the ensemble member forecast positions. Note that the observed landfall location of the eye (see Figure 2) was within the ensemble distribution, although the ensemble landfall was approximately 12 hr later than observed. The probabilistic intensity product (lower panel) shows a considerable spread among the members, particularly beyond 84 hr, just prior to landfall. These products can be extremely valuable to assess the uncertainty in both track and intensity forecasts, and NRL is currently developing these capabilities and products further.

The COAMPS-TC system was also run in a fully coupled mode, interactive with NCOM, during the Office of Naval Research-sponsored Interaction of Typhoon and Ocean Project (ITOP) during the summer and fall of 2010. These forecasts were run on the Cray XT5 at the Navy DSRC. An example of a fully coupled COAMPS-TC forecast for Typhoon Fanapi is shown in Figure 6. The NCOM ocean model was applied with a 10-km horizontal resolution in

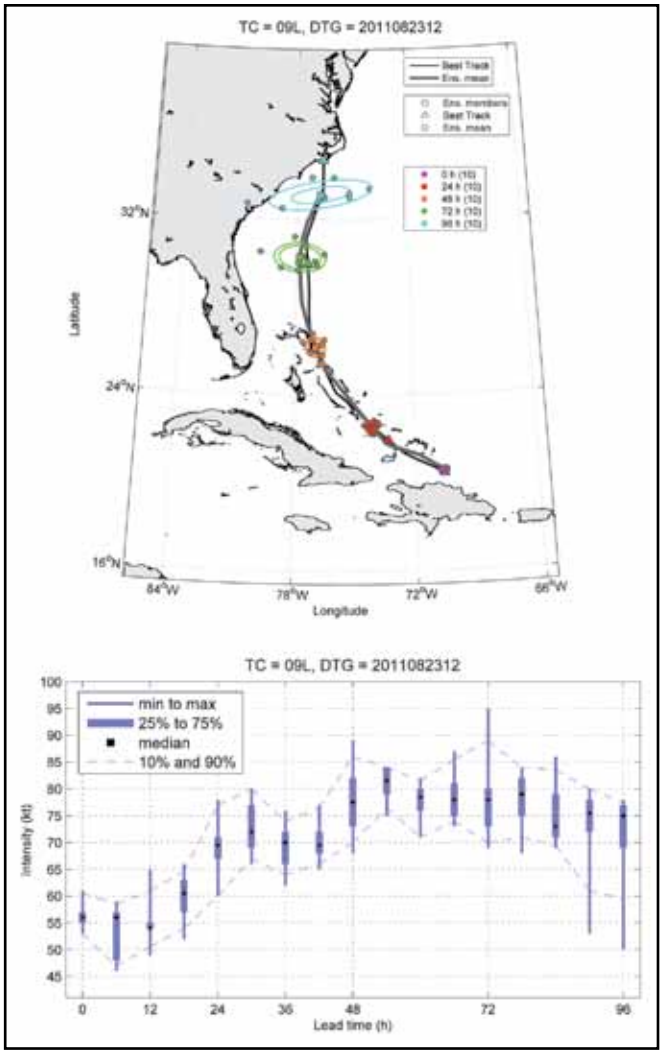

Figure 5. Probabilistic products from COAMPS-TC ensemble for Hurricane Irene corresponding to track (top panel) and hurricane intensity (bottom panel). Real-time forecast initialized at 1200 UTC 23 August, which is approximately 4 days prior to landfall. Probabilistic track product shows TC position from individual ensemble members every 24 hr and ellipses that encompass 1/3 and 2/3 ensemble distributions. Intensity (knots) shown as function of forecast lead time (hours) and median, minimum, maximum, and 10% and 90% distributions shown as denoted by legend

this example. The COAMPS-TC predicted track (red) from a 90-hr real-time forecast valid at 0600 UTC 19 September 2010 is quite close to the observed or best track (black). The sea surface temperature, shown in color shading, indicates significant cooling was predicted by COAMPS-TC during the passage of Fanapi, due to enhanced mixing by the strong near-surface winds of the typhoon. The predicted cooling of the sea surface temperatures of 2-4°C is in agreement with estimates from in situ and remote sensing observations in this region.

A joint Navy/Air Force Hurricane Hunter program is in its first demonstration phase this year with Airborne Expendable Bathythermographs (AXBTs) being deployed from WC-130J hurricane reconnaissance aircraft in order to improve the initialization and validation of coupled models such as COAMPS-TC. Over 30 AXBTs were

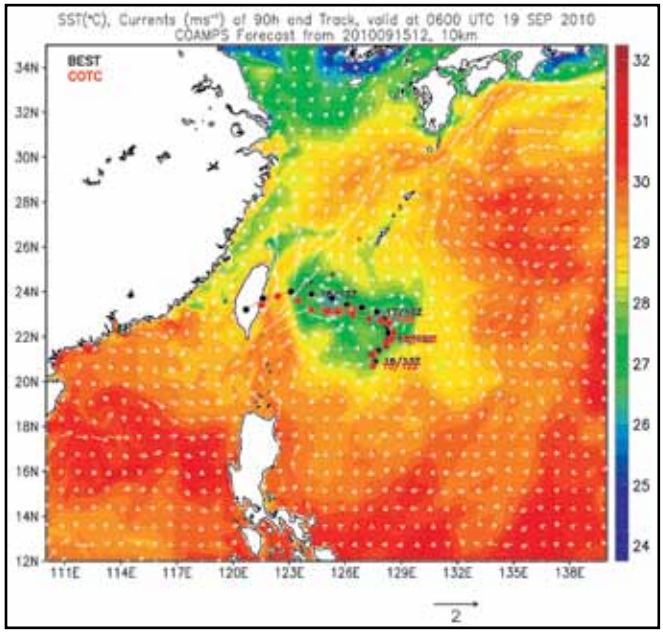

Figure 6. COAMPS-TC predicted track (red) for Typhoon Fanapi from 90-hr real-time forecast valid at 0600 UTC 19 September 2010 and observed best track (black). Sea surface temperature (°C) shown in color shading and indicates significant cooling during passage of Fanapi. Surface currents shown by white vectors

deployed from Air Force Hurricane Hunter aircraft in Irene as it approached landfall on Cape Hatteras. These new observations may help document the existence of ocean mixing (and cooling of the sea surface) along the storm track and in coastal regions that may have prevented Irene from further intensification.

Prediction of tropical cyclone track and particularly intensity remains one of the greatest challenges in meteorology today. The results of this research highlight the promise of high-resolution deterministic and ensemble-based approaches for tropical cyclone prediction using COAMPS-TC. The development and testing of COAMPS-TC has required large computational resources and has only been possible through the support of the DoD HPC Challenge Program, which has already led to significant improvements in the Navy's effort to improve tropical cyclone intensity prediction (COAMPS-TC was the best numerical model for intensity in the Atlantic during the recent 2010 season). In addition, this research will lead to new capabilities in the form of mesoscale TC ensemble forecasts, providing the Navy with probabilistic forecasts of tropical cyclone intensity and structure for the first time. It is also expected this research will help motivate new field campaigns that focus on the key measurements needed to further advance our understanding of the convective structure and dynamics of these systems as well as provide validation.

Acknowledgments

We acknowledge the support of the Office of Naval Research (ONR) Program Element (PE) 0602435N and PMW-120 PE 0603207N, as well as the National Oceanic and Atmospheric Administration (NOAA)-sponsored Hurricane Forecast Improvement Project (HFIP). We also appreciate support for computational resources through a grant of Department of Defense High Performance Computing Modernization Program time from the Navy DoD Supercomputing Resource Center at Stennis Space Center, Mississippi. COAMPS® is a registered trademark of the Naval Research Laboratory.

Characteristics of Bluff Body Stabilized Turbulent Premixed Flames

By Alejandro M. Briones, University of Dayton Research Institute, Dayton, Ohio; Balu Sekar, Air Force Research Laboratory, Turbine Engine Division, Combustion Branch; and Hugh Thornburg, High Performance Technologies, Inc., Reston, Virginia

Introduction

Requirements for rapid increase in thrust for takeoff and climb calls for additional thrust-producing devices. The afterburner concept is an auxiliary burner located behind the turbine section and forward of the exhaust nozzle that operates by injecting fuel to the hot exhaust leading to further combustion and extra thrust.

In such a system, the flames are typically stabilized by using an array of bluff body flameholders that create a recirculation zone. In a flameholder stabilized flame scenario, the flameholders are arranged in a single plane perpendicular to the flow direction and spaced either regularly or irregularly in either lateral dimension. The flameholders provide robust fluid recirculation zones that allow turbulent flames to uniformly attach and spread across the duct. The combustion products exit through a converging/diverging nozzle with extensive film cooling and a variable throat area located downstream of the afterburner exit. The afterburner may experience two types of instabilities: static and dynamic. The static stability refers to the ability of the flameholders to sustain a flame without blowing out. The dynamic stability refers to the unsteady character of the flame and often occurs at discrete frequencies spontaneously excited by feedback between the unsteady heat release rate and generally one of the natural acoustic modes of the combustor. The high-frequency dynamic (combustion) instability (i.e., 120–600 Hz) is named screech and is attributed to a combination of factors including flameholder geometry, fuel spray injection sites, blockage, non-uniformity of fuel/air ratio, evaporation rates, and ignition process. A canister liner is typically

used in afterburners to reduce screech in order to avoid deterioration and failure of the engine. The low-frequency dynamic instability (i.e., 50–120 Hz) is named rumble and is coupled with the fuel and air supplies and its interaction with the unsteady flow field. Due to the destructive nature of screech and rumble, considerable efforts from engine manufacturers are oriented to understand their foundation for further reduction and suppression.

The fundamental understanding of the flame/flow field dynamics in such an environment is extremely difficult. Therefore, insight into the fundamental mechanisms responsible for afterburner's static and dynamic instabilities can be obtained from single flameholder studies with prevaporized fuel and premixed mixture in order to isolate the turbulence-chemistry interactions from the complex physical processes. The basic structure of the flow field generated by a flameholder in non-reacting and reacting conditions is well-known. The flameholder generates a flow field composed by boundary layers, separated shear layers, and a wake [1]. There are two hydrodynamic instabilities associated with the shear layer and wake. The shear layer or Kelvin-Helmholtz (KH) instability is a convective instability related to the amplification of disturbances, leading to vortex roll-up and symmetric paring of the separated shear layers. The wake or Bérnárd/von Kármán (BvK) instability leads to an asymmetric shedding of vortices from the opposite side of the flameholder and sinuous wake structure.

The purpose of this investigation is to enhance the understanding bluff body stabilized turbulent flames. Computationl Fluid Dynamics (CFD) has earned a key role in the design and development of high performance gas turbine combustion systems both as a pretest analysis tool to predict static and dynamic instabilities, as well as a post test analysis methodology to propose modifications to address instabilities that are identified during development testing and employed in this work. Since the non-reacting and reacting flow past a flameholder in a typical combustor is unsteady, modeling of combustion needs to be truly an unsteady process that requires a comprehensive numerical model for accurate prediction of flame structure, propagation characteristics, and ignition/extinction phenomena, applicable over a wide range of high Reynolds number (Re) operating conditions. Therefore, the specific objectives of this investigation are as follows: (1) to examine the reliability of three-dimensional URANS and LES for predicting the von Kármán vortex street Strouhal number (St_{vK}) of reacting two-dimensional flameholder geometries (for example, square and rectangular); (2) to compare the effect of various URANS models on predicting St_{vK}; (3) to examine the effect of aspect ratio (AR) on St_{vK} for rectangular flameholders immersed in non-reacting and reacting flows; and (4) to compare the various flow fields obtained with the several flameholder geometries. This research will enhance the understanding of the shear layer and wake zone dynamics of flames anchored to a bluff body.

Numerical Approach

The three- and two-dimensional planar governing equations of continuity, momentum, energy, species, and turbulence are solved using the coupled pressure-based solver of FLUENT [2]. The temperature- and species-dependent thermodynamic and transport properties are given in [3]. Turbulence is modeled using either the Realizable k-ε URANS [4], SST k-ω URANS [5], or the constant Smagorinsky-Lilly Large Eddy Simulation (LES) formulation [2]. Since the flow is incompressible, the viscous heating terms and kinetic energy are neglected in the energy equation. Both the standard wall functions and the enhanced wall treatment are used to model and resolve the viscous laminar sublayer, respectively. The governing equations are discretized using a second-order upwind scheme.

Results And Discussion

The following sections discuss validation of our numerical models and the flow characteristic of square and rectangular flameholders:

A. Validation of Numerical Model and Grid Independence.

The Realizable k-ε and k-ω RANS, and LES models have been validated previously [6, 7, 8]. The LES model and reduced mechanisms have also been validated in triangular and v-gutter bluff body stabilized premixed flames in terms of velocity, temperature, and product species profiles [9]. The model was able to predict comparable flame profiles obtained in experiments for stable, dynamically unstable, and blowout conditions [8, 9]. Additional validation is presented in [10].

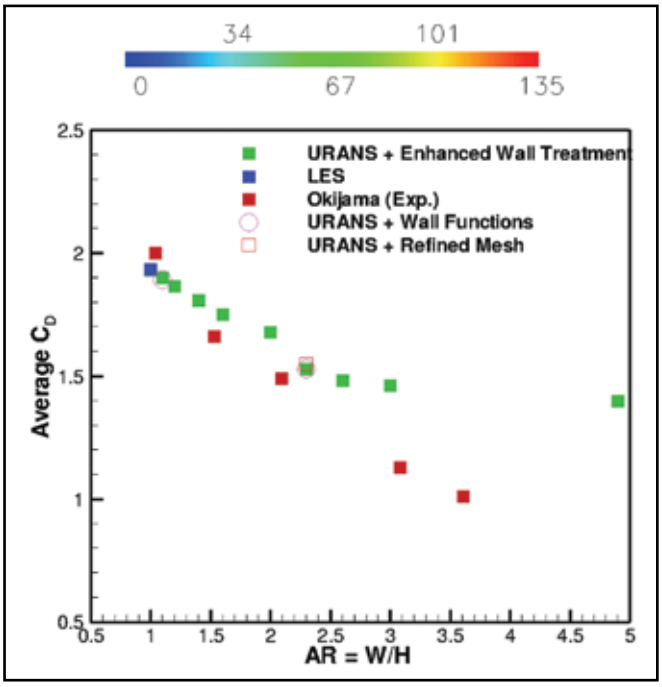

Figure 1. Average drag coefficient (CD) as a function of aspect ratio (AR = W/H) for rectangular cylinders. The experimental measurements are taken from Okajima [11]

For a parametric study involving turbulent reacting flows, it is computationally expensive to conduct grid independence studies. However, two meshes were used to model the cases corresponding to 3d–AR2.3–kε–c–125 and 3d–Triang–kε–h–125, respectively, in order to demonstrate that the numerical results were independent of mesh size [10].

B. Square and Rectangular Flameholders. We now discuss the characteristics of the various square and rectangular flameholders.

i. Reacting Flow

Figure 2, Figure 3, and Figure 4 present the instantaneous flame structures for AR = 1.0 (2d–AR1.0–kε–h–125; AR = 2.0 (2d–AR2.0–kε–h–125; and AR = 2.3 (2d–AR2.3–kε–h–125, respectively [10], in terms of heat release rate, temperature, and spanwise vorticity contours. The approaching flow velocity is 52.1 m/s. The flow velocity decreases as it approaches the flameholder. The flow is stagnant at the front of the flameholder. The pressure decreases from the center of the flameholder foreface towards the leading edges. The flow separates at the leading edges, and a recirculating bubble is formed in both the upper and lower surfaces of the flameholder. Due to thermal expansion, the flow also accelerates at a streamwise direction located near the flameholder's leading edge. The temperature increases from 300 K to nearly 2500 K, corresponding to T_b/T_u = 8.33. As mentioned before, previous investigations have shown that with increasing T_b/T_u, the BvK instability is inhibited and KH instability is dominant [12,13,14]. For AR = 1.0 the flame is no longer fluctuating, and both BvK and KH instabilities are suppressed with the large T_b/T_u. Therefore, St_{vK} is zero and any St related to the shear layers are also zero. This type of flame is ideal for afterburners since there is neither static nor dynamic instability. Moreover, the heat release rate contours indicate that the flame is attached to the flameholder's leading edge, and the high temperature region is found to cover both the flameholder's upper and lower surfaces and extend all the way to the wake and downstream locations. The preheat zone is thin near the flameholder's leading edge and thickens further downstream. The spanwise vorticity contour is symmetric with negative and positive sign in the upper and lower shear layers, respectively. The von Kármán vortices are suppressed by a combination of flame-induced baroclinic torque, gas expansion, and vortex diffusion [15,16,17]. This result is in agreement with that of Erickson et al. [12], Nair and Lieuwin [13], and Kiel et al. [14] who stated that large T_b/T_u, suppresses von Kármán vortex shedding. Now,

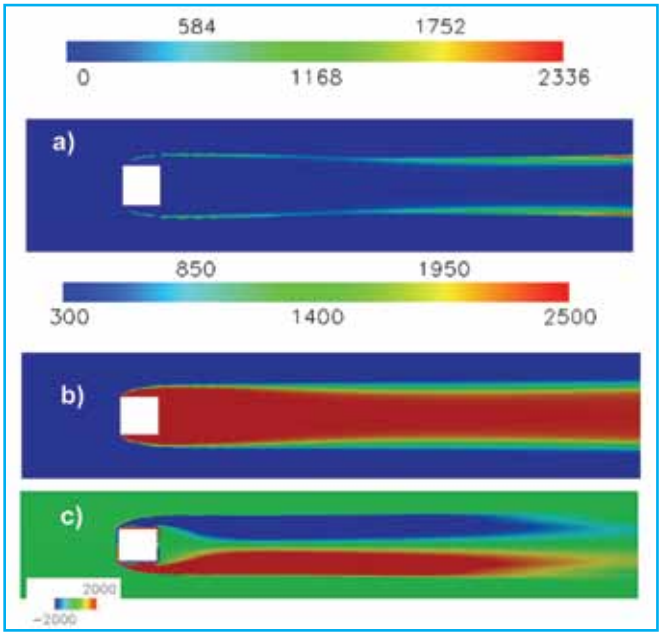

Figure 2. Instantaneous (a) heat release rate, (b) temperature, and (c) spanwise vorticity for the square flameholder (AR = 1.0). The units are in m/s, W/m^3, K, and s^{-1}, respectively

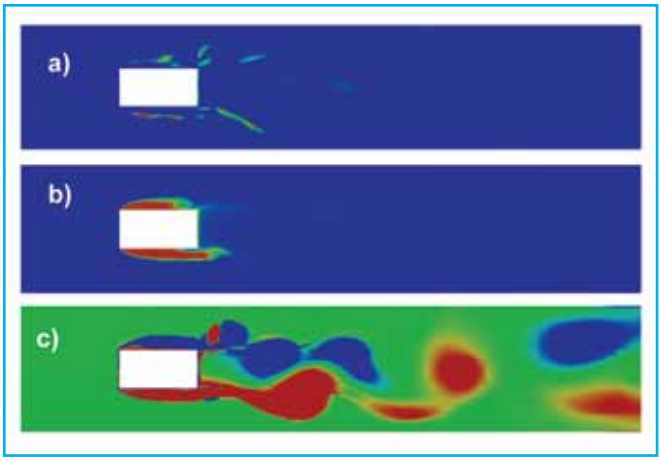

Figure 3. Instantaneous (a) heat release rate, (b) temperature, and (c) spanwise vorticity for the rectangular flameholder with AR = 2.0. The units are in m/s, W/m^3, K, and s^{-1}, respectively. The scales and legends for the contour plots are the same as those in Figure 2

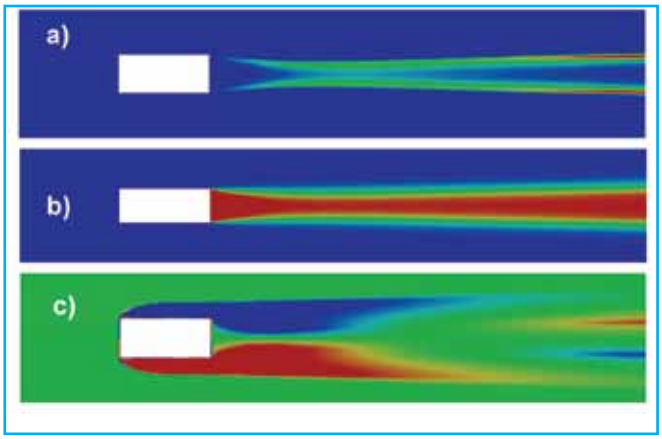

Figure 4. Instantaneous (a) heat release rate, (b) temperature, and (c) spanwise vorticity for the rectangular flameholder with AR = 2.3. The units are in m/s, W/m3, K, and s-1, respectively. The scales and legends for the contour plots are the same as those in Figure 2

for AR = 2.0, the heat release rate contour indicates that the flame exists only on the upper and lower surfaces of the flameholder. At this condition (as indicated by the spanwise vorticity contours), the BvK instability is pronounced, and the flame cannot survive behind the flameholder. The von Kármán vortices lead to quick mixing in the wake, and the hot products from combustion are well mixed. Therefore, the high temperature region is only possible adjacent to the flameholder's upper and lower surfaces within the separation bubbles. This flame exhibits large-scale wake disruption as those described by Lieuwen and coworkers [13]. Although this result is consistent with the statement that von Kármán vortex street plays a role in static instability, this result suggests that BvK instability can be found even at high T_b/T_u of ~8.33. Nevertheless, on the wake T_b/T_u is on the order of unity (cf. Figure 3), which leads to a flow field similar to that of its corresponding non-reacting condition. Finally, for AR = 2.3, the heat release rate contour indicates that the flame is now anchored to the trailing edge of the flameholder. The wake appears to be shorter than that corresponding to the flame anchored on the square cylinder (cf. Figure 2). The high temperature region only exists in the wake and not on the horizontal surfaces of the flameholder. The preheat zone thickness increases from the trailing edge to downstream locations, and then it decreases again further downstream. The spanwise vorticity contours indicate there is no von Kármán vortex shedding.

ii. Effect of Combustion on the Shedding Frequency

Figure 5 presents the von Kármán street shedding Strouhal number (St_{vK}) as a function of aspect ratio (AR) for all non-reacting and reacting square and rectangular bluff bodies presented in [10]], measured and computed St_{vK} as a function of AR for rectangular cylinders. The experiments indicate that St_{vK} decreases from AR = 1.0 with value of St_{vK} = 0.13 to AR = 2.8. The latter bluff body exhibits two dominating St_{vK} with values of 0.056 and 0.152. Then, St_{vK} decreases with further increase in AR. The non-reacting URANS simulations indicate that St_{vK} decreases from AR = 1.0 to AR = 2.0. The predicted St_{vK} nearly matches the experimental values. Further increase in AR abruptly raises St_{vK}. The rectangular bluff body with AR = 2.3 exhibits a St_{vK} value of 0.17, which is several times larger than that expected. Therefore, our numerical simulations with Realizable k-ε URANS clearly underpredict the critical AR. Several other simulations were tried for the non-reacting flow past a bluff body with AR = 2.3 as indicated in this figure, including a refined mesh, wall functions, two-dimensional simulation at lower Re = 20,000, and the SST k-ω URANS. They all predicted approximately the same St_{vK}. LES simulation was not tried because it is computationally expensive. It appears from this plot that it is possible to predict St_{vK} using only two-dimensional URANS simulations for non-reacting flows. St_{vK} definitely decreases from AR = 1.0 because the bluff body width (W) becomes comparable with the separation bubble length, which in turn reduces backflow. At a critical AR, the separation bubble is accommodated on the sidewalls of the bluff body, backflow is reduced, and then St_{vK} increases again. Further decrease in AR reduces St_{vK} again. We believe that the RANS models underpredict the critical St_{vK} for transition because the length of the separation bubble is underpredicted. The characteristic St_{vK} corresponding to the reacting flows discussed in the previous section are also shown in the graph below. When the flow for AR = 1.0 is ignited, the flame becomes statically and dynamically stable with St_{vK} = 0.0 (cf. Figure 2). Now, when the flow for AR = 2.0 is ignited, the flame is dynamically unstable exhibiting von Kármán vortex street (cf. Figure 3). The characteristic St_{vK} associated with this flame is 0.058, which is not far off from those predicted and measured values in the non-reacting condition. Finally, for the flame anchored on the bluff body with AR = 2.3, St_{vK} = 0.0 since it is also statically and dynamically stable (cf. Figure 4). No flame can potentially survive at St_{vK} greater than that predicted by its non-reacting counterpart. In summary, the presence of a flame will definitely alter the flow field and tend to reduce BvK instability, and von Kármán vortex street might play a role in blowout as suggested by previous investigations [12,13].

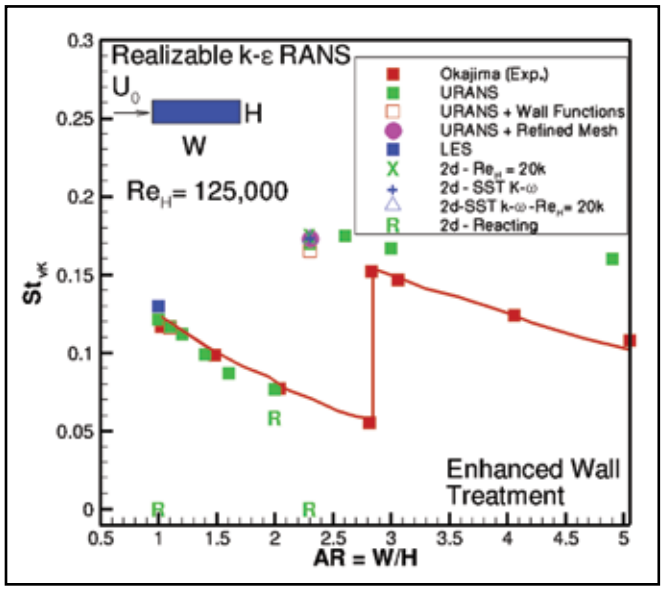

Figure 5. von Kármán Strouhal number (St_{vK}) as a function of aspect ratio (AR = W/H) for rectangular cylinders. The experimental measurements are taken from Okajima [11]

Conclusion

Two- and three-dimensional reacting flows past square and rectangular cylinders are simulated using either Realizable k-ε URANS, SST k-ω URANS, or LES with either two-step global chemistry or 44-step reduced kinetics mechanism for C3H8-air combustion with temperature- and species-dependent thermodynamic and transport properties. The chemical-turbulence interaction is modeled using EDC. Important conclusions are as follows:

29

For the square and rectangular cylinders, the flow separates from the side walls, forming two recirculating bubbles. Two shear layers are then formed at the trailing edges of the bluff body, shedding axisymmetric vortices. For aspect ratios (AR) less than 2.3, there is backflow from the wake. The length of the bluff body (W) is smaller than the separation or reattachment length for AR<2.3. With gradual increase in AR = W/H by increasing W, backflow is diminished, and the von Kármán Strouhal number (St_{vK}) decreases. For 2.0<AR<2.3, W becomes the size of reattachment length, and St_{vK} jumps to a higher value. Separation again occurs at the trailing edge for AR = 2.3. Further increase in AR decreases St_{vK} again.

2. Experimental [11] results report that StvK first decreases linearly from AR = 1.0 until it reaches a minimum value at AR<2.8. The bluff body with AR = 2.8 exhibits two dominating frequencies (one low and other high). Therefore, there is a frequency jump. Further increase in AR decreases St_{vK} again. The predictions using Realizable k-ε URANS accurately match the measurements for $1.0 \leq AR \leq 2.0$. The predicted StvK jump occurs somewhere between 2.0<AR<2.3. With further increase in AR, the simulations predict the decrease in St_{vK}. The model qualitatively matches the experimental results. Quantitative discrepancies are only found at AR≥2.3. This might be because the Realizable k-ε URANS underpredicts the separation bubble size.

3. Experimental [11] results reported in the literature state that the drag coefficient (CD) decreases first rapidly from AR = 1.0 to AR = 1.5, and then CD decreases nearly linearly with further increase in CD for a constant Reynolds number (Re). The numerical results with Realizable k-ε URANS also indicate that CD decreases with increasing AR. The predictions match well in the region of 1.0≤AR≤2.3. With further increase in AR, the predictions deteriorate. This might be due to boundary-induced disturbances, which in turn are due to the proximity of the bluff body to the exit plane. It could also be due to overprediction of the static pressure and underprediction of the reattachment length that leads to pressure recovery.

4. The effects of mesh size, URANS model, and Re on St_{vK} were investigated in a two-dimensional non-reacting flow for a rectangular bluff body with AR = 2.3. The results clearly show that none of the aforementioned conditions had significant effect on St_{vK}. In addition, a LES calculation was performed for the flow past a square cylinder. The LES and URANS predicted similar St_{vK}. These results suggest that two-dimensional flows with URANS models (Realizable k-ε and SST k-ω URANS) are sufficient to calculate St_{vK}.

5. The flame promotes static and dynamic stability for AR = 1.0 and 2.3, whereas the flame is dynamically unstable for AR = 2.0. The latter flame exhibits large-scale wake disruptions, and the flame only exists in the separation bubble region. This flame also exhibits a von Kármán flow pattern, indicating that this instability plays a role in blowout. The flame anchored at AR = 1.0 is attached to the leading edge of bluff body, whereas the flame anchored at AR = 2.3 is attached to the trailing edge of the bluff body. The flames seem to anchor on the most downstream separation location.

6. The triangular cylinder only exhibits flow separation at the trailing edges, forming two shear layers that shed axisymmetric vortices. A flame is positioned in the flow field after igniting the mixture and reaching limit cycle. LES predicts a smaller recirculation length than the Realizable k-ε URANS. LES predicts a flow field in which Bérnard/von Kármán (BvK) instability is suppressed, whereas URANS predicts a flow field with both Kelvin-Helmholtz (KH) and BvK instability. The predicted St_{vK} is 0.36 for both URANS and LES, suggesting again that LES is not necessary for this global prediction. The predicted St_{vK} for the non-reacting condition is, however, lower (i.e., 0.23), which compares well with the experimental value of 0.25.

Acknowledgments

This material is based on the research sponsored by the Air Force Research Laboratory under agreement number FA8650-10-2-2934. Acknowledgments are expressed to the DoD High Performance Computing Modernization Program for providing the computational resources. Special thanks to Joseph Zelina for his insightful comments. The help of Fernando Browning and Abd-Almalek Awad in postprocessing computed data is also greatly appreciated. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

References

[1] Prasad, A., Williamson A., 1997, "The Instability of the Shear Layer Separating from a Bluff Body," J. Fluid Mech. 333, 375-402.

[2] ANSYS FLUENT 12.0, 2009, Theory Guide.

[3] Briones, A.M., Sekar, B., 2010, "Characteristics of Multi-Cavity Trapped Vortex Combustors," ASME Turbo Expo 2010: Power for Land, Sea, and Air, GT2010-22151.

[4] Shih, T.-H., Liou, W.W., Shabbir, A., Yang, Z., Zhu, J., 1995,

"A New k-Eddy-Viscosity Model for High Reynolds Number Turbulent Flows - Model Development and Validation," Computers Fluids, 24, 227–238.

[5] Menter, F.R., 1994, "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," AIAA J., 32, 1598–1605.

[6] ANSYS FLUENT 12.0, 2009, in Validation Guide, "Turbulent Flow in a Transition Duct."

[7] ANSYS FLUENT 12.0, 2005, in Advance Fluent Training, "Turbulent Flow Examples."

[8] Gokulakrishnan, P., Kwasi, F., Klassen, M., Roby, R., Soteriou, M., Kiel, B., Sekar, B., 2009, "LES-PDF Modeling of Flame Instability and Blow-out Stabilized Flames," 45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, AIAA-2009-5409.

[9] Gokulakrishnan, P., Bikkani, R., Klassen, M.S., Roby, R.J., Kiel, B.V., 2008, "Influence of Turbulence-Chemistry Interaction in Blow-out Predictions of Bluff-Body Stabilized Flames," 47th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, AIAA-2008-1179.

[10] Briones, A., Sekar, B., Thornburg, H., 2011, "Characteristics of Bluff Body Stabilized Turbulent Premixed Flames," Proceedings of the ASME Turbo Expo 2011: Power for Land, Sea, and Air, GT2011- 45089, June 6-10, 2011, Vancouver, British Columbia, Canada.

[11] Okajima, 1990, "Numerical Simulation of Flow around Rectangular Cylinders," J. Wind Eng. Ind. Aerodynamics, 33, 171-180.

[12] Erickson, R.R., Soteriou M.C., Mehta, P.G., 2006, "The Influence of Temperature Ratio on the Dynamics of Bluff Body Stabilized Flames," 44th AIAA Aerospace Sciences Meeting and Exhibit, AIAA-2006-753.

[13] Nair, S., and Lieuwen, T., 2007, "Near-Blowoff Dynamics of a Bluff-Body Stabilized Flame," J. Propulsion Power, 23, 421-427.

[14] Kiel, B., Garwick, K., Lynch, A., Gord, J.R., Meyer, T., 2006, "Non-Reacting and Combusting Flow Investigation of Bluff Bodies in Cross Flow," 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, AIAA-2006-5234.

[15] Ghoniem, A.F., Chorin, A.J., Oppenheim, A.K., 1982, "Numerical Modeling of Turbulent Flow in a Combustion Channel," Phil. Trans. Royal Soc. London A, 304, 303-325.

[16] Soteriou, M.C., Ghoniem, A.F., 1994, "The Vorticity Dynamics of an Exothermic Spatially Developing Forced, Reacting Shear Layer," Proc. Combust. Inst., 25, 1265-1272.

[17] Coats, C.M., 1996, "Coherent Structures in Combustion," Prog. Energy Combust. Sci., 22, 427-509.

[18] Sjunesson, A., Henriksson, R., Lofstrom, C., 1992, "CARS Measurement and Visualization of Reacting Flows in Bluff Body Stabilized Flame," 28th SAE, ASME, and ASEE, Joint Propulsion Conference and Exhibit, AIAA-1992-3650.

Faster Tracker: Simulating Satellite Tracking Using Parallel Computing

By Andrew Lindstrom,University of Hawaii at Hilo — Mentors: Carl Holmberg, Maui High Performance Computing Center, DoD Supercomputer Resource Center (MHPCC DSRC), Maui, Hawaii, and Dr. Paul Schumacher, Air Force Research Laboratory (AFRL)

The task of tracking and detecting satellites orbiting the Earth is performed using a number of ground-based sensors around the world, each one tasked as an isolated system. Objectives may include tracking specific, known satellites, or detecting previously unknown objects. Both types of observations are used to update a central catalog of satellites. This data collection and catalog updating process is largely manual, and it may not scale sufficiently to handle the growing number of objects to be tracked. However, networking the sensors together and tasking them from a single facility may obtain the required scalability. This concept can be tested using software models of the existing sensors. The first step of this project is to use an existing MATLAB model in its alpha stage of development and from it create a version that will take on a larger data set without sacrificing turnaround time.

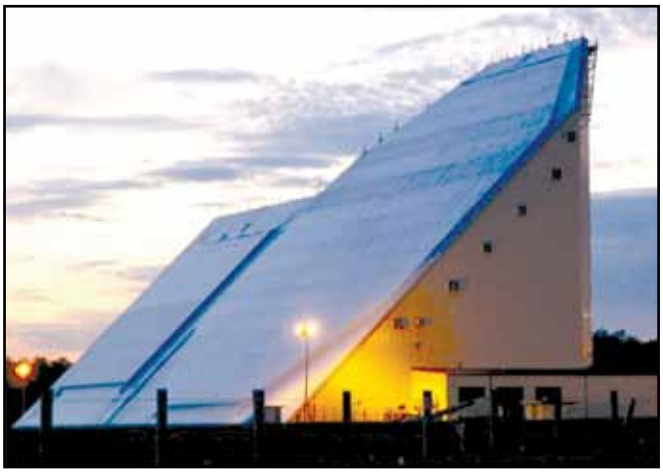

Figure 1. FPS-85 Space Track Radar

The existing model simulates the FPS-85 radar. The model was analyzed to determine which types of performance optimizations would likely provide the desired throughput using the computational resources at MHPCC, while within the time constraints of a summer internship. Initial development was performed on a quad-core, Windows-based PC using a 10-satellite test set. Early optimizations included removal of the user interface, real-time plotting, and client/server socket communication logic. They were superfluous to the batch processing target environment and by far the largest cycle consumers in the original workstation environment. Simulation data were automatically dumped to a file at the end of the run.

The results of additional code profiling suggested that a data parallel computing approach showed the most promise for added performance, using the MATLAB Parallel Computing Toolbox (PCT). Using the PCT made the parallelization process easier than anticipated, the primary effort being to document the original code and identify the segments that were both significant CPU cycle consumers, and suitable for parallel processing. On a symmetric multiprocessing system, the PCT "parfor" construct turns a

standard serial for-loop into a data-parallel loop distributed over the available CPUs and cores.

The majority of processing occurs within a set of nested loops, the outer loop handling the simulation time-steps, the inner stepping through a list of satellites. Fortunately, the steps in either loop were independent of the results of the previous steps. In the interest of minimizing model reengineering, the time-step loop was parallelized. This required only minor rearrangement of the output data structure for compatibility with the "parfor" parallel data handling.

The work then moved to MHPCC's Linux-based *Mana* cluster. The modified code's initial scalability testing on *Mana* used the additional quad-core CPU preset on each node. The only required MATLAB code modifications were for file i/o paths.

The model has shown a potential for scalability thus far not completely realized. (Using the existing data set of 3700 time-steps, the run time was reduced by 2.5 times when moving from a single core to four, both on the Windows workstation and a single *Mana* node. The code showed an ability to scale when using eight cores on the *Mana* node. Enlarging the data set by a factor of 10 resulted in a run time roughly equivalent to the previous four core tests. The PCT only provides SMP parallelism, and does not support clustered, multinode processing. Further scalability testing required the MATLAB Distributed Computing Toolbox (DCT) to span cluster nodes, and there was not sufficient time to become familiar with its use.

Future goals of this project will include testing against a larger data set on a larger SMP system and implementing an approach to fully exercise the *Mana* cluster using the DCT. Ultimately, the goal is to link multiple models over a network, their tasking automatically set by a simulated resource manager. The original PC-based radar simulator would benefit from additional optimization of the user interface and data displays. Reworking the simulation to process each satellite in parallel may provide better performance and more flexibility than time-step parallelism. To prove useful in a production environment, the model should be tested against a major fraction of the entire space object catalog.

Table 1
Model Run Times, Before and after Optimization

Time (Sec)	Cores	Notes
10,424	1	Original model
859	1	Plots disabled
724	1	GUI removed
319	4	'parfar' added
304	8	10x enlarged data set

Prior to beginning this project, I had no experience with HPC and little with MATLAB. I have learned much regarding the purpose and implementation of application while gaining insight into MATLAB's value as an engineering tool. I would like to thank the MHPCC and the Akamai Workforce Initiative internship program for the opportunity to participate in this project.

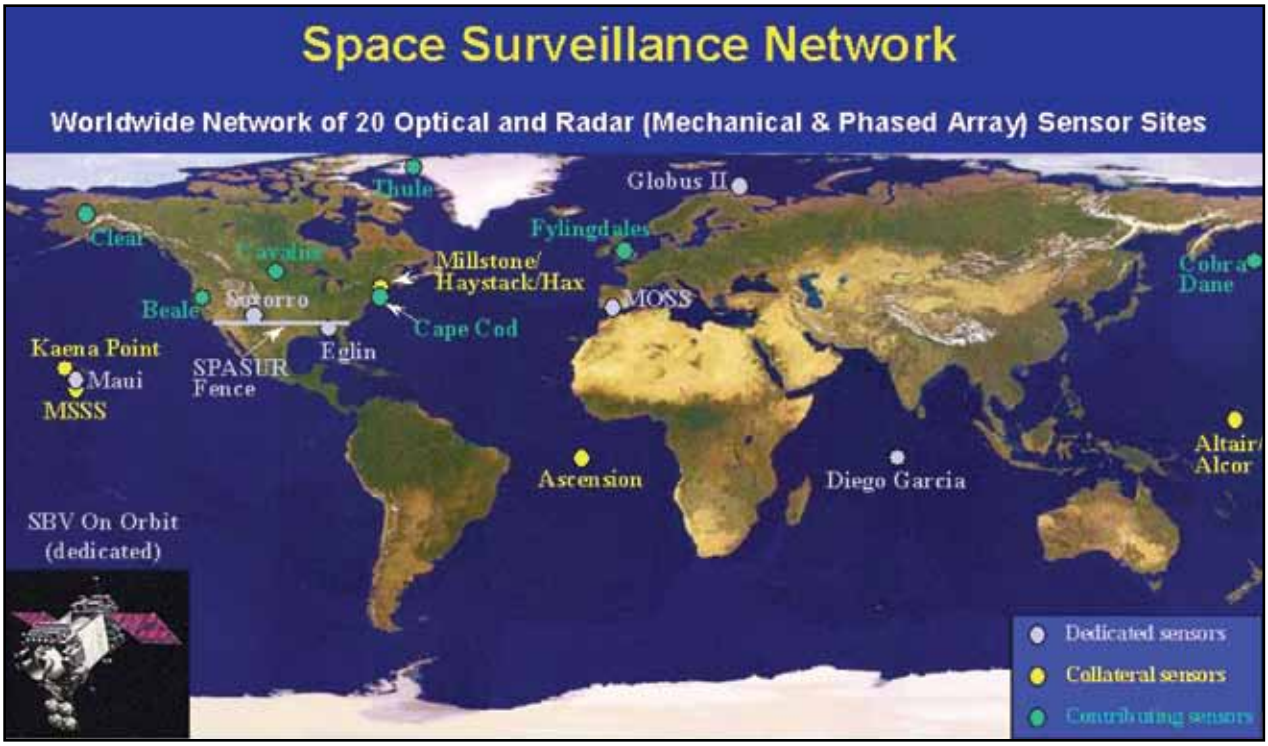

Figure 2. Space Surveillance Network

Optimization of Nanocrystalline Silicon Carbide Ceramics Through Atomistic Simulations

By Drs. Bryce D. Devine, Charles Cornwell, Jabari Lee, and Charles Welch, U.S. Army Engineer Research and Development Center, Information Technology Laboratory, Vicksburg, Mississippi

Introduction

High-performance silicon carbide (SiC) ceramics produced using the most current processing techniques and optimized using computational modeling have the potential to be marvelous structural materials. SiC exhibits a Young's modulus that is twice that of steel, a compressive strength that is three times that of very high strength steel, yet with a density that is comparable with aluminum. The material is resistant to corrosion, wear, and fatigue and performs remarkably well at high temperatures. SiC is composed of silicon and carbon—two cheap and naturally abundant elements that can be produced through fairly low-energy processes. The main limitation of the material is its poor fracture toughness and associated brittleness, and its relatively low tensile strength as compared to compressive strength. These characteristics have precluded its widespread use in structural applications since its commercial development in 1893. With an improvement in its fracture toughness and tensile strength, SiC could replace high-strength steels in many applications at a fraction of the weight and possibly at a fraction of the cost over the lifetime of service.

Improving the fracture toughness is an exceedingly difficult task that has challenged ceramicists for decades. In the past few years, several technological advances have rekindled hope that this challenge may be met. It has been well established that the mechanical properties of materials, such as ductility and tensile strength, vary with the grain-size distribution and, in general, a smaller grain-size distribution leads to improved performance. As the mean grain size in a material is reduced to below the nanoscale, materials typically exhibit an increase in hardness and strength. This effect continues with decreasing grain size until a critical size is crossed. At extremely fine grain sizes, materials get softer with decreasing grain size. The crossing point typically corresponds to a transition from dislocation-accommodated plasticity to grain boundary sliding. The phenomenon is often not associated with ceramics, which for the most part exhibit limited dislocation activity. However, simulations of nano-indentation in polycrystalline SiC suggest a crossing point at about a grain size of about18 nm. At extremely small grain sizes of <5 nm, simulations also reveal an unusual increase in both fracture toughness and tensile strength. Indications of a crossing point and the unexpected effect in extremely small grains suggest that optimization of the grain-size distribution is one possible route to improving the performance of SiC.

Another means of improving fracture toughness currently under investigation is the reinforcement of the ceramic matrix with a tensile material. Experimentally, SiC composites have been produced using SiC whiskers, polymers, carbon fiber, and most recently, carbon nanotubes—all with mixed results. The improvement in mechanical properties of a composite depends on the properties of the two constituents and the strength of the bonding between the constituents. Composites formed with extremely strong tensile materials, such as carbon nanotubes, have demonstrated improvements in fracture toughness, but with a commensurate decrease in overall strength. Careful optimization is required to balance the two properties.

Computational modeling is playing an integral role in current efforts to develop ceramic materials. Optimization of both the grain size and composition requires a sophisticated knowledge of the atomic-scale interactions that control sintering and deformation. Atomistic simulations can provide substantial insight into these atomic level processes. However, due to computational demands of such simulations, modeling of polycrystalline ceramic systems typically has been limited to non-ideal grain sizes, high strain rates, or to early stages of sintering phenomenon. Models with grain-size distributions that are comparable with experimental materials consist of tens of millions of atoms. Furthermore, simulating diffusion-limited processes, as is the case in sintering or failure under slow strain rates, requires excessively long simulation times. With the resources available through high performance computing (HPC), the scope of molecular dynamics (MD) simulations can be enlarged to begin to model these experimentally relevant systems. As examples, the following sections describe results from ongoing large-scale MD simulations exploring the atomic-scale mechanisms that govern sintering and mechanical failure in nanocrystalline SiC.

Simulations of Early Stage Sintering of Nanocrystalline SiC Ceramics

Sintering is the process where a compacted granular material is heated to a sufficiently high temperature to cause the particles to fuse together. The temperature has to be high enough to drive the consolidation reaction, but not so high as to completely amorphize the material. A temperature 0.8 of the melting temperature (T_m) is a typical upper limit corresponding to about 2100ºC in SiC. At these elevated temperatures, grain growth occurs rapidly to the point where a nanoscale grain structure is difficult to achieve. A significant challenge to producing ceramic composites with small grain-size distributions has been developing a sintering process that fully consolidates the solid while preserving the fine grain size and the reinforcement material. Recently, Field-Assisted Sintering Technology (FAST) methods, also known as Spark Plasma Sintering, where an electric field is applied across the specimen during sintering, has greatly alleviated much of these challenges. The field induces rapid sintering at lower temperatures, which inhibits grain growth and allows for the inclusion of temperature-sensitive materials, such as carbon nanotubes (CNTs). As a new technology, not much is known yet about how the field affects the sintering mechanisms. In fact, the actual consolidation mechanism is unknown, which relegates progress to the plodding speed of trial and error.

Sintering, in general, involves the exchange of material between particles and between the particle interiors with their surfaces. Consequently, the relative rates of surface, grain boundary, and point-defect diffusion determine the overall sintering rate. Current efforts seek to determine which diffusion processes enhance and limit sintering at various temperatures and under the influence of an external field.

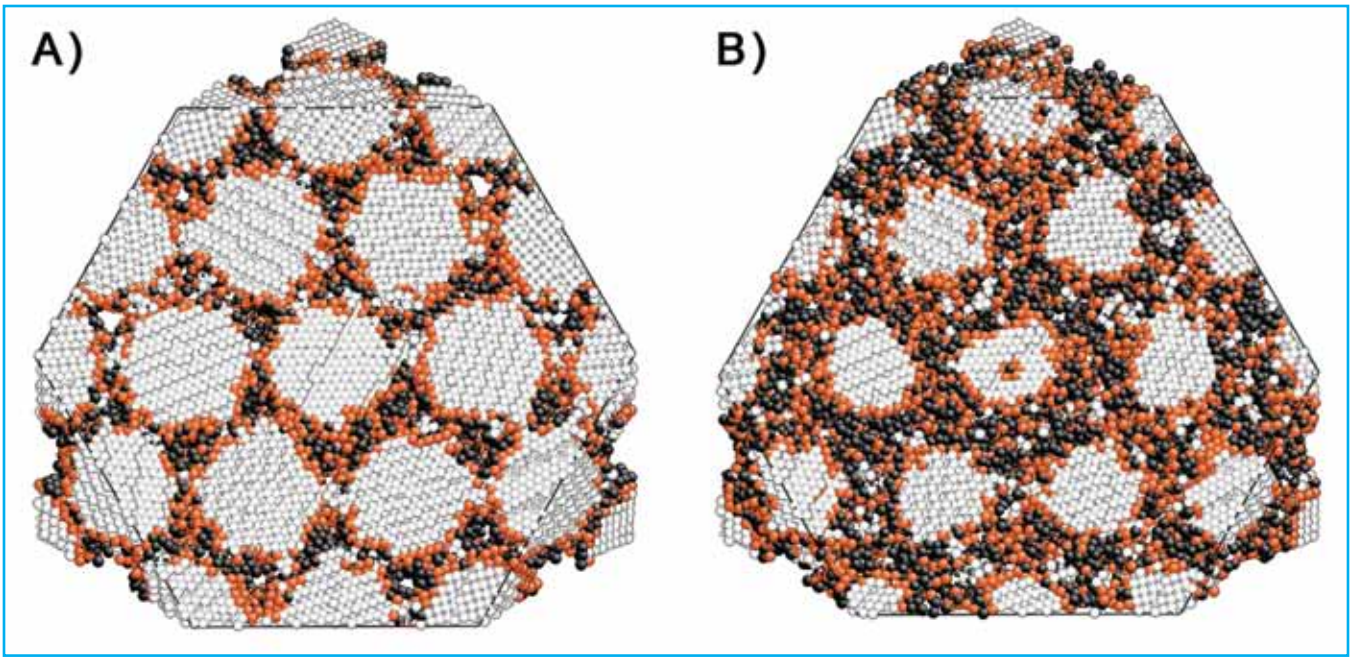

Figure 1. Microstructure observed in early stage of sintering 5-nm SiC nanocrystals. Images show cross section through close-packed plane of system after sintering at 0.7 T_m (A) and 0.8 T_m (B) for a time of 0.5 ns. Black atoms indicate under-coordinated grain boundaries and defects, orange atoms represent dislocations, and grey atoms are crystalline SiC

So far, simulations have revealed that extremely small grain sizes of <5 nm consolidate at an extremely fast rate. Below 5 nm, particles are able to rotate to form low-energy interfaces, which in turn leads to increased neck formation between particles and an increase in the densification rate. Rotation is inhibited with larger particle sizes, in which case the sintering reaction proceeds more slowly. This is shown in Figure 1A, which shows a cross section of the simulated system sintered for 0.5 ns at 0.7 T_m. Black atoms in the figure indicate under-coordinate atoms at grain boundaries and vacancies, orange atoms indicate dislocations, and grey atoms are crystalline SiC. Grey areas between crystals indicate transgranular crystal growth across low-energy boundaries. Plots of the densification rate at temperatures ranging from 0.5 to 0.8 T_m with a 5-nm grain size are shown in Figure 2. The initial portion of the curve at 0.5 T_m shows an increased rate of densification that corresponds to neck formation. Simulations with grain sizes of 10 nm and larger do not show significant particle rotation and consequently exhibit less neck formation in the initial stages of sintering.

A pronounced increase in densification rates is observed at 0.8 T_m as shown in the bottom curve in Figure 2. The rate increase also coincides with a significant increase in vacancy formation and diffusion through the crystals. Vacancies can be observed in Figure 1B, which shows the microstructure of the system after 0.5 ns at 0.8 T_m. Vacancies show up as black atoms surrounded by an orange-colored ring of atoms in the interior of the crystals. The effect is delayed with increased crystal size. As shown in Figure 3, which plots the densification histories of various nanocrystals at 0.8 T_m, larger diameter crystals exhibit a decreased rate and delayed time before the rapid densification is observed. It has not yet been determined whether the increased defect formation rate is the cause or a consequence of the increased densification rate. However, vacancies are readily formed in the larger diameter systems, suggesting that the rate is proportional to the longer diffusional paths across larger diameter crystals rather than a decrease in the vacancy formation rates.

Simulations of High Strain Rate Deformation of Nanocrystalline SiC

SiC readily forms screw dislocations at elevated temperatures during sintering. Previous MD simulations have shown that dislocation loops are formed during plastic deformation at standard temperatures, indicating a ductile deformation mode

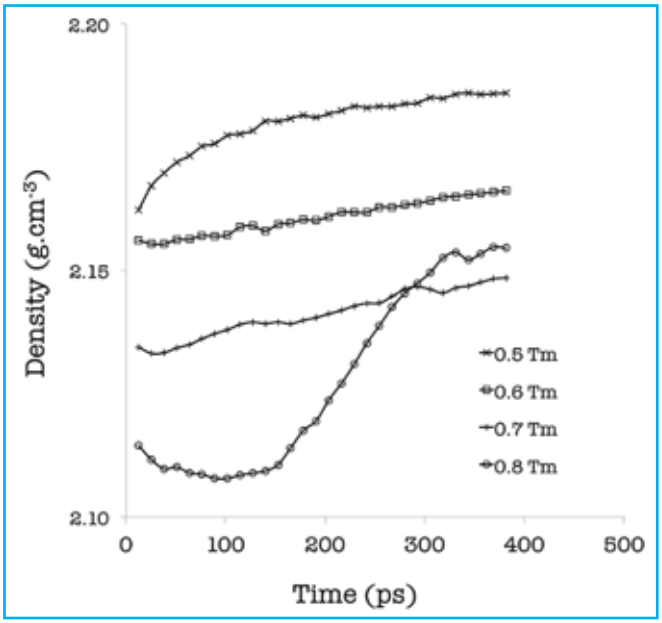

Figure 2. Densification rates observed during sintering of 5-nm nanocrystalline SiC. Curves shift downward for each temperature due to thermal expansion of system. Increased rate observed during first 100 ps at 0.5 T_m attributed to neck formation. Rate increase at 0.8 T_m coincides with an increase in vacancy formation and diffusion through the crystals

that may be optimized to improve toughness. Currently, atomistic simulations are being conducted to explore effects due to grain size on the material response in the high rate and quasi-static strain regimes.

Initial simulations on a textured, columnar polycrystalline system have shown that at high strain rates, fracture occurs before formation of the lowest energy dislocation. Some plastic deformation does occur, but it is accommodated entirely by grain boundary sliding.

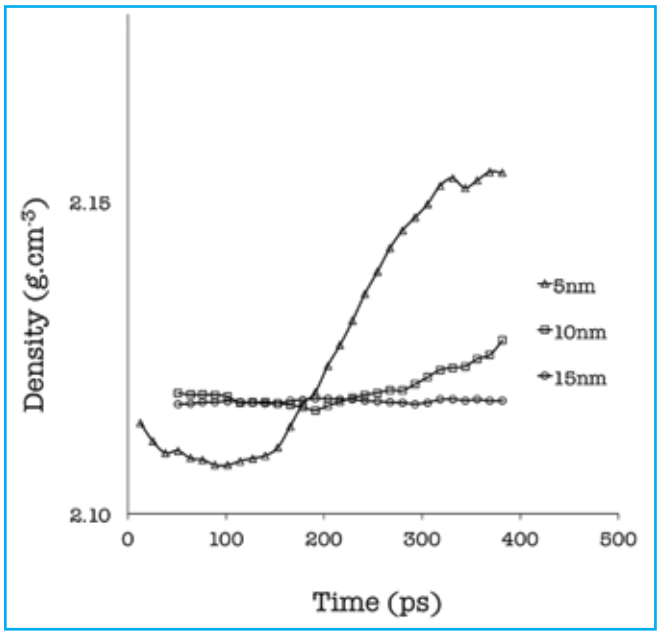

Figure 3. Densification rates observed during sintering of SiC at 0.8 T_m for 5-, 10-, and 15-nm-diameter crystals

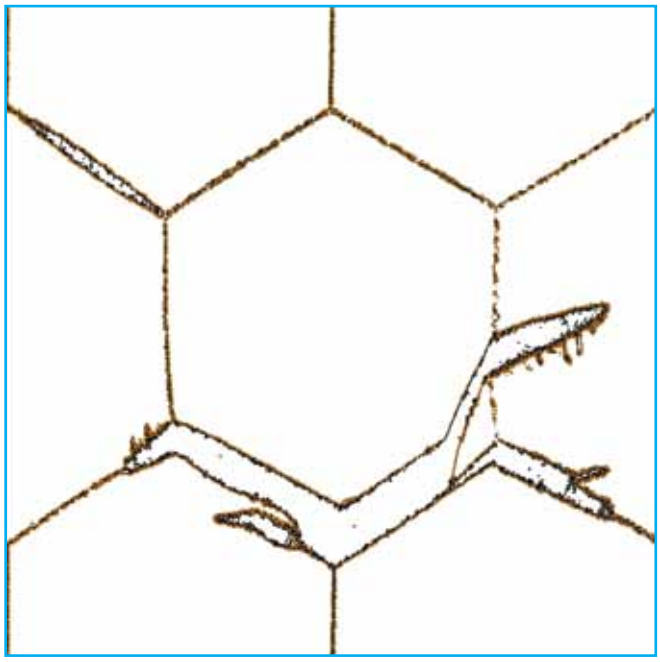

Figure 4. Fracture in 75 nm columnar SiC nanocrystals. Crack progression inhibited by dislocation (orange) formation in crystals with diameters above 50 nm. Black indicates atoms at grain boundaries and surfaces

Simulations also reveal a surprisingly complex fracture response to high strain rates. Cracks tend to progress in a transgranular fashion when the grain boundaries are well formed. During transgranular fracture, crystals only exhibit shear in one plane (100). Cracks in other planes are attenuated by dislocation formation at the crack tip. In small crystals of <10 nm, dislocation formation is inhibited, resulting in more frequent transgranular shearing. In crystals of greater than 50 nm in diameter, the formation of dislocations inhibits crack progression even in the shear plane, leading to ductile-like fracture. Figure 4 illustrates the fracture behavior in a 75 nm crystalline system. For clarity, only atoms that are at surfaces and grain boundaries (black) or in dislocations (orange) are depicted. In Figure 4, fracture consists of a mix of transgranular and intragranular cracking. The fuzzy edges of the cracks indicate dislocation formation.

Conclusions

Although this project is still in its initial stages, the atomistic simulations are already revealing key properties of nanocrystalline SiC ceramics that can assist in optimizing both the sintering process and the design of the composite. Simulations of sintering indicate that rapid densification rates correlate directly with vacancy diffusion. Current work is seeking to determine whether a similar mechanism is responsible for rapid sintering observed during field-assisted sintering. Simulations of mechanical response suggest that grain boundary modification is a good approach to increase the ductility and toughness of SiC materials. Previous simulations have shown that small grain sizes may lead to an increase in both toughness and strength. However, controlling the sintering process with such small grain sizes will be challenging given the rapid rate at which small grains sinter.

Simulations of the size required by this project are only tractable with the HPC resources. In both the sintering and mechanical response studies described here, large system sizes and long simulation times are required to reveal phenomena relevant to the actual experiments. Large-scale atomistic simulations are being used to both guide material design and the design of the ceramic sintering process, and are laying the foundation for rapid developments in structural ceramics.

Acknowledgments

This study gratefully acknowledges funding support from the U.S. Army Engineer Research and Development Center research programs "Nanoscale Studies of Polycrystalline Materials with Emphasis on Ceramics Synthesis" and "Multi-Scale Understanding and Engineering of Grain Boundaries in Ceramics," and the allocation of computer time from the High Performance Computing Challenge Project "Molecular Dynamics Simulations to Underpin the Design and Development of High-Performance Carbon-Nanotube-Based Filaments, Membranes, and Coatings." We thank the Chief of Engineers for permission to publish this article.

Air Force Research Laboratory
DoD Supercomputing Resource Center

From the Director's Desk – Frank Witzeman

Last year around the date of Supercomputing 2010 (SC10), the AFRL DSRC was receiving the DoD HPCMP Technology Insertion 2010 (TI-10) Cray XE6 system designated *Raptor*. This new system, the largest in the Defense Department, consists of 43,712 compute cores (2.4 GHz AMD Magny Cours), 87 TB memory, and 1.6 PB disk storage. *Raptor* was successfully integrated into our facility in near record time, and it was available to early access users in January this year. On February 16, 2011, *Raptor* went into full production mode, just over 3 months from its arrival. During the early access capabilities demonstration phase, the largest user application consumed 43,008 cores. Successful scaling across 20,000-30,000 cores was also demonstrated for a number of applications that can take advantage of *Raptor's* architecture, including the extremely fast and robust Gemini interconnect. Despite some early system issues, *Raptor* has become a stable and suitable leadership-class supercomputer for the DoD HPCMP, capable of achieving over 400 TFlop/sec.

You may recall from previous *HPC Insights* issues that the DoD HPCMP is introducing the HPC Enhanced User Environment (HEUE) project to the user community. The primary function of HEUE is to provide an enhanced user interface to HPC systems and advanced capabilities to analyze and store data. HEUE consists of three main components: (1) Utility Server – US, (2) Center-Wide File System – CWFS, and (3) Storage Lifecycle Management – SLM. Late last year, the AFRL DSRC received the US and CWFS systems and started integrating them into our environment. The US is an Appro cluster with 88 mixed-mode nodes, where 44 of the nodes are used for standard compute (2.3 GHz AMD Opteron 8-core processors, 16 cores per node), 22 of the nodes are used for large memory (possibly shared) applications (32 cores and 256 GB per node), and the remaining 22 nodes are used for graphics or general-purpose computing on GPU-based (NVIDIA Tesla M2050) architecture. The CWFS is a Panasas solution with PAS8 storage shelves holding 1.36 PB of raw data storage, over 1 PB usable. Finally, the SLM component is a software and hardware solution based on General Atomics' Nirvana Storage Resource Broker to associate metadata to files for retention/archiving purposes.

On August 1, 2011, the US and CWFS at the AFRL DSRC entered production status. The HEUE environments at the other DoD HPCMP Centers were put into production on different dates. Several enhanced features include secure remote desktop services for data visualization, remote batch job management for applications submitted to and returned from the large HPC systems, and interactive pre- and postprocessing of data associated with the HPC applications. The CWFS is configured such that the users' home directories for the US are located there, and the files are also visible to the HPC systems' login nodes. The CWFS really represents a midterm repository for user HPC data to reside for about 30 days before archiving or deletion. This will help alleviate some of the issues with the short-term data storage on the HPC systems' scratch or workspace, as well as prevent rapid automatic default archiving of nearly all user data. It is anticipated that the SLM capability will be fully functional by the SC11 timeframe.

As we've ramped up our efforts to bring *Raptor* and HEUE into our services portfolio, we've also put the finishing touches on our existing facility upgrade and followed the progress being made on our new building. To be consistent with green energy initiatives, we've incorporated free cooling loops, water-cooled chillers, and water-cooled HPC compute racks where possible. Our goal is to reduce dependence on massive amounts of cold air needed to dissipate heat from our data centers and to leverage our Midwest climate to lower electricity usage. We continue to request help from the HPC vendor community to do their part to discover more innovative and effective ways to reduce energy consumption from the multiple sources including boards, blades, racks, and so on. Even though our new building (available in March 2012) will have sufficient space and capacity to host air-cooled HPC systems, we're demanding that only water-cooled systems be installed there. The entire building is high-level Leadership in Engineering and Environmental Design (LEED) certified, and we must continue to keep energy costs to a minimum. Bringing chilled water directly to the HPC racks or to associated heat exchange units is therefore the preferred approach, and raising the chilled water temperature is also highly desirable. Together with HPC vendors, we can ensure that an energy-efficient roadmap is developed and followed so that we can reduce our costs and save valuable resources while growing our HPC capabilities for many years to come!

Frank Witzeman
Director, AFRL DSRC

Army Research Laboratory
DoD Supercomputing Resource Center

From the Director's Desk – Dr. Raju Namburu

Welcome to the 2011 fall issue of *HPC Insights*. Things have been busy here at the Center with identifying and converting an existing building at Aberdeen Proving Ground (APG) to accommodate Technology Insertion - FY12 (TI-12) computers. This facility is envisioned to be scalable for supporting new HPC systems in FY12 and beyond, and will be the home of the ARL DSRC.

As most of you know, the ARL DSRC systems were housed in two buildings 150 m apart. U.S. Army S&E expertise in HPC, advanced networking, and software application development evolved in these two APG buildings under various Army HPC programs, including the current DoD HPCMP, starting with the *ENIAC* in 1946. The ARL DSRC developed innovative design approaches in converting these two old buildings to house state-of-the-art HPC systems and maintained them exceptionally over the last two decades. As the DSRC systems started growing, scalability of space and associated facilities started posing technical challenges and additional costs. MG Nikolas Justice, Command General of the Research Development and Engineering Command (RDECOM) and his management team, RDECOM Deputy Director Gary Martin; ARL Director John Miller; Communications-Electronics Research, Development, and Engineering Center Director Ms. Jill Smith; and Computational and Information Sciences Directorate Director Dr. John Pellegrino swiftly gathered future requirements for the ARL DSRC to address space and power needs for FY12 and beyond. RDECOM and the APG command worked with the U.S. Army Test and Evaluation Command and Program Executive Office Integration (PEO-I) and proposed a few existing buildings as possible options to house the ARL DSRC.

The ARL DSRC staff collected all the pertinent data and developed comparative analyses between the two proposed buildings. We will highlight some of the pertinent features of these two buildings. Building candidate 1 is a single-story structure with an 18-ft ceiling and an overall area of 26,000 ft^2. This building is currently used as an office space and an ARL advanced computing and scientific visualization laboratory. The ARL DSRC proposed a plan to build a 5500 ft^2 raised-floor computing space to house TI-12 computers with the flexibility to expand and a provision to support 12 MW of power. Building candidate 2 is a three-story structure with approximately 20,000 ft^2 for each floor and a sheltered three-story-high bay area. This building is currently used as an office space by PEO-I, with a number of cluster computers and networking experimentation labs. Similar to candidate 1, the proposed plan for candidate 2 is to use a 5500 ft^2 raised-floor computing space for TI-12. The RDECOM management team will make a final decision on the facility before September 30, 2011, in order for the ARL DSRC to move forward.

The ARL DSRC is one of the tenants at APG. As a tenant, the ARL DSRC at APG played an important role by leveraging synergy between AMC/RDECOM S&E and Army T&E in accelerating design, acquisition, and fielding of new technologies. For example, synergy between S&E and T&E at APG with HPC helped to field Frag Kit 6 in Theatre in under 4 months. Similarly, the ARL DSRC classified computing capability supported the AMC/RDECOM Ground Combat Vehicle Program and a number of Army Program Executive Office Integration Program Management acquisition programs.

The Base Realignment and Closure process at APG will provide new opportunities for HPC, including C4ISR, Cyber Defense, Data to Decisions, etc. RDECOM vision in establishing scalable HPC facilities for the ARL DSRC will play an important role not only in taking advantage of synergy between various old and new tenants at APG, but also in accelerating/fielding Army/DoD mission-critical requirements and emerging technologies.

We continue to prepare for the arrival of new HPC capabilities and increased computational power in FY12 and beyond. The increase in computational power and envisioned scalable facilities will bring challenges and new opportunities to our users and staff. Along with the ARL DSRC staff, we are looking forward to taking on new challenges and providing the best computational resources and services to solve research challenges for our HPC user community.

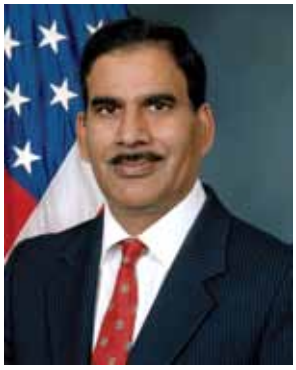

Dr. Raju Namburu
Director, ARL DSRC

Energy Aware Scheduler Saves Resources

By Mike Knowles, ARL DSRC Site Lead, Lockheed Martin

Background

PBSPro is the scheduler employed across the HPCMP to schedule required resources for customer job demands. This scheduler attempts to meet all job requirements and maximize usage of HPCMP resources. This capability can often keep the systems 80 to 90 percent busy, as some resources are constantly in transition and being reserved for the next highest priority job. As it is not feasible, or advisable, to maintain systems at 100 percent utilization (due to extended expansion factor and job start-up time concerns), it would be attractive to reduce power requirements on the portion of machines that are not actively being used. PBSPro has an optional module that is oriented toward power monitoring and control. This module has the capability to power off some HPC nodes and systems that are currently not required or being reserved in anticipation of being allocated to a pending job. Many commercial sites already take advantage of this capability, especially in a periodic job schedule mode of operation, where peak resources are only needed during specific times of the week/month. Estimates indicate that this simple mechanism could save millions of kilowatt-hours (kWH) during the year across the Program (10 percent power reduction on ARL HPC systems would yield approximately 1.4M kWH power reduction per year). Associated reduced cooling requirements would yield further power reductions. The overall goal is to retain or exceed current system utilization numbers while significantly reducing site power requirements.

Phased Implementation

The ARL DSRC in conjunction with Altair and Instrumental are currently evaluating this Energy Aware Scheduler (EAS) capability. This effort consists of further development and integration of a robust PBSPro mechanism to work on compatible HPCMP assets to control and coordinate power to resources that are not currently in use. The EAS, while architecture specific, could be employed on most of the computing assets of the HPCMP. In order to integrate the EAS modules into production, a phased process will be employed. The process involves working with the EAS in a simulation mode to understand system-specific scheduling methods as they relate to the available system control methods. The EAS will be configured for each architecture to use existing control mechanisms. There are also many parameters associated with the EAS per architecture that will affect overall behavior. These parameters are associated with the system and scheduler configuration and include node size, node disk configuration (i.e., are disks present on nodes), idle time when node is available to be powered off, time prior to job start where node needs to be available, node boot time, node boot time integration requirements, last node downtime, and many others. Likewise, each architecture has sets of commands to control node availability and power profile. Some nodes may actually be able to be put into sleep mode instead of full power down. The first phase of the EAS effort is an attempt to capture all associated relevant system/scheduler parameters and system commands to run a long-term simulation to understand how parameter changes would affect the behavior of the EAS. Both the test and development systems (TDS) and production systems will have significant simulation time and detailed analysis to fully understand the effects of the EAS configuration. EAS decisions and resulting pseudo command invocation will be logged during the simulation, and overall effectiveness will be evaluated. This phase of the project will comprise the bulk of the time spent per architecture and system. Once the EAS configuration reaches appropriate system and scheduler effectiveness levels, the EAS will be moved into production first on the TDS systems, then ultimately to the production HPCMP assets. At every step of the way, EAS principals will be working with systems and site personnel to ensure overall reliability of HPCMP assets and to ensure minimal disruption to user jobs.

One of the by-products of the EAS effort is to evaluate node reliability concerns. The diskless configurations of many of the HPCMP HPC assets are targeted towards increasing system reliability, as disks historically decrease system reliability. The power cycling of nodes causes some concerns about the EAS, especially nodes with disks. However, due to the overall node and component reliability increases, controlled environment of the HPCMP systems, increasing Program power consumption, and increase in power pricing, the EAS may prove to be a cost-effective way to reduce overall system and site costs.

The current implementation of the EAS runs asynchronously from the PBSPro scheduler. Altair is investigating mechanisms to further integrate EAS components into the scheduler to shorten the feedback loop. Also, the ability to allow reservations to remain whole (i.e., not become degraded when a node is off due to EAS request) is being investigated. Most likely this situation will be indicated as a new state in the scheduler that has to be considered during scheduler job and system evaluation.

The EAS effort will be closely coordinated with and vetted through the HPCMP Workload Management CoP (Communities of Practice). This group of Program personnel has coordinated the Program-wide implementation of PBSPro, the ARS (Advance Reservation System), the SLB (Shared License Buffer), the CWJM (Center-Wide Job Management) component of HEUE, and other Program initiatives in the last few years. The WMC will work in close conjunction with EAS efforts to ensure job throughput and that system reliability remains high, while overall system/site/Program power requirements are reduced. Interested parties can coordinate thoughts, suggestions, and contributions by e-mail at wmt@hpcmp.hpc.mil.

ARL Making Room for More Horsepower

By John Lazorisak

A building at Aberdeen Proving Ground that once provided the United States Army with horsepower will be providing a new form of horsepower for both ARL and the HPCMP. ARL in conjunction with the HPCMP will be adding 5500 square feet of new computing space for future HPC systems. The new facility, which will be located in an already existing building at Aberdeen Proving Ground, will provide future expansion for both ARL and HPCMP HPC capabilities as well as leveraging new technologies to reduce energy costs.

The building, which will house the new HPC site, was originally built in 1918 by the Army as a stable and is approximately 26,000 square feet. The stable was then used for working horses and mules helping in the relocation of artillery and shells at Aberdeen Proving Ground and elsewhere. The building was later renovated into a commissary (no relation to horses) and then into office space that now includes the ARL Advanced Computation Scientific Visualization Laboratory. An office space section of the building approximately 5500 square feet in size will be renovated to a raised floor computing space. Being originally designed as a stable, the building layout is long and thin with 18-foot-high ceilings, making it a perfect fit for renovation into a raised floor computing environment. The renovations include demolition of the current office space area, refurbishment of the area with the installation of a 30-inch raised floor, and installation of new and efficient power and cooling systems. An existing mezzanine area will also be refurbished with glass walls to allow meeting attendees to look out over the computing complex. In addition, the area surrounding the building provides ample area for the mechanical and electrical plant necessary to support the facility, including space for chillers, generators, and containerized UPS systems.

The facility will support up to 12 megawatts of power (15,000 KVA at an 80 percent power factor) or approximately 16,000 horsepower (1 horsepower = 746 watts) capability. The power will be sourced from an existing 34.5 kilovolt transmission line located near the building and distributed through five 3000 KVA transformers into the facility. The 12 megawatts will provide the power necessary for the HPC assets, including the required power for backup and cooling systems.

To provide an uninterruptable and conditioned power source, the facility will be outfitted with five 1200 KVA UPS systems. The UPS system will be based on new flywheel technology instead of conventional battery UPS systems. A flywheel UPS does not rely on batteries to store electrical energy, but rather stores electrical energy as kinetic energy in dense, spinning discs. A flywheel UPS system can be up to 99 percent efficient and require less maintenance, space, and cooling compared with a conventional battery-based UPS system. In the event of a utility power failure, the flywheel UPS system will provide enough energy to ride through short outages and provide sufficient time for the backup generator system to engage and provide power.

The backup generator system will consist of five 2500 KW diesel generators mated to the facility power system through an automatic transfer switch system. In the event of a main power failure, once the generators have started and are ready to engage the power load, the automatic transfer switch gear will switch over to the generators from the UPS and provide electricity to the facility for approximately 24 to 48 hours.

The cooling plant will consist of six 300-ton air cooled water chillers, each outfitted with air economizing coils, five 10,000-gallon chilled water storage tanks and 16 computer room air handlers (CRAHs). The current design of the chiller plant is an N+1 configuration, allowing any one of the four chillers to be offline for maintenance or other issues without impairing the cooling system. The 10,000-gallon chilled water storage tanks will provide backup cooling in the event of a power failure much like a UPS provides electrical power during a power outage. The computing space is to be outfitted with 16 CRAHs plumbed into the chilled water plant to transfer heat from inside the computing facility. The computing space will also support future chilled water cooling applications, such as water cooled rack doors or directly cooled HPC systems.

Planning and design of the facility has already begun, with demolition and construction to begin soon. The expected completion date of the facility is March of 2012. The renovated facility will house the upcoming TI-12 and TI-14 systems, the MANET DHPI system, as well as providing crucial HPC swing space.

This facility renovation is only a fraction of the size of the entire building and will alleviate the need for near-term computing space. It will provide the capability for ARL and the HPCMP to expand in the near future and provide potential for future expansion.

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

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

The solution of partial differential equations (PDE) accounts for a majority of the computing time on ERDC supercomputing systems. The equations of fluid motion are a good example; they're solved in applications ranging from regional models of storm surge to high Reynolds number flow past an airfoil. The equations of electromagnetism are another good example; they're solved in applications ranging from antenna modeling of integrated shipboard systems to high-energy microwave devices. Quantum chemistry applications solve approximations to Schrodinger's PDE.

Scientists have been solving PDE on supercomputers for over 30 years. There have been tremendous advances in variety and complexity of models and in the size of the computational meshes that represent the underlying physical domain. It is now possible to solve finite element problems with billions of elements, over highly irregular physical domains. Yet, despite these advances, the underlying numerical methods are still fairly basic. Exascale computing experts recently concluded that the predominant approach involves[1]

"…first-order-accurate operator-splitting, semi-implicit and explicit time integration methods, and decoupled nonlinear solution strategies. Such methods have not provided the stability properties needed to perform accurate simulations over the dynamical time scales of interest. In most cases, numerical errors and means for controlling such errors are understood heuristically at best. The solutions may … be stable but [may] contain significant long-time integration error."

The predominant approach lags the state of the art in numerical methods because of the time and effort required to modernize a code. In practice, new codes with more advanced numerical techniques are developed to replace older ones. But the integration of fully implicit time integration and fully coupled nonlinear solution methods require investments in numerical analysis staff and software development. Such investments can be difficult to justify for small code development projects, so this is clearly an area where scientific computing middleware plays an important role in assisting code evolution.

Open-source numerical libraries, supported by a large user community, offer an affordable and sustainable model for integrating new methods into scientific codes. The PETSc linear algebra library is a good example, one that is currently supported on HPCMP systems. PETSc (Portable Extensible Toolkit for Scientific Computation) is a suite of data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations. It provides the user with many alternative methods for solving linear systems of equations. The ParMETIS mesh partitioning library is another good example. ParMETIS is an MPI-based parallel library that implements a variety of algorithms for partitioning unstructured graphs, meshes, and for computing fill-reducing orderings of sparse matrices. It helps the user divide their mesh into subdomains in a way that minimizes the amount of communication between subdomains. How many of the DSRC staff are able to advise and participate with users interested in integrating these libraries into their codes?

All parallel PDE solvers require communication between subdomains. In finite element codes, communication is needed where an element is shared between two or more subdomains (i.e., processors). Whether it's an explicit time-stepping code or an implicit steady-state solution of a linear system, the communication between subdomains occurs at regular intervals. This communication represents the overhead of a parallel algorithm.

Amdahl's law clearly constrains the achievable speedup on a problem according to the fraction of overhead. However, we tend to solve larger problems as more processors become available. So long as the overhead grows more slowly than the number of processors, we enjoy practical

Dr. Robert S. Maier
Director, ERDC DSRC

[1] Modeling and Simulation at the Exascale for Energy and the Environment, Report on the Advanced Scientific Computing Research Town Hall Meetings on Simulation and Modeling at the Exascale for Energy, Ecological Sustainability and Global Security, p. 88, http://www.sc.doe.gov/ascr/ProgramDocuments/ProgDocs.html.

scalability. But if one is interested in simply computing faster and not solving larger problems, the potential of petascale and exascale computing is more difficult to realize. Hence, there is a need to overlap communication and computation, to "hide" the overhead.

Modern parallel architectures provide hardware and software layers that allow memory transfers between nodes even while CPUs continue to process numerical calculations. But the compilers, programming techniques, and middleware that might allow users to take advantage of the overlap are not yet well developed. This need is well-known and has been the subject of much research over the past decade. Programming methods for overlapping communication and computation are slow to find their way into general practice. It is an advanced topic in MPI and OpenMP, and not usually taught in MPI classes. It is also an investment in code development and may require revising computational kernels. There's some hope that parallel global address space compilers will eventually automate the overlap, just as compilers were eventually designed to automatically vectorize codes. But that's not really soon enough to help move our users into the petascale computation realm in the next several years.

There's a role for the DSRC in providing some expertise in open-source numerical libraries, and programming techniques for overlapping communication and computation. My experience has been that the host laboratories view the DSRC as the Center of such expertise. How well are we playing that role? In today's budget environment, this has to be a more effective partnership between the DSRC and PETTT. Brad Comes once said he could recall when the PETTT on-sites were much more a part of the DSRC activity. I think he meant that the DSRC played a bigger role in vetting their projects, making sure they were relevant to parallel computing. While all of our PETTT activities today are worthwhile, few of them actually focus on code scalability. This is an area where we can work together to more firmly establish the DSRC as the Center of expertise in parallel computing.

Aerial view of construction at ERDC Information Technology Laboratory earlier this year. Insert is more recent view showing office space and auditorium being built around supercomputer floor space.

Maui High Performance Computing Center
DoD Supercomputing Resource Center

From the Director's Desk – David Morton

There has been a revolution over the past few years in supercomputer capabilities. Commodity technologies have driven down the cost of hardware, so capabilities are available that would have been unimaginable just a few years ago. One of the biggest challenges is making these incredible resources available to users in ways that are conducive to solving complicated problems faster and easier. In accordance with the theme for SC11, "Connecting Communities Through HPC," the MHPCC DSRC is employing new tools to make supercomputer resources easily available in novel ways to a broad range of users.

In particular, the MHPCC DSRC has been directed by the DoD HPCMP to create and deploy a web-based portal environment. This is a "Software-as-a-Service" (SaaS) model that removes barriers to HPC resources and data. I sometimes explain the concept as the HPC analogy to Google Apps. Web services technologies are used for the implementation serviced, and access to HPC resources is controlled by a job submission/queuing/scheduling service that ensures optimal system utilization.

Although the traditional research and test community is well served through the customary batch computing approach, a larger DoD science and engineering community (i.e., mainstream acquisition) is left underserved today. The portal approach addresses the needs of this class of users. As the portal's model of "HPC Software-as-a-Service" is implemented, a key goal of the Program will be to provide outreach to this community.

Additionally, a new program that has been exploited at the MHPCC DSRC is the implementation of Dedicated Support Partitions (DSPs). A DSP provides dedicated processors to a specific project for a significant period of time to accomplish work that could not otherwise be done in a shared resource environment. The asset that is made available to the user is basically a dedicated cluster of a set size. This asset can be a 24 x 7 asset or only made available during requested times. The majority of this work employs the Dell Quad Core Xeon Cluster *(Mana*) with 9216 cores at the MHPCC. All project leaders with an active RDT&E computational project are eligible to submit a proposal. Further, all application software development efforts, large-scale weapons system test support, and other activities requiring substantial dedicated time on HPCMP resources that cannot be serviced through normal batch processing, interactive processing on the new utility servers, nor the Program's advance reservation service will be considered. There are a number of Principal Investigators who are capitalizing on this new program.

The MHPCC DSRC has partnered with the Pacific Command (PACOM) to implement the Energy Efficient Computing (E2C) Joint Concept Technical Demonstration (JCTD). The E2C JCTD is a proposed $13M effort (including a $7M portal and E2C supercomputer) to retrofit a model DoD efficient Data Center at MHPCC. The 2-year (FY12-13) effort will develop and demonstrate a model set of energy efficient design recommendations applicable and tailorable to DoD legacy data centers.

As an Outreach Initiative, MHPCC supported 2011 summer internships for undergraduate and graduate students attending the DoD military academies and other universities to enhance their knowledge of computational methodologies in HPC. Programs supported by the MHPCC DSRC included HPCMP Military Academy Internships, AFRL/DE Scholars Program, University of Hawaii, Akamai-Center for Adaptive Optics (CfAO), and the Maui Economic Development Board (MEDB) Ke Alahele Intern Program. Universities represented were the U.S. Air Force Academy, U.S. Military Academy, Princeton University, University of Southern California, Michigan Tech, Western Oregon University, and the University of Hawaii. Research projects were individually designed for each intern and their major area of study. HPC was employed in each research effort. Project examples included image processing, algorithm development, computational fluid dynamics, radar modeling and simulation, and other HPC applications.

David Morton
Director, MHPCC DSRC

GNET: Visualizing Real-Time sFlow Data on the DREN at the MHPCC DSRC

By Randy Goebbert and Ross Matoi, Maui High Performance Computing Center DoD Supercomputing Resource Center; and Phil Dykstra, DoD High Performance Computing Modernization Program Office

Introduction

The Graphical Network Exploration Tool (GNET) project at the Maui High Performance Computing Center DoD Supercomputing Resource Center (MHPCC DSRC) is creating tools to collect, analyze, store, and visualize (in real-time) network monitoring data. Providing detailed layer 2 through 7 network monitoring and in-depth traffic analysis is essential in defending against threats to network availability and security. Improved visibility into traffic flows and routes allows for greater control and more effective management of the network. The primary source of data for GNET today is "sFlow." The multivendor sFlow standard was created and is maintained by an industry-wide consortium and provides a method for monitoring high-speed switched and routed networks. sFlow is widely available throughout the product lines of most major manufacturers. The sFlow implementation provides a method for monitoring links at 10 Gb/s+ without impacting the performance of sFlow-enabled switches or significantly impacting network loading. There are over 70 sFlow-enabled switches on the DREN (Defense Research and Engineering Network) today. With the implementation of the Joint Sensor program, more sFlow-enabled switches are coming online. sFlow scales well to the DREN high-speed multisegment topology. Its packet-based sampling scheme randomizes sampling to avoid synchronizing with any periodic patterns in the network traffic. Switches sample packets and send UDP sFlow datagrams back to a central collector. Multiple tools can consume sFlow; DREN currently uses commercial packages from inMon and SolarWinds, neither of which provide the real-time geographic worldwide view developed for GNET.

GNET Implementation

At the MHPCC DSRC, the initial implementation included the demonstration of sFlow collection and visualization on a limited subset of DREN sFlow data. The team was able to create an end-to-end system that receives, parses, stores, and displays the real-time sFlow. The output data consist of Keyhole Markup Language (KML) files suitable for display on Google Earth. KML is an XML-based notation for geographic annotation that is a standard of the Open Geospatial Consortium. The visualization concepts using KML and Google Earth originated with the SNMP-based Planet-DREN project, which was demonstrated at the DREN Networks Conference in 2006. An important design goal was to limit the client-side dependencies to commonly installed, freely available programs. As such, the GNET tools will be available throughout the DREN. As implemented, the only client-side dependency is the installation of Google Earth.

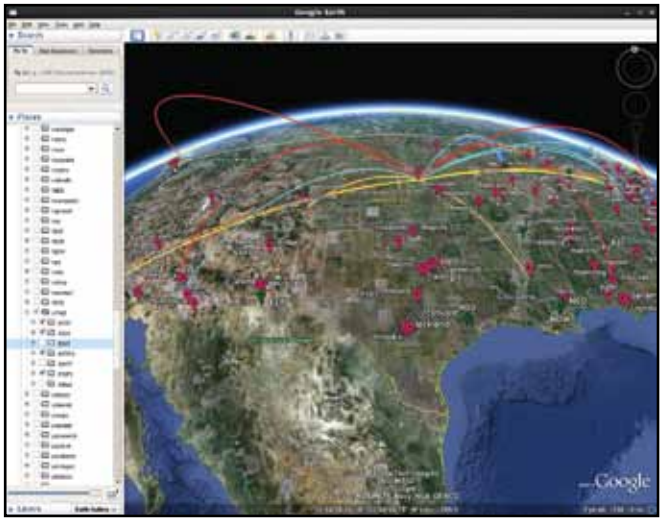

Figure 1. GNET DREN data flow visualization

The end-user is only required to download a single small KML file containing network links to the actual visualization data. Viewable data categories are presented in the "places" column on the left of Google Earth along with the time span of the sampled data. The link data are represented by arcs whose height is proportional to data volumes and color coded based on traffic type. The flows may be filtered by all DREN sites or by common traffic types. Data flows may also be filtered by highest volume end points for each data type (e.g., DNS, HTTPS, SMTP, POP3, etc). The data are updated continuously on Google Earth via the KML network link mechanism. Other Google Earth mechanisms allow additional information to be displayed. For example, a window containing further information may pop up based on a mouse click on one of the arcs. When an individual site is clicked, graphs are generated displaying statistics for that site, including

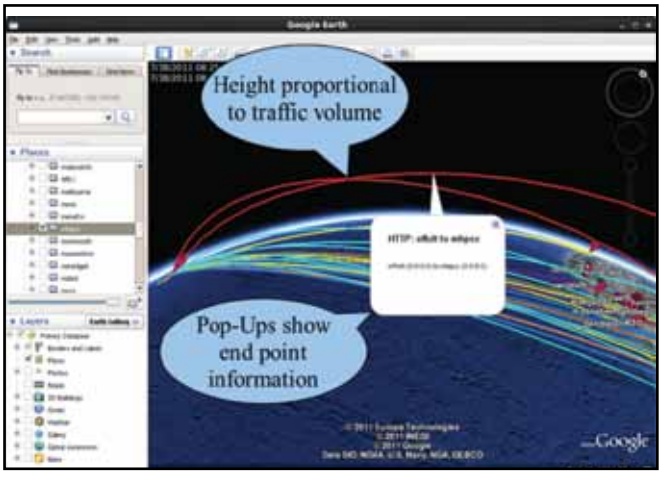

Figure 2. GNET data displays

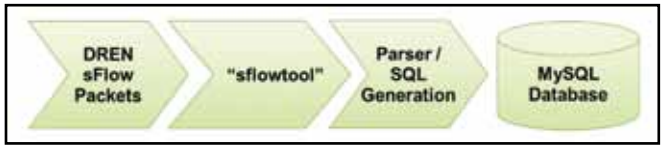

Figure 3. sFlow Ingestion

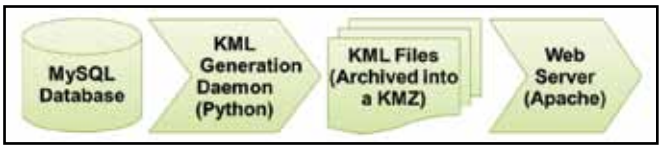

Figure 4. KML Generation and Service

volume by protocol/port and top "talkers" for the site. For an individual site, columnar data are generated and overlaid on Google Earth representing input/output traffic.

Starting with "sflowtool," a freely available program that accepts and interprets the sFlow datagram, Python scripts were created to parse the sflowtool output and insert sFlow data into a MySQL database. These scripts match IP addresses against known DREN subnets that are cross-correlated with site locations in the database.

On the extraction side, a daemon implemented in Python was written to periodically retrieve data from the database, translate the latest data into KML files, and compress those files as KMZ archives. The files are then made available for clients via an Apache web server.

Future GNET Directions

The GNET system is evolving to support route profiling and peering optimization, provide decision-making information for congestion management, and to help understand the application mix on the DREN network. The system can connect to and display other data sources, including the Joint Sensors for security information and events and to the DREN Active Measurement Program for performance. The "Host sFlow" project exports physical and virtual server performance metrics using the sFlow protocol. This makes sFlow an alternative to SNMP for monitoring hosts, allowing GNET tools to receive, process, and display system status across the DREN network.

With GNET, all sFlow data are stored in a relational database. In the future, other data sources may be unified in this database. This provides a mechanism for non-real-time data exploitation for security and audit trail analysis. Custom queries of specific time frames from within Google Earth are planned to provide specific views of historic data flows. These visualizations can take the form of animated KML "movies" that compress longer time windows.

As the full sFlow stream is turned on, and as additional sFlow sources come online, the system must be scaled to handle the expected rates and volumes of data—both on the database insertion side and the KML generation component. Ultimately, a distributed database and query solution may be implemented. For data sources originating outside of the DREN network, a correlation of IP addresses to external geographic regions can be depicted on Google Earth. The team is also researching advanced display techniques such as HTML5/WebGL implementations.

HPCMP Portal Initiative

By David M. Morton, MHPCC DSRC Director

The HPCMP is increasingly interested in ways to promote access and use of HPC. Enhancing effective access and use of HPC resources includes (1) improving agile access to HPC computational and data storage resources by current users, (2) increasing transparent and user-friendly access to these resources by client-based "communities of practice" that have not been historical HPC users, and (3) establishing new ways that HPC can increase project cycle effectiveness and efficiencies in these user communities.

At the direction of the HPCMPO, the MHPCC DSRC has taken a leadership role in expanding HPC support to DoD science and engineering organizations through the use of a web-enabled portal.

"Initially, the MHPCC will focus on the development of the necessary infrastructure needed followed by a transition phase where the MHPCC will start to offer on-demand "portal" –based services to include Matlab and CREATE applications. Over the long-term, additional applications will be added," said Cray Henry, HPCMP Director.

MHPCC is working with the Army Research Laboratory (ARL) DSRC, the USACE Engineer Research and Development Center (ERDC) DSRC, the Computational Research and Engineering Acquisition Tools and Environments (CREATE) Team, and HPCMP leadership to integrate ongoing activities. Initial planning has been completed, and a development path has been approved within AFRL and the HPCMP communities to provide portal capabilities available to users in a timely and cost-effective manner.

HPC resources have been demonstrated to be of great value in supporting science, engineering, and business enterprises. However, historically, HPC use has been confined to specialized groups and has not expanded into other sectors where the derived value of HPC could be significant. This has been due, in part, to stringent access requirements to HPC resources, lack of application software, and limited access to specialized talent and cost constraints. Activities such as the CREATE initiative have focused on overcoming these constraints and

"mainstreaming" HPC capabilities to support clients such as DoD acquisition programs. Increasing access and ease of use of HPC applications has the potential to dramatically expand the customer base both for CREATE applications as well as third party applications such as MATLAB. The portal initiative technical goals are as follows:

- Provide a secure unified access point with single sign on.
- Support an integrated framework with access to decentralized components allowing HPC "jobs" to be run on available HPC resources in response to and in support of applications being served through the portal.
- Provide web or web-like centralized interface for users that require no software installation on user workstations and will work across multiple security enclaves.

Fundamentally, the success of the HPCMP portal initiative will be based on the value that users derive from access and ease-of-use of HPC resources. To this end, the HPCMP portal initiative will include a robust outreach component. The expediency with which these resources can be delivered, user experience, and the ease of use of HPC technology are vital. Ideally, the portal will attract many new users by providing access to HPC through features like the MATLAB drag-and-drop portlets. Successful implementation of the process flow and data models, which help organize and put into automatic configuration

control of the codes, data input, and data output, should attract new and existing CREATE users. With careful integration of the full suite of CREATE applications and the integration of ancillary open-source and COTS tools, the portal will evolve to be the users preferred application delivery choice. The rich set of collaboration tools provided by the selected framework will also help ensure the success of the portal.

MHPCC DSRC Energy Efficient Computing Initiative

By Captain Joseph Dratz, MHPCC DSRC Technical Director

Escalating energy consumption is straining the ability of military data centers to deliver cost-effective support to the warfighter in an environmentally sound manner. Since October 2009, the MHCC DSRC has partnered with several Department of Energy labs through the Federal Energy Management Program (FEMP) to conduct four comprehensive data center efficiency assessments. Over those 2 years, MHPCC has significantly reduced energy consumption through segregation of hot and cold isles, increasing the chilled water set point and installing more efficient chillers and CRAH units. In order to push legacy data center efficiency further, MHPCC partnered with PACOM to propose the Energy Efficient Computing (E2C) JCTD. Planned for execution in FY12 and 13, E2C will demonstrate the ability to retrofit a legacy data center with infrastructure, hardware, and integrated control software improvements to achieve energy efficiency levels seen currently only in newly purpose built efficient data centers. E2C is an ambitious initiative with goals that include reducing the Center's Power Usage Effectiveness (PUE) ratio to an extremely aggressive target of 1.1, achieving Leadership in Energy and Environmental Design (LEED) certification and minimizing grid power requirements through the integration of separately developed renewable power sources. In addition, the MHPCC DSRC initiative will investigate more efficient parallel algorithms to improve the utility of HPC systems and thereby reduce energy costs. Further, E2C will provide a tailorable set of efficiency plans and guidelines to DoD legacy data centers that can be implemented based on available funding.

The E2C technical activities will focus on the following:

- **Software Development** – the software component of the initiative involves two distinct activities: (1) measurement and management of energy consumed by infrastructure and (2) reduction of energy consumption by the HPC systems within the data center down to the chip level. This will involve integrating available COTS center management solutions and analysis of specific HPC user codes to determine the potential for reducing energy consumption through efficient node use.
- **Infrastructure Enhancement** - New technologies are developed continually, and the increasing cost of energy is making some previously infeasible options surface as

viable candidates. Though E2C is still in the initial planning stages, the following options are feasible:

- *Cooling Technology* – Integration of efficient chillers for air cooling and passive warm water-cooling for water-cooled systems.
- *Power Delivery* – Power conditioning, magnetic energy recovery switching (MERS), thermal camera monitoring of PDUs, and offline UPS integration.
- *Energy and Thermal Management* – Comprehensive real-time monitoring of water, plenum, and rack temperatures.
- *Renewable Power* – Integration of planned Maui Solar Initiative photovoltaic DC power.

The successful execution of the E2C JCTD will involve the close coordination of a number of partner organizations and integration of several affiliated projects. Current E2C collaborators include PACOM/J8, ARL DSRC, ERDC DSRC, Pacific Northwest National Laboratory (PNNL), and Lawrence Berkeley National Laboratory (LBNL). The success of E2C largely depends on maintaining a close working relationship with the HPCMPO Technology Insertion process, the MHPCC Research, Development, Deployment & Management contractor, and the key organizations delivering the Maui Solar Initiative. There are inherent challenges in executing a rapid demonstration on Maui with key personnel spread all over the U.S., but the team is mitigating the risk with a 6-month planning process.

Though ambitious, the E2C JCTD will demonstrate technologies and procedures that will substantially decrease the amount of energy consumed at MHPCC. These technologies will be applicable and made available to other DoD HPC and IT sites. Energy is becoming an increasing costly component to the operation of large data centers, and technology exists to significantly reduce the cost and environmental footprint of operations.

Navy DoD Supercomputing Resource Center

From the Director's Desk — Tom Dunn

As the fall comes to a close, the Navy DSRC continues to look forward to the arrival of new, computationally expansive HPC systems in 2012. In the meantime, we continue to be involved in numerous projects that enhance the DoD HPCMP infrastructure and its availability to our users. The Center has successfully put into production the Appro Utility Server (US) and Panasas Center-Wide File System (CWFS).

The Navy DSRC staff recently conducted a large-scale HPCMP-wide analysis of current tape infrastructure versus new tape technologies to determine a timely and cost-efficient storage technology refresh plan. This study aggressively pursued new tape technologies with the goal to limit the required investment in new tape solutions by maximally leveraging the Program's prior T10KA and T10KB media investment and minimizing T10KC investments to allow market conditions to reduce costs to the HPCMP. These newest tape drives double the transport speed, and the tapes themselves can store five times as much data.

More than 15 DoD HPCMP supercomputers serve nearly 4,000 users and generate multiple petabytes of data each year. For several years now, the Navy DSRC has taken on the challenge of providing secure, redundant storage of the HPCMP data. The Center excels at gathering, storing, and protecting these data at a remote storage facility that provides more than 17 petabytes of backup storage, with the capability to store a total of 102 petabytes of data.

We're also providing leadership in the Enterprise System Monitoring (ESM) initiative, establishing system monitoring protocols within existing software to provide comprehensive, automated high performance computing systems checks. The ESM effort also provides a means of gathering numerous system metrics that assist Centers in fine-tuning the HPC infrastructure to best suit HPCMP users' needs.

The importance of real-time high performance computing, particularly for the validation of experimental and forecast models, is well-known to the Navy DSRC and the meteorology and oceanography (METOC) community it serves cycles to on a 24x7 basis. This importance is highlighted in Dr. James Doyle's success story, "Tropical Cyclone Prediction Using COAMPS-TC," on page 2 of this edition of *HPC Insights.* That team's real-time modeling and prediction of the intensity of Hurricane Irene as she raked the East Coast of the U.S. was performed on the Navy DSRC Cray XT5, *Einstein,* using scheduled, highly available resources.

Finally, we look forward to the implementation of the new Defense Research and Engineering Network (DREN III) contract. We anticipate that the Navy DSRC wide area network (WAN) bandwidth capacity will increase significantly. We, along with the DREN team, have recently taken interim steps to increase the Navy DSRC WAN bandwidth by 25 percent for the benefit of our user community.

Tom Dunn
Director, Navy DSRC

USNS Pathfinder is one of the six Oceanographic Survey Ships that are part of the 26 ships in Military Sealift Command's Special Mission Ships Program

DoD HPCMP Users Group Conference 2011

By Rose J. Dykes, Technical Writer, U.S. Army Engineer Research and Development Center DoD Supercomputing Resource Center, Vicksburg, Mississippi

DoD HPC users and personnel from the High Performance Computing Modernization Program Office and its five Supercomputing Resource Centers gathered in Portland, Oregon, for the 21st Users Group Conference on June 20-23, 20ll.

Keynote Speaker Dr. Jeffery P. Holland, Director, U.S. Army Engineer Research and Development Center and Director of Research and Development, U.S. Army Corps of Engineers, Vicksburg, Mississippi, presented "From Discovery to Acquisition: Maintaining Technological Superiority for the Nation's Defense Through HPC" on the first day of the Conference.

Dr. Cynthia Dion-Schwarz, Director, Information Systems and Cyber Security, Assistant Secretary of Defense, Research and Engineering, Washington, D.C., served as the Keynote Speaker on the second day.

Featured Speakers were Cray J. Henry, Director, High Performance Computing Modernization Program, Lorton, Virginia; Steve Wallach, Chief Scientist, Cofounder, and Director, Convey Computers, Richardson, Texas; Dr. Wilfred R. Pinfold, Director, Extreme Scale Programs, Intel Corporation, Intel Labs, Hillsboro, Oregon; and

Dr. Holland presenting the Keynote Address

Dr. Thomas P. Gielda, Chief Technology Officer, Caitin, Inc., Fremont, California.

Over 400 community members attended the conference. Five concurrent computational science tracks covered over 15 topic areas over a 4-day period. The most modern innovations in HPC were presented at the nighttime Poster Session.

Hero Awards were presented to six people, who are featured on the inside back cover.

The following photographs present more insight into the Conference.

ANNOUNCEMENTS

DoD HPCMP Hero Awards for 2011

Presented by Cray Henry at the DoD HPCMP 2011 Users Group Conference, Portland, Oregon, June 22, 2011.

Long Term Sustained
Hugh Thornburg
DoD HPCMPO PETTT

Long Term Sustained
Patricia Hall
NAWCAD 5.4.1
NAVAIRSYSCOM

Up and Coming within the HPCMP
Thad Irby
NRL, Battlespace Environments Institute (BEI)

Technical Excellence
Anders Wallqvist
U.S. Army Medical Research and Material Command, BHSAI

Innovative Management
Jennifer Rabert
Navy DSRC

Technical Excellence
Timothy Campbell
NRL, Battlespace Environments Institute (BEI)

DoD HPCMP Users Group Conference 2012

New Orleans, Louisiana, June 18-21, 2012,
The Roosevelt New Orleans

DOD HPC MODERNIZATION PROGRAM

DoD High Performance Computing Modernization Program

DoD Supercomputing Resource Centers
Networking/Security • Software Applications Support

SUPERCOMPUTING FOR THE WARFIGHTER

The HPCMP expands problem-solving capabilities for researchers and scientists by providing a suite of computational capabilities and services to address modern military and security challenges.

AFRL • ARL • ERDC • MHPCC • NAVY SC11 Edition