HPCinsights

DoD High Performance Computing Modernization Program

Fall 2010
SC10 Edition

EVOLVING TECHNOLOGY

AFRL • ARL • ARSC • 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
Joan Henley
Chuck Abruzzino

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

ARSC DSRC, Fairbanks, AK
Debra Damron
Mary Haley

ERDC DSRC, Vicksburg, MS
Rose J. Dykes

MHPCC DSRC, Maui, HI
Jeff Schmidt
Betty Duncan

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

HPCMPO, Lorton, VA
Deborah Schwartz
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 1

HPC at Work

Responding to Crises—At Home and Abroad 2
Integrated Molecular Design of High Performance Infrastructure Materials 7
Potential Energy Surface Mapping of Energetic Materials Using Coupled Cluster Theory 14
Atomic "Triple-Dog" Dare: Using HPC Resources to Help Create Dipolar Matter 16
Parametric Studies for Superconducting Quantum Interference Filters 19
State-of-the-Art Computational Fluid Dynamics Flow Solver Aids X-51A Research 20
Geo-Temporal Visualization of Information Collected from Large Databases at MHPCC 23
Investigation of Three-Dimensional Internal and External Flow Separation 25
Optimization of Finite Element Codes for Stimulated Brillouin Scattering in Photonic Crystal Fibers 30

DoD Supercomputing Resource Centers

AFRL DSRC

From the Director's Desk – Frank Witzeman 32
AFRL DSRC Prepares for Cray XE6 Supercomputer 33

ARL DSRC

From the Director's Desk – Charles J. Nietubicz 34
ARL TI-09 System *Harold* Dedicated in June 29 Ceremony 35
2010 HPCMP JEOM Faculty Workshop 36

ARSC DSRC

From the Director's Desk – Frank Williams 37

ERDC DSRC

HPCMP Needs an Exascale Computing Initiative Now 38
Practically Green 39

MHPCC DSRC

From the Director's Desk – David Morton 41
Interns from Universities and Military Academies Receive Research Experience at the MHPCC DSRC 42
Maximizing Performance of the WRF Model on *Mana* 43

NAVY DSRC

From the Director's Desk – Tom Dunn 45
What's New with Navy DSRC Systems 46
Learn About the Automated Compliance Checking Project 47

DAAC

DAAC Wins Department of Energy OASCR Award 48

Announcements

DoD HPCMP Hero Awards for 2010 49
DoD HPCMP Users Group Conference 2011 49

About the Cover: The collage on the cover depicts how advances in material development, nanotechnology, and high performance computing affect all aspects of combat readiness and give the warfighter a competitive edge on the battlefield. Advances in material development technologies are necessary to achieve the performance improvements required for future battlefield systems.

DoD High Performance Computing Modernization Program

By Cray Henry, Director

Cray Henry
Director, HPCMP

The Department of Defense High Performance Computing Modernization Program (HPCMP) entered into its seventeenth year in 2010, successfully achieving our mission to accelerate the development and transition of advanced defense technologies into superior warfighting capabilities. To meet our mission, we acquire, deploy, operate, support, and maintain best-value supercomputers; we develop and deploy parallel software; and we operate the Defense Research and Engineering Network (DREN), enabling effective access to supercomputers and to distributed S&T/T&E computing environments.

In 1993, the HPCMP delivered over 180 gigaFLOPS (one thousand million FLoating-point OPerations per Second) of computing capacity. In 2010, the Program delivered over 1,158,000 gigaFLOPS of computing capacity, a factor of over 643,333 percent improvement! Today, the DoD has access to some of the world's most powerful supercomputers and to a variety of computing architectures, chosen to best meet the identified requirements of the Department. We continuously educate the RDT&E workforce with the knowledge needed to employ computational modeling efficiently and effectively, while promoting collaborative relationships among the DoD computational science community, the national computational science community, and minority-serving institutions.

This past year, we have had several accomplishments of note. The HPCMP enabled rapid solutions to critical DoD research and engineering problems, which would have been difficult or impossible to solve without computational technology, by delivering 933 million computing hours of capability to 4200 scientists and engineers working on 400 projects. Among these projects are Design of Very High-Strength Carbon Nanotube Fibers, Integrated Docking Pipeline for the Prediction of Large-Scale Protein-Protein Interactions, Virtual Prototyping of Directed Energy Weapons, high-fidelity computational modeling to support compatibility, certification and clearance process for various aircraft-store configurations, the two-way Coupled Atmosphere-Ocean Model, and predicting ship response to underwater explosions.

We enable technology demonstrations and distributed test events to occur throughout the DoD by delivering 30 gigabytes per second across the DREN, supporting reconfigurable connections to over 200 research and test facility locations. This was all done in a cost-effective manner, at the rate of approximately 26 cents per computational hour, which is about 28 percent of current commercial rates!

Our major focus continues to be on operation and maintenance of our six DoD Supercomputing Resource Centers (DSRCs), operation and security of the DREN, development of software applications, and training and education services to our users.

This issue of *HPC Insights* is published in time for the SC10 Conference. It is particularly fitting, as experts in the world of supercomputing are gathered to share ideas and accomplishments, and look to the future of broadening the capabilities of their world and sharing them with the user community.

The HPCMP strengths are outlined in the articles that demonstrate our ongoing efforts at providing the latest computational resources to the DoD and the computational science and engineering community.

Responding to Crises—At Home and Abroad

By Christine Cuicchi, Computational Science and Applications Lead, Navy DoD Supercomputing Resource Center

The Navy DoD Supercomputing Resource Center (DSRC) is no stranger to assisting the Navy in supporting the local and global community in times of trouble, as the operational modeling team at the Naval Oceanographic Office (NAVOCEANO) is often called to provide service above the normal call of duty. During a typical day, the NAVOCEANO team runs a series of ocean circulation and wave forecast models on Navy DSRC high performance computing (HPC) resources to support Navy operations worldwide on a 24x7 basis. These models include the Global Navy Coastal Ocean Model (G-NCOM) that provides a daily 3-D prediction of ocean structure from pole to pole, surface to bottom at a horizontal resolution of 1/8 degrees (14 km or 7.5 nautical miles). These models have been used to assist humanitarian efforts in previous crises and, most recently, to provide assistance during the oil spill recovery period after the Mississippi Canyon 252 incident that occurred on April 20, 2010.

Extraordinary efforts that the operational modeling team has supported in the past and present using Navy DSRC HPC systems are as follows:

2001: NAVOCEANO modeling community produces high-resolution and time-dependent weather and ocean forecast products on Navy DSRC HPC systems in support of the U.S. Navy's effort to recover the Japanese fishing ship *Ehime Maru* off the southern coast of Oahu, Hawaii.

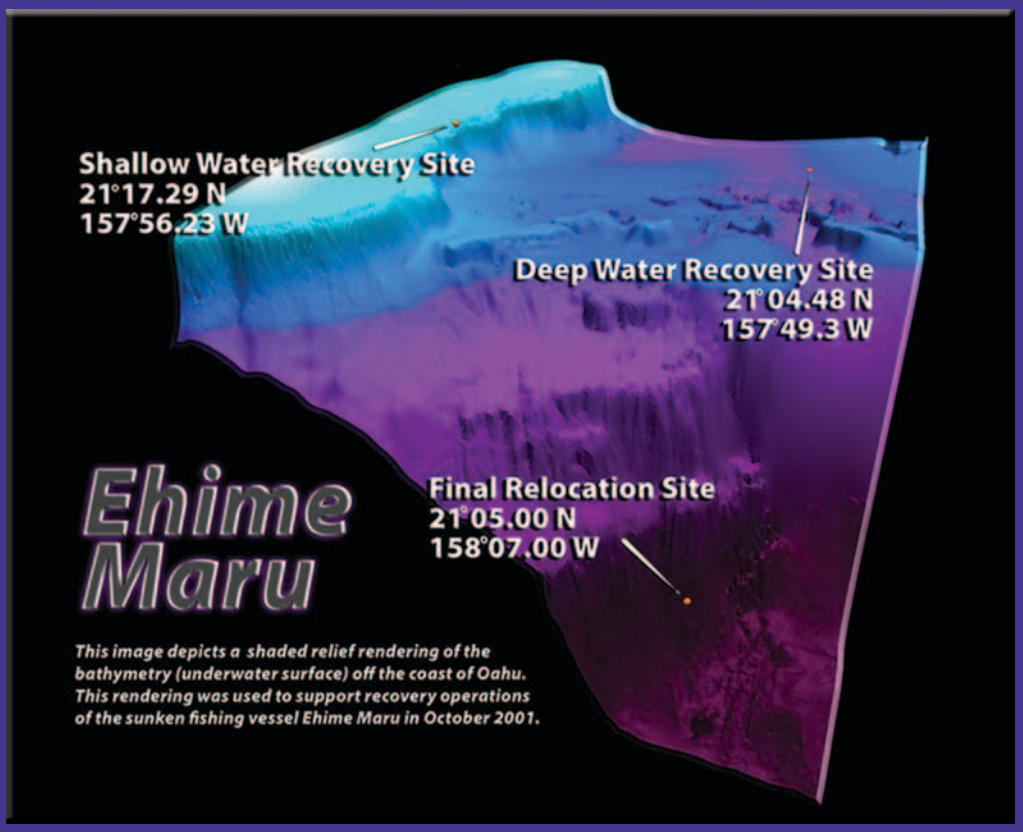

Figure 1. Bathymetry surrounding the Ehime Maru recovery site

2004: In the aftermath of the Indonesian Tsunami, NAVOCEANO oceanographers turn to Navy DSRC HPC capabilities when asked to assist in the tsunami rescue effort. Computational models are used to forecast conditions that affected rescue efforts and recovery operations in Sumatra, Sri Lanka, and the Maldives. Coastal currents, waves, and surf forecasts were used to plan the landing of rescue parties and supplies, and the location of the westward current off the northwest tip of Sumatra (at about 6°N on Figure 2) was used to search for victims who had been swept offshore.

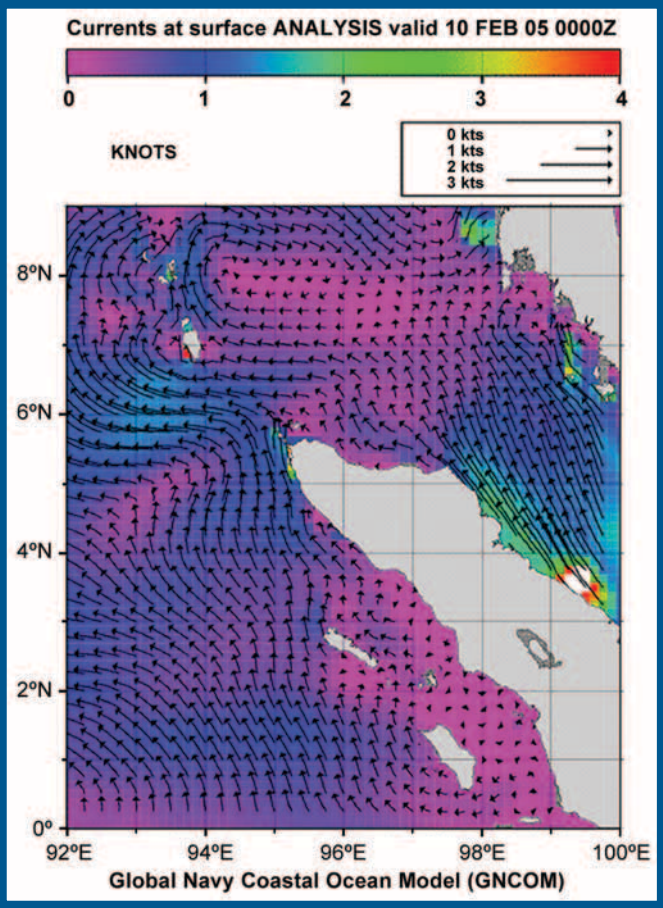

Figure 2. G-NCOM Surface Current Analysis near island of Sumatra, February 2005

2005: Hurricane Katrina devastates the U.S. Gulf Coast. Despite significant to catastrophic damage to their homes in Mississippi and Louisiana, DSRC personnel work around the clock to get the Center's HPC systems—which remained available throughout the storm—reconnected to the Defense Research and Engineering Network (DREN) via satellite-based communications within just a few days of the hurricane's landfall. The Center's team ensures that the NAVOCEANO modeling community is able to compute and deliver their time-critical oceanographic products to the Navy's fleet without significant delay, and that the High Performance Computing Modernization Program (HPCMP) users are able to use the systems as well.

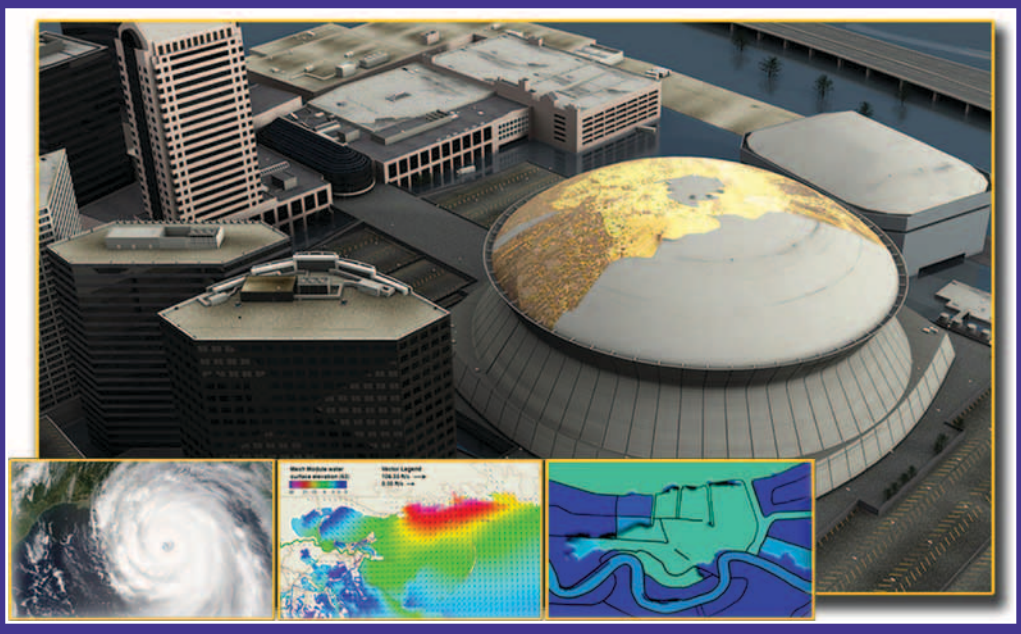

2009: G-NCOM is used by the NAVOCEANO team to provide hindcast drift estimates for the area of the Air France Flight 447 crash in the Atlantic Ocean near Brazil. These estimates helped to determine the possible paths of aircraft debris postflight and suggest where the plane went down (it was never found). NAVOCEANO has provided similar forecasts for a number of military and civil aircraft recovery efforts including the 2007 Adam Air loss in Indonesia (the black box was found) and a private plane near Trinidad in 2009 (also recovered)

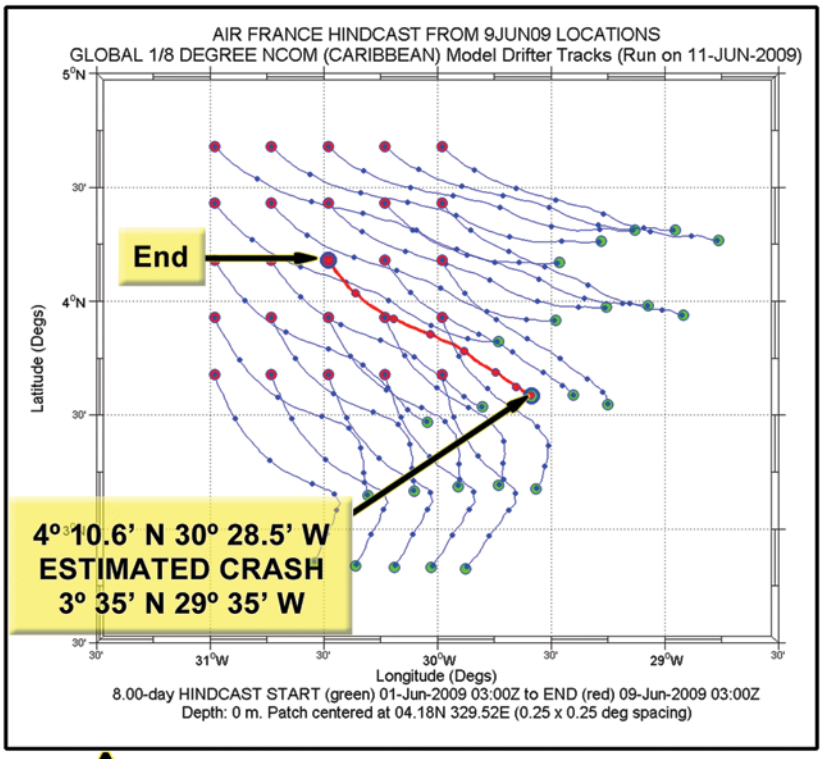

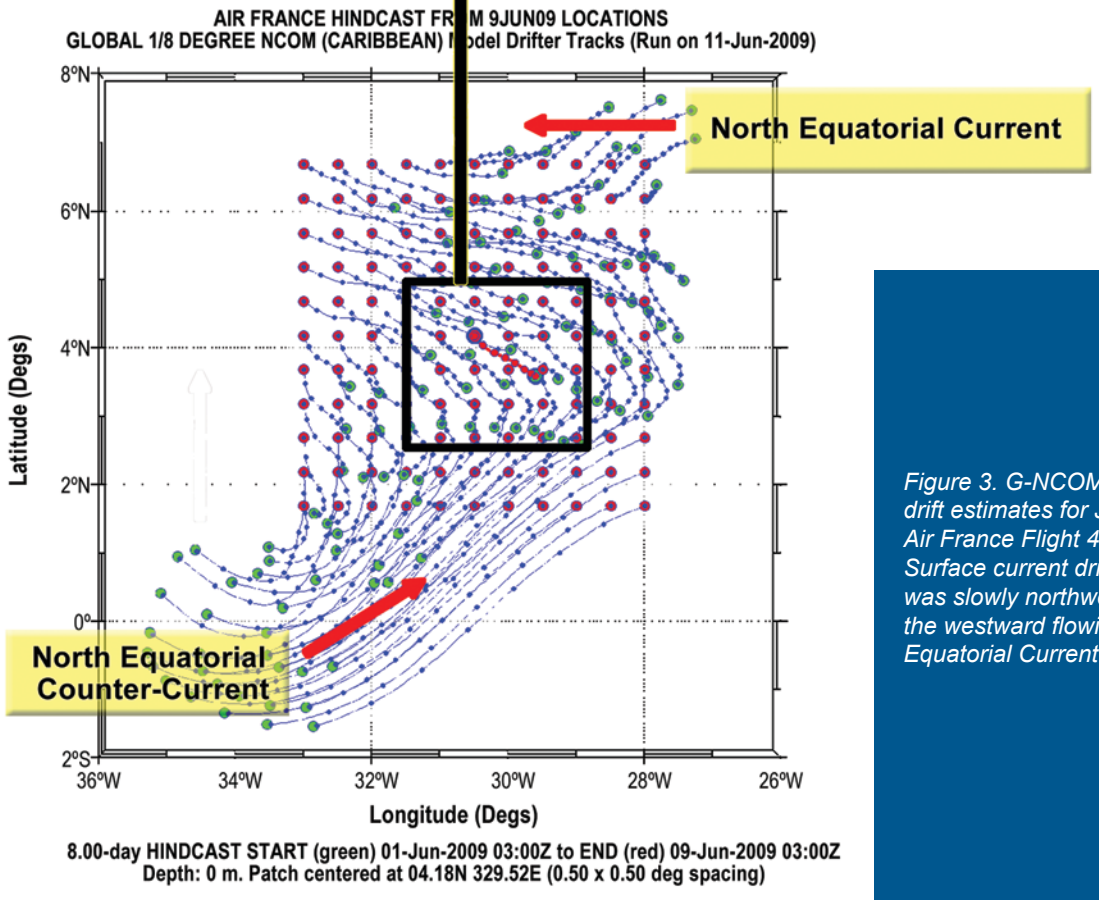

Figure 3. G-NCOM hindcast drift estimates for June 2009 Air France Flight 447 crash. Surface current drift in this area was slowly northwestward into the westward flowing North Equatorial Current (NEC)

2010: In the Gulf of Mexico, ocean current fields from G-NCOM become an integral part of the National Oceanic and Atmospheric Administration's (NOAA) Office of Response and Restoration (OR&R) oil spill trajectory forecasts for the Mississippi Canyon 252 incident. In special cases, NAVOCEANO is often tasked with providing higher resolution coastal forecasts, and for this incident, has implemented a 1/36 degree (3 km or 1.7 nautical miles). American Seas NCOM is short order. The AMSEAS-NCOM has been declared operational and is currently running on the Navy DSRC HPC systems.

Figure 4. Surface current analysis and sea surface temperatures in the Mississippi Canyon 252 area, June 2, 2010. Clockwise flowing Loop Current is "pinching off" and will move slowly westward.

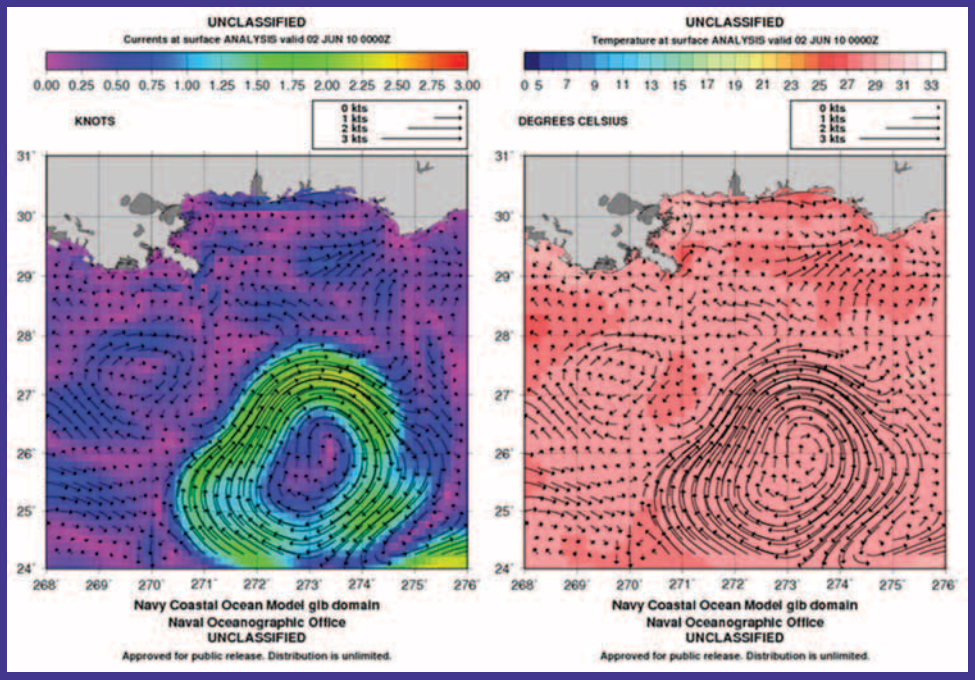

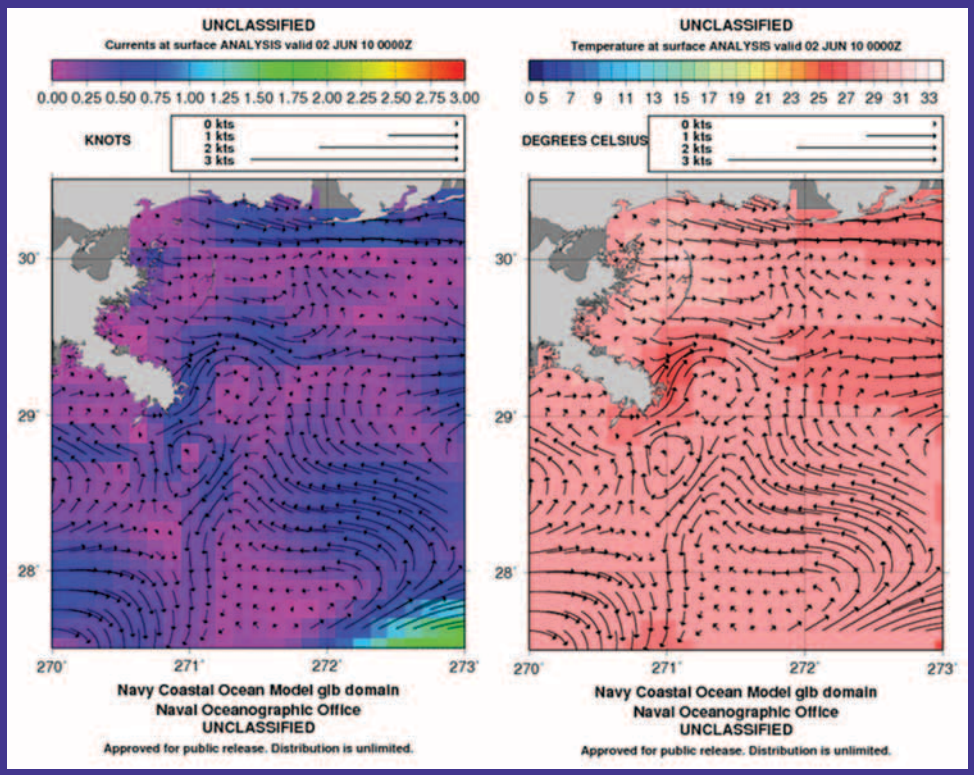

Figure 5. Close-up of Gulf of Mexico and Mississippi Sound surface currents and sea surface temperatures.

2010: (Continued).

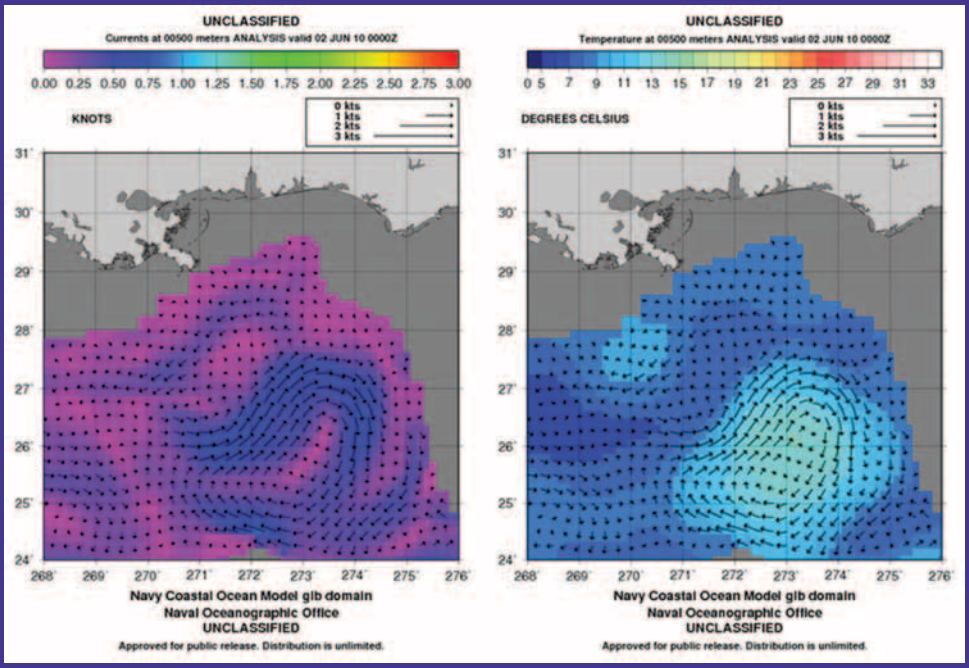

Figure 6. Shown are sea temperature and currents at 500 meters (1650 ft). Model fields in Figures 4, 5, and 6 are being used by the NOAA Office of Response and Restoration to predict the movement of oil at various levels in the ocean

In addition to our technical response, the parent command of the Navy DSRC, the Naval Meteorology and Oceanography Command (CNMOC), has made available time and training to allow employees to volunteer with local oil spill response efforts. Navy DSRC personnel have recently taken advantage of this opportunity to assist with local wildlife rescue and recovery efforts, and with the distribution of food and supplies to families affected by the spill.

Navy DSRC and NAVOCEANO METOC Team

Pictured with *Albert*, a Cray XT5 test system, are Eileen Jones, Christine Cuicchi, and Dr. Frank Bub. The three communicate on a daily basis to ensure that the operational modeling forecast products developed on the Navy DSRC HPC systems are delivered to the Navy fleet several times a day every day of the year.

- **Christine Cuicchi** is the Computational Science and Applications Lead for the Navy DSRC and serves as the Center's liaison to the METOC modeling community. She began her 13-year career in the HPCMP as a member of the Computational Science and Engineering Group at the ERDC DSRC from 1997-2000, after which she moved to the Navy DSRC.
- **Dr. Frank Bub** is the Technical Lead of the Naval Oceanographic Office (NAVOCEANO) Ocean Prediction Department's Model Operations and Global Modeling Group. After a 20-year career as a Navy oceanographer, Dr. Bub spent 13 years as part of the research faculty at the University of New Hampshire and the University of Massachusetts Dartmouth. He has been a part of the NAVOCEANO team since 2002.
- **Eileen Jones** is the Model Operations Team Lead within the Model Operations and Global Modeling Group and serves as the team's liaison to the Navy DSRC. A former contractor with the Navy DRSC, she has been with the NAVOCEANO team for 7 years.

Integrated Molecular Design of High Performance Infrastructure Materials

By Dr. Charles F. Cornwell, U.S. Army Engineer Research and Development Center Information Technology Laboratory

Materials underpin almost all technology. The pace of material development had reached a plateau. Today this is changing because of nanotechnology. At the U.S. Army Engineer Research and Development Center (ERDC), we are using nanotechnology to develop high-performance infrastructure materials, materials with 5 to 10 times the strength, strength-to-weight ratio, or stiffness-to-weight ratio of existing infrastructure materials. Research efforts such as ours will have a profound effect on society over the next 2 decades.

Nanotechnology is defined as the study and control of matter on an atomic and molecular scale. Exploiting nanotechnology requires atomic-molecular level material design and the ability to synthesize these materials with atomic-molecular level precision. Building materials with nanoscale precision poses formidable theoretical, computational, and experimental challenges to developing advanced materials. Recently developed experimental techniques allow researchers to build, manipulate, and test the properties of objects with nanometer-scale dimensions. Advances in theoretical and computational techniques, coupled with high performance computing (HPC), allow the properties of these materials to be predicted prior to actually synthesizing the material. They also allow us to understand the causes for material weakness, and the computational methods can guide the design of materials to overcome these weaknesses.

In the past, the scientific disciplines have dealt with relatively different length scales: chemistry focusing on the bonding in and between atoms, solid state physics dealing with how these atoms assemble themselves, and metallurgy with the everyday manifestation of those assemblies. At the nanometer scale, these boundaries among the disciplines begin to blur. Advances in theoretical, computational, and experimental techniques are converging at the nanometer scale. Now that the scales of various disciplines are intersecting, cross-disciplinary efforts are increasingly more productive. ERDC has initiated a cross-disciplinary program that leverages the DoD HPC capabilities with ERDC's modeling and simulation and experimental capabilities.

High performance computing plays a key role in advancing our understanding of the behavior of new materials. Realistic numerical simulations provide both quantitative and qualitative insights into the properties of materials at the nanometer scale, and developments are occurring to allow us to extrapolate this understanding to the macroscale. The models and simulations are sometimes thought of as taking middle ground between theory and experiment, with commonalities to both. The calculations are based on our theoretical understanding of the interactions between atoms and molecules. However, the results are often analyzed in much the same way as experiments. Classical theory relies on analytical and numerical approximations that reduce the answer to complex problems to a model tractable in a closed-form solution. This process restricts analytical solutions to highly idealized problems that are difficult to interpret in terms of experimental results. Numerical methods and computers solve problems of a more complex nature by replacing the need for approximation with sophisticated and effective calculation schemes. Advances in experimental techniques and the ability to work with increased levels of computational complexity extend the range of problems that can be simulated to encompass real systems. Theoreticians are no longer restricted to working only with idealized models that are amenable to closed-form solutions but can deal directly with detailed and highly realistic problems that can aid in the understanding and extrapolation of experimental results.

The ERDC program exploits this new capability and relies heavily on a tight integration between simulation and experiment to design and process new materials with nanometer-scale precision. There are three components to the program: (1) atomistic modeling and molecular design, (2) material synthesis and nanoscale diagnostics, and (3) advanced structural concepts and macroscale (continuum) material modeling. The atomistic simulations are conducted at the ERDC Information Technology Laboratory (ERDC/ITL), and the simulations are run using the DoD High Performance Computing Modernization Program (HPCMP) resources at ERDC. The ERDC Construction Engineering Research Laboratory (ERDC/CERL) conducts the material synthesis research component of the program in order to build the new materials. A companion research effort by the ERDC/ITL and the Natick Soldier Research Engineering and Development Center (Natick) develops new structural concepts in anticipation of the unique properties of the advanced materials being developed. One such concept, called an Inverse Triaxial Structural Element (ITSE), would use the novel high-strength materials for base camp structures.

In this strategy, the three components are integrated, and material design and development proceeds in an iterative fashion. The physics-based material models are validated using the results from material synthesis and nanoscale diagnostics. The physics-based material models then investigate the properties and performance of novel material designs. Macroscale material models use material synthesis information to analyze the manufacturing process. They also use the material properties obtained from experiment and simulation to design and investigate the performance of advanced structural concepts.

In recent years, significant progress has been made in modeling the evolution of materials during processing. These advances have primarily been in the area of application of continuum methods where empirical constitutive relations are used to describe how a material will behave. These methods can predict the final shape of a specimen following deformation processing and its temperature history. However, these models have had little impact in optimizing material properties because they do not quantitatively describe, predict, or manipulate the internal structure of the material.

Material properties such as strength, stiffness, friction, corrosion, and heat transfer are strongly influenced or controlled by phenomena at the nanometer scale. Traditional macroscopic empirical studies shed little light on these phenomena, leaving out the fundamental insight needed to understand material behavior caused by the interaction of atoms and molecules.

The initial research effort of the program is to design and develop a laboratory sample of a 1-million psi (6.89 GPa) carbon-nanotube-based fiber. Under this program, ERDC is expanding its synthesis and nanoscale diagnostics capabilities in nanotechnology research (Figure 1). In 2008, the Synthesis Team in collaboration with Professor Strano's group at the Massachusetts Institute of Technology Institute of Soldier Nanotechnologies (MIT-ISN) discovered the major reason for the termination of carbon nanotube forest growth during the synthesis process. Through a series of refinements to the synthesis process (Modified Ferrocene Catalytic Chemical Vapor Deposition), multi-walled carbon nanotube forests were grown to a length of 3.5 mm. This is possibly the current record within the DoD community, and equivalent to the world

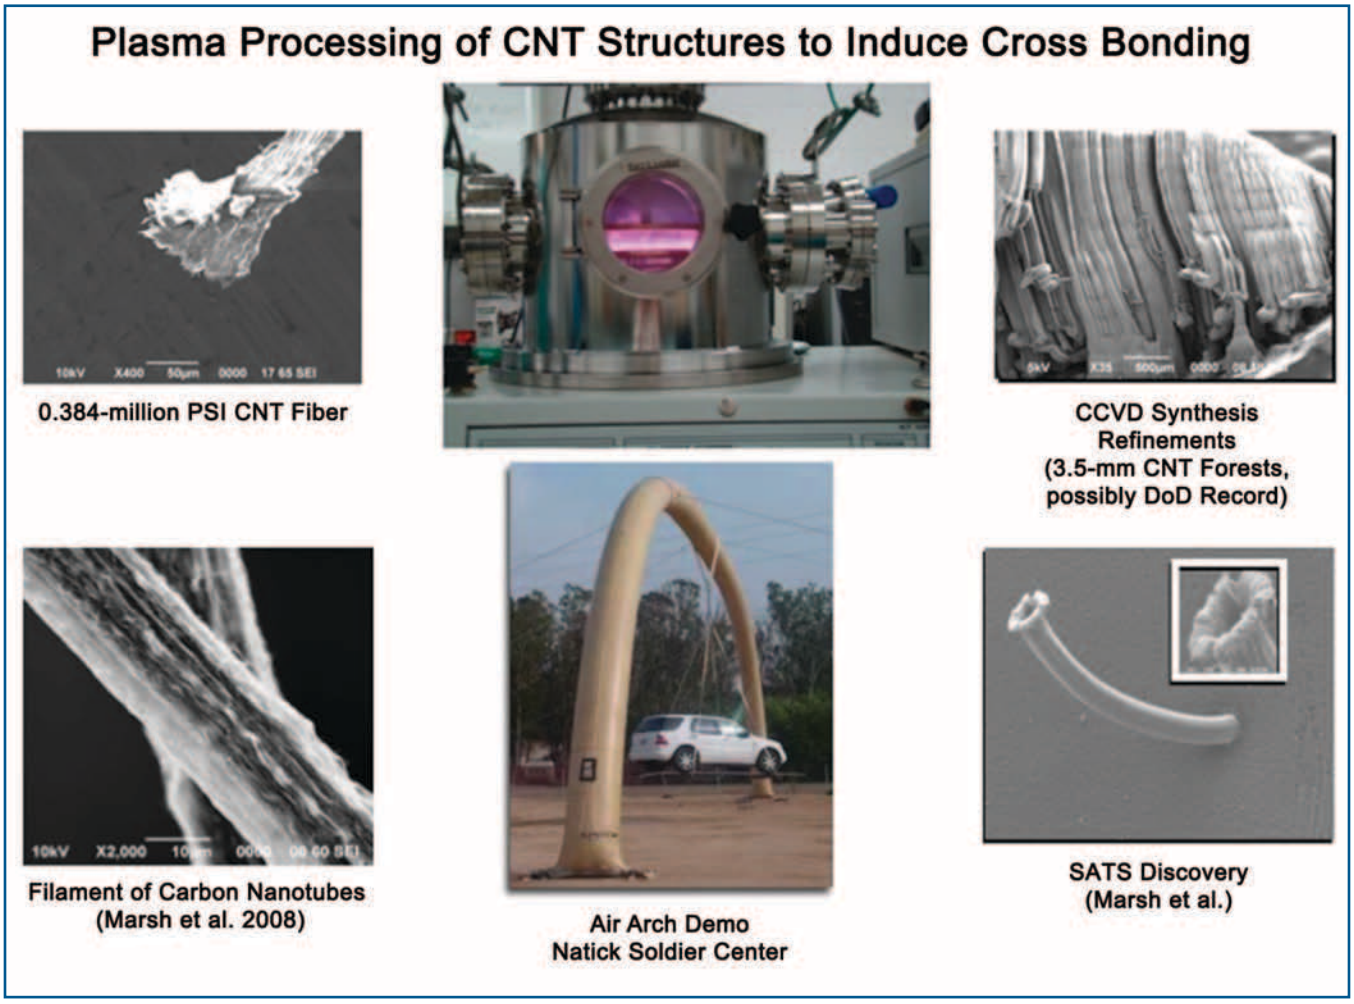

Figure 1. Plasma processing of carbon nanotube structures to induce cross bonding

record of a few years ago. The Synthesis Team also discovered a new form of carbon nanotube structure, the self-assembled tube structure (SATS), in January 2009. In June 2009, methods of assembling carbon nanotubes into structures were demonstrated using single-stranded DNA-ligase. In a further collaboration with Professor Strano's group at the MIT-ISN, the Synthesis Team used the dielectrophoresis synthesis method to produce carbon nanotube filaments that had a tensile strength of 0.384 million psi. These filaments, while extremely strong, will likely be further strengthened via interstitial carbon atom cross bonding. These material discoveries were characterized and validated by the Synthesis Team through nanoscale imaging and spectroscopy techniques with recently acquired ERDC equipment or equipment being used by the Synthesis Team at the University of Illinois.

The Modeling and Simulation Team provides support to the Synthesis Team in their efforts to develop a 1-million psi carbon-nanotube-based fiber. The simulations complement and help guide experiments by assisting with the interpretation of experimental results and pointing out the most promising molecular design avenues to explore. The Material Synthesis and Nanoscale Diagnostics Team at ERDC/CERL provides the initial input for the simulations. The simulations use experimental data as the starting point to validate the physics-based material models. The Atomistic Modeling Team uses atomistic simulations to develop and evaluate new molecular design and materials synthesis techniques. Through atomistic simulations, we have identified both successful and unsuccessful molecular designs to meet project goals, saving considerable time and resources from the traditional build-test, rebuild-retest cyclic method of material development.

Carbon nanotubes form the basic building block of the fibers. A carbon nanotube can be viewed as being a tube rolled from a graphene sheet, with the role direction and the circumference of the roll defining the chirality vector. There are three general chiralities of carbon nanotubes, as shown in Figure 2: armchair, zigzag, and chiral. Chirality determines many of their properties such as tensile response and electrical properties.

Atomistic simulations were used to determine the effect of defects on the strength and elastic modulus of carbon nanotubes. Figure 3 shows the resultant tight binding molecular dynamics- (TBMD-) predicted stress-strain curves and moduli for the pristine (5,5) nanotube along with the three defect cases. In all cases, these single-walled nanotubes (SWNTs) exhibited nonlinear stress-strain response, followed by brittle failure at a critical strain. The pristine nanotube showed a

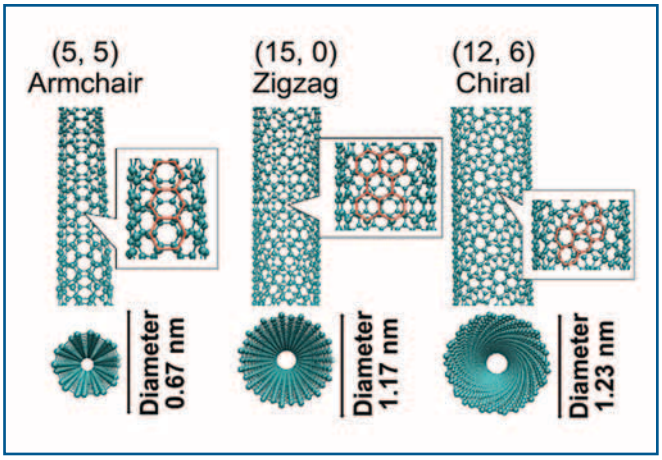

Figure 2. Three general types of carbon nanotubes: armchair, zigzag, and chiral

zero-strain Young's modulus of ~ 160.0 million psi (1100 GPa), a rupture stress of about 15.5 million psi (110 GPa), and critical strain of about 17 percent. This modulus is about five times that of steel. The Stone-Wales defect, which has a 6-carbon atom ring replaced by a 5-carbon atom ring, showed nearly the same stress-strain relationship, but failed at a lower critical strain of about 15 percent. In contrast, the nanotubes with vacancy pairs failed at 11 percent strain for the clustered vacancy pair, and 13 percent strain for the dispersed vacancy pair. The TBMD-calculated stress at failure for these single nanotubes ranges from 11.0 to 15.0 million psi (75.8 GPa to 103.4 GPa). The results (Figure 3) showed that even carbon nanotubes with common molecular defects have much greater tensile strength than the program goal of 1 million psi.

While the strength and stiffness of carbon nanotubes are extremely high, fibers composed of aligned carbon nanotubes are far weaker. The Simulations Team performed simulations of neat fibers of aligned carbon nanotubes to determine the stress-strain characteristics of the fibers. The simulations showed that the fibers fail because the relatively weak van der Waals forces between molecules allow carbon nanotubes to slip past each other before the shear forces reach the intrinsic breaking strength of the carbon nanotubes. To overcome this limitation, covalent interstitial carbon atoms, that is, a carbon atom inserted between the strands of the fiber, were introduced. The interstitial carbon atoms form chemical bonds between the strands of the fibers to facilitate load transfer between the carbon nanotubes and thus prevent slipping. Such chemical bonds can be produced by irradiating carbon nanotubes fibers with high-energy particles. However, if the radiation is too energetic or applied for too long, the mechanical integrity of the carbon nanotubes is compromised. The Atomistic Simulations Team conducted simulations to test the feasibility of producing chemical bonds to

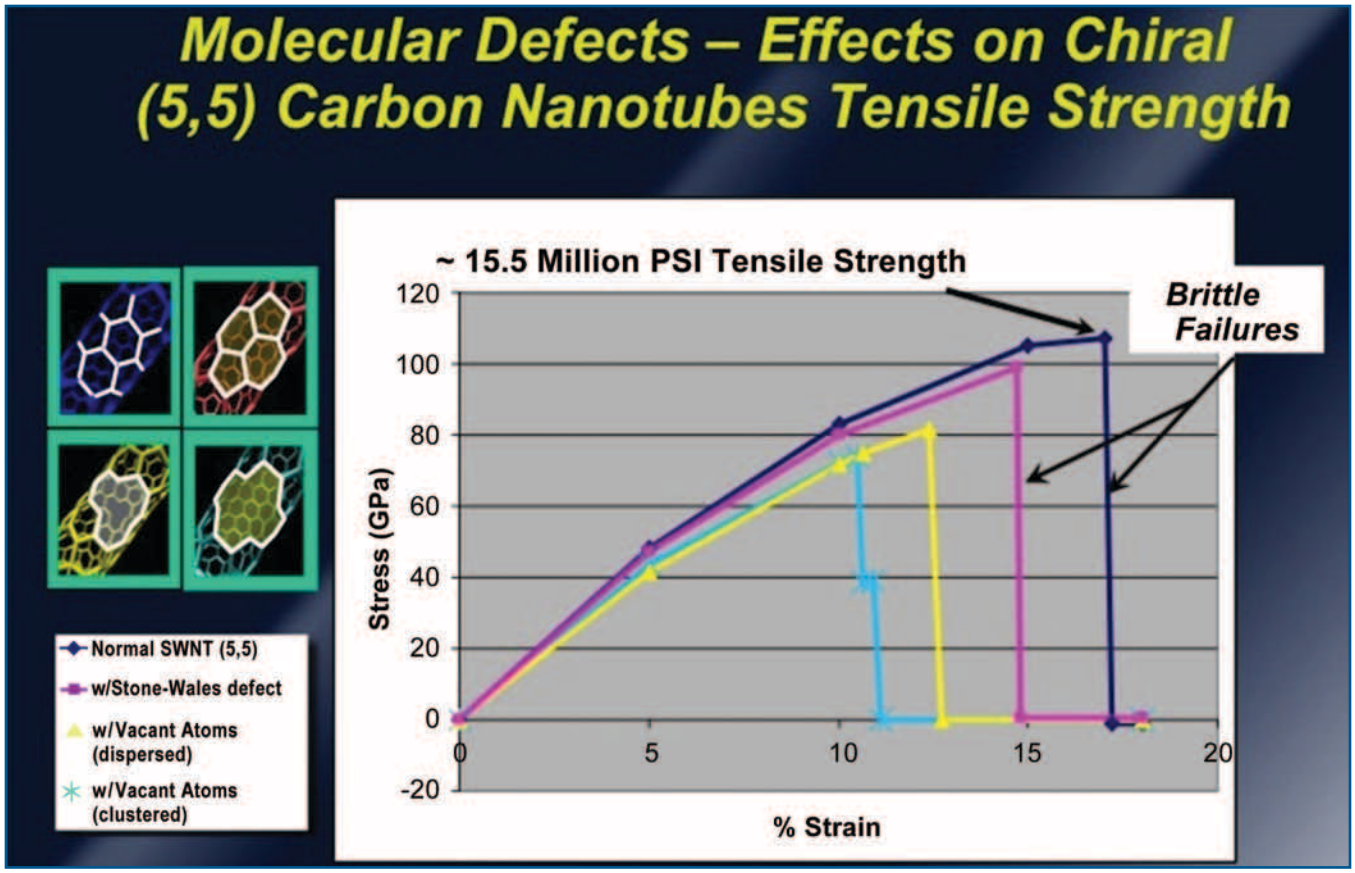

Figure 3. Tensile stress-strain curves predicted for pristine (5,5) and defective SWNTs: pristine tube (solid), Stone-Wales defect, two vacant dispersed, and two vacant atoms clustered

facilitate load transfer between the carbon nanotubes and thus strengthen the fibers.

High-fidelity digital fiber models were constructed based on experimental information to evaluate the optimal fiber length and density of cross-link atoms. Experimental observations indicate that the strands of a carbon nanotube fiber occur in hexagonal close-packed (HCP) fibers of parallel tubes of similar radius. To better understand the failure mechanics of carbon nanotube fibers, we construct fibers of parallel (5,5) carbon nanotubes, arranged in an HCP configuration. Figure 4 shows the fibers from several angles and the distribution functions used to select the carbon nanotube length. The longitudinal axes of the fibers are initially centered on the z-axis. The strands consist of parallel carbon nanotubes placed end to end and are arranged parallel to one another. Each fiber has 19 strands arranged in an HCP configuration with a gap between strands equal to 3.33 Å. Experimental results indicate that the carbon nanotubes within the fiber are randomly distributed. That is, there is no correlation between the z-coordinate of the different tubes, and they most probably have random azimuthal orientations. The fibers are constructed using a random distribution of carbon nanotube lengths, and each carbon nanotube is given a random rotation about its longitudinal axis between zero and 2π. The simulations were run at 300 K.

To avoid extremely short or long carbon nanotubes, the range of acceptable lengths are restricted to the average carbon nanotube length plus or minus the full width at half maximum of the Gaussian distribution function. In particular, we wanted to ensure that none of the carbon nanotubes spanned the entire length of the fiber. The cross-linking atoms were given a uniform distribution along the length of the fiber. The distribution of cross-links in the xy-plane showed a preference for forming bonds along a line representing the point of closest approach of the strands. The smallest fiber had 300,575 atoms, and the largest fiber had 2,421,163 atoms.

Figure 4 shows a carbon nanotube fiber with one HCP layer and carbon atoms forming chemical bonds between adjacent carbon nanotubes. Note that the atoms between the carbon nanotubes and the atoms that form the caps of the carbon nanotubes have a higher strain energy, as depicted by the color scale where blue represents low-strain energy and red represents regions of higher strain.

Results indicate that including cross-link atoms between the carbon nanotubes in the strands increase the load transfer between the carbon nanotubes and prevent

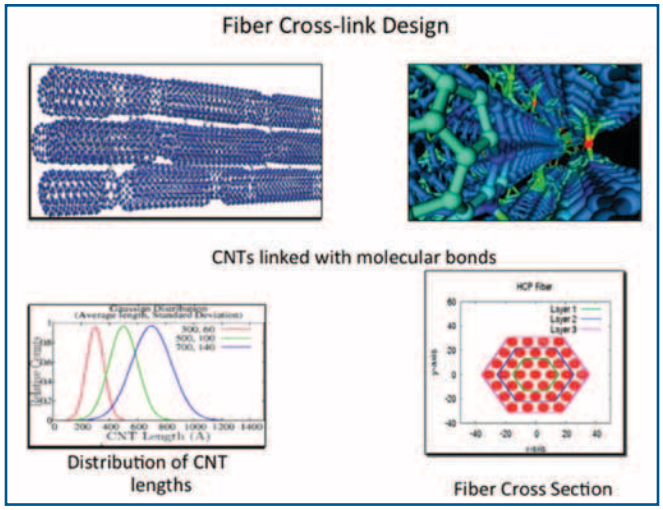

Figure 4. Fiber Cross-link Design: Carbon nanotube fiber from several perspectives. Distribution function gives some indication of the carbon nanotube lengths in the fibers

them from slipping. This increases the initial elastic modulus and yield strength of the fibers. The simulations show an increase in the initial elastic modulus, yield strain, and yield strength with an increase in the average carbon nanotube length (e.g., contrast group C results compared with group B in Figure 5) and cross-link concentrations with the values approaching an upper limit for the longer fibers and higher concentrations of cross-links. Visual inspection of the simulations shows that for the shorter fibers and lower concentrations of cross-links, the fibers fail at the cross-links. For the longer fibers and higher concentrations of cross-links, the fibers fail because the carbon nanotubes begin to pull apart. The yield strength predicted for the pristine (5,5) carbon nanotube (Figure 5) is 110 GPa. From the yield strength of the (5,5) carbon nanotube in Figure 5, it can be seen that the load transfer provided by the cross-link atoms allows the fiber to retain a substantial portion of the yield strength of the constituent carbon nanotubes in the fiber.

The stress-strain relationships exhibited in Figure 5-C probably represent the maximum strength one can expect to obtain from a carbon-nanotube-based fiber and indicate that carbon nanotube fibers with covalent cross-linking bonds can support stresses well in excess of 1 million psi (6.89 GPa). The actual initial elastic modulus and strength of the fibers will depend on the nature of the interaction between the fibers and other factors such as the number and distribution of cross-link atoms and defects that are produced during processing.

A variety of experimental and computational results found in the literature elucidate the effects of cross-linking on the elastic properties of SWNT fibers and multiwalled nanotubes (MWNTs) (Garg and Sinnott, 1999; Peng et al., 2008; Pregler and Sinnott, 2006; Vodenitcharova et al., 2007; Liu and Qin, 2005; Ni and Sinnott, 2001; Huhtala et al., 2004). In 2008, Peng et al. reported results for experimental measurements of single-shell failure for MWNTs that display fracture strengths of 100 GPa and fracture strains that are close to theoretical predictions of structures containing only small defects such as vacancies or Stone-Wales defects. Pregler and Sinnott (2006) examined the effects

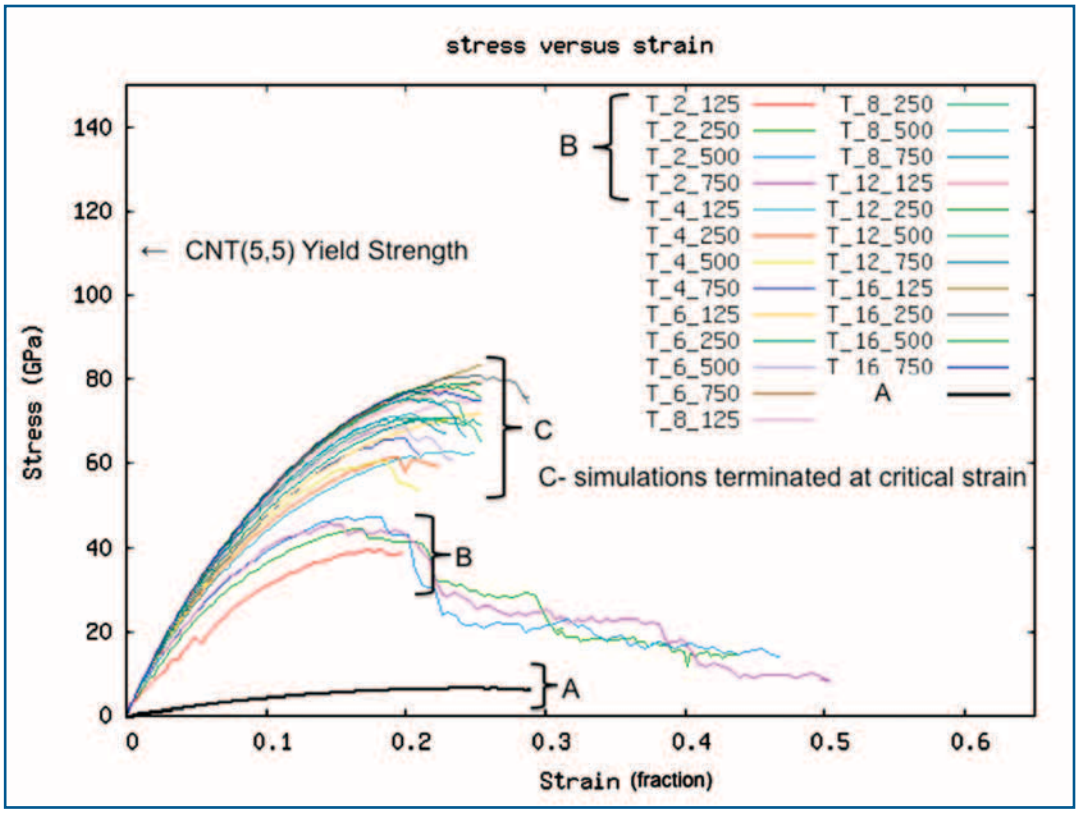

Figure 5. Experimental carbon nanotube tests: Stress-strain curve results. All of the simulations not in group A or B are from group C

of CF^3, Ar, and electron beam irradiation to promote cross-linking between the shells in MWNTs. The results show that the type and distribution of cross-links have a significant effect on the load transfer between carbon nanotubes. Their results indicate how low-irradiation intensity can reduce failure in MWNTs while allowing them to retain their structural integrity. The trends in experimental measurements and theoretical calculations are beginning to converge and indicate that irradiation is a viable method for producing cross-links between carbon nanotubes to promote load transfer while allowing them to retain their structural integrity.

Based on the findings of the Simulation Team, the Synthesis Team is now conducting experiments to optimize the creation of interstitial bonds between carbon nanotubes by electron beam irradiation (Figure 1). The MD simulations revealed that the required length of the carbon nanotubes to achieve the design goal is significantly less than current growth capabilities, so the challenge is the creation of high enough density of interstitial carbon atom bonds. Ultimately, the ability to produce carbon nanotube fibers with a minimum number of defects while optimizing the number and distribution of cross-links will play a major roll in determining the properties of carbon nanotube fibers. With current trends in experiment and calculations, it is reasonable to expect that fibers with tensile strengths in excess of 1 million psi (6.89 GPa) will be available in the foreseeable future. The expectation is that the simulation results are qualitatively correct and can provide insight into the stress-strain relationship of fibers constructed from an ensemble of carbon nanotube fibers.

The Advanced Structural Concepts Team uses macroscale material modeling to investigate the performance of novel structural concepts such as the ITSE. A schematic diagram of the ITSE is shown in Figure 6. The ITSE consists of an outer fabric or membrane that contains (i.e., confines) a granular material, such as sand, and a central inner pliable bladder that is pressurized and places the surrounding granular material under a confining pressure, thus increasing shear strength and moduli of the granular material, and thus the entire structural system. Experimental tests show significant load-carrying capabilities from this simple construct using existing high-strength fabric and moist sand.

The ITSE exhibits higher strengths and stiffness with higher confining pressures. The Advanced Structural Concepts Team uses numerical analysis to predict the enhanced performance for the ITSE if the outer Vectran® fabric is replaced by a carbon nanotube superstrength membrane. The carbon nanotube membrane would support a higher internal pressurization and thus increase the load-carrying capabilities of the ITSE over that used in the current ITSE laboratory-testing (Figure 7) program. Data required to perform the numerical analysis is derived from the fiber simulations and used to generate the stress-strain relationship for the material model for the outer fabric of the ITSE.

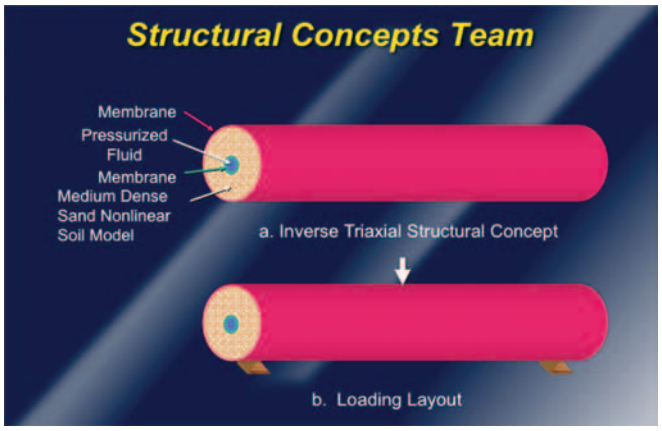

Figure 6. Schematic diagram of ITSE structure concept

In Figure 5, the data for the simulation labeled "T_16_250" are scaled to give a maximum tensile strength of 1 million psi (6.89 GPa) and is shown in the plot labeled "A." The resulting stress-strain curve in Figure 5-A lies well below the stress-strain curves for group B and group C in Figure 5. However, the data in Figure 5-A can be compared with the current stress-strain data used for the current high-strength material, Vectran®, in Figure 8. These data (Figure 5-A) are used as the input for the ITSE carbon-nanotube-based fabric simulation (Figure 8). The results of this simulation can then be compared against the experimentally validated simulation where the outer cylinder of the ITSE was made from a high-strength braided textile, Vectran®, to investigate the performance enhancements that can be achieved with the use of novel high-strength carbon-nanotube-based fabrics. These numerical modeling efforts using a 3-D nonlinear finite element method of analysis are currently underway.

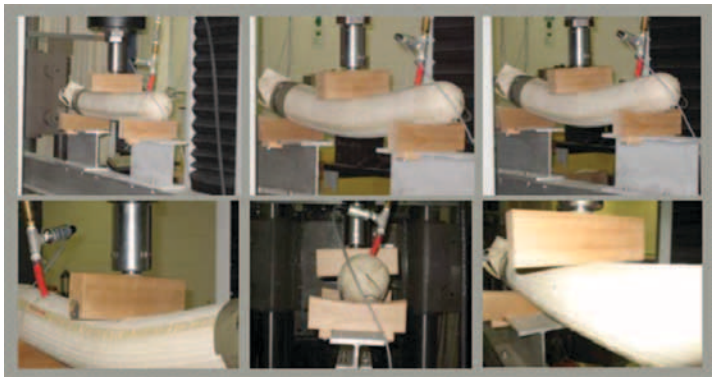

Figure 7. Experimental testing of the ITSE structure. Data from these tests are used to calibrate the ITSE models

In summary, advances in experimental, theoretical, and computational techniques, coupled with HPC resources, are changing the way the modeling and simulation community provides input to engineers and the way engineers design new materials and test advanced structural concepts. There is an advantage to integrating modeling and simulation with experimental programs. This new paradigm provides a holistic approach to material design with the potential for accelerating the development of new materials. It also provides a method for evaluating the performance of these materials and allows the potential benefit of these materials to be tested in the application of advanced structural concepts.

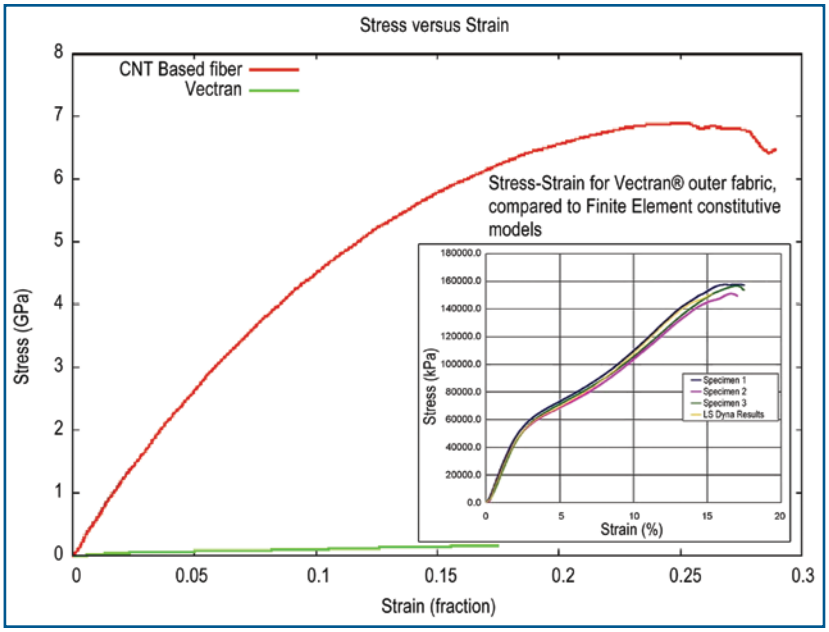

Figure 8. Stress-strain response of carbon-nanotubes-based fabric compared with Vectran. Inset shows experimental results for Vectran outer fabric, compared with finite element constitutive models

We gratefully acknowledge funding support from the U.S. Army ERDC research program "Molecular-Designed Carbon Nanotube Filaments, Membranes, and Coatings"; the HPC Challenge Project "Molecular Dynamics Simulations to Underpin the Design and Development of High-Performance Carbon-Nanotube-Based Filaments, Membranes, and Coatings"; and the allocation of computer time from the DoD HPCMP.

Carbon Nanotube Technology for Military Engineering Research Team

ERDC Research Team: Dr. Bob Welch (Program Manager Carbon Nanotube Technology for Military Engineering), Dr. Bob Ebeling (Team Leader-Structural Concepts), Dr. Charles Marsh (Team Leader-Material Synthesis), Dr. Charles Cornwell (Team Leader-Molecular Dynamics Simulations), Rick Haskins, Dr. N. Jabari Lee, Dr. Robert Kirgan, Dr. Anthony Bednar, Clint Arnett, Thomas Carlson, Dr. Kevin Abraham, Pete Stynoski, Jeffry Grey, Lauren Lilly, Micael Edwards, Brad Newcomb, Dr. Jeff Allen, Tim Garvin, Jon Georges, and Ben Masters

Natick Collaborators: Claudia Quigley, Karen Buehler

NASA Collaborators: Dr. Richard Jaffe (NASA Ames), Dr. Mike Meador (NASA Glenn)

Rice University Collaborators: Nobel Laureate Robert Curl, Prof. Matteo Pasquali, Prof. Robert Hauge

MIT Collaborators: Prof. Mike Strano, Dr. Jae-Hee Han (PhD post-doc w. Michael Strano), Dan Heller, Monica Usrey

DoD PETTT Collaborators: Carnegie Mellon University: Prof. Anthony Rollett; University of Florida: Prof. Susan Sinnott and Prof. Simon Phillpot; Clemson University: Prof. Steve Stuart; High Performance Technologies, Inc.: Dr. Jean Blaudeau, Dr. James Lill, James Larentzos

About the author: Dr. Charles F. Cornwell is a research scientist at ERDC working as a research physicist in the Computational Analysis Branch, Information Technology Laboratory. He holds a PhD in physics from Florida Atlantic University. Before coming to ERDC, Dr. Cornwell participated in the User Productivity, Enhancement, and Technology Transfer and Training (PETTT) program, which is part of the DoD HPCMP designed to enhance efficient use of the computing facilities provided by the HPCMP.

Potential Energy Surface Mapping of Energetic Materials Using Coupled Cluster Theory

By DeCarlos E. Taylor, Army Research Laboratory

DoD Significance

The experimental formulation and characterization of energetic materials (EM) and armor ceramics (AC) used for soldier protection is an expensive process requiring a large investment in researcher time and funding. Experiments often pose significant risks to researchers through unexpected detonation or accidental contamination; therefore, extreme (and expensive) measures must be taken to ensure researcher safety. As a result, computer simulations are becoming more prominent in research and development programs within the Department of Defense (DoD). The Army Research Laboratory (ARL), Naval Surface Warfare Center (Indian Head), and Air Force Research Laboratory scientists routinely use computer simulations to evaluate the properties and performance of materials using a combination of quantum mechanics (QM) and molecular dynamics (MD) simulation. Through computer modeling, the performance of materials spanning a range of environmental conditions can be studied at a comparatively cheap cost compared with experimental determinations. Computer modeling ultimately expedites the design, testing, and fielding of improved materials providing increased protection to the soldier.

Quantum Mechanics

At the atomic scale, the electronic degrees of freedom within a material can be characterized through solution of the time-independent Schrödinger equation $H\Psi = E\Psi$, where H is the system Hamiltonian and E is the energy eigenvalue corresponding to the wave function Ψ [1]. Solution of this equation yields a wide array of useful chemical information including the atomic structure, thermodynamic data, and spectroscopic signature of the system. Although guised in a rather simple representation, the Schrödinger equation is a notoriously difficult equation to solve, and a successful analytic solution has thus far only been achieved for one system, the hydrogen atom. In practice, the Schrödinger equation is normally solved by expanding the unknown wave function Ψ in terms of a finite set of known functions (basis set), and the set of expansion coefficients is determined algebraically [2]. The solution of the quantum mechanical equations is hopelessly impossible without a computer, and early practitioners of QM were justifiably pessimistic about the application of QM in chemistry [3]:

The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.

– P.A.M Dirac

However, with the advent of the computer and the incredible processor speed coupled with large amounts of available memory, QM has now become routinely applicable to a wide variety of molecules and materials. Further, with access to DoD Supercomputing Resource Centers (DSRCs) and their massively parallel compute clusters, DoD researchers can now do state-of-the-art QM calculations that were previously impossible.

One such state-of-the-art QM method that provides accurate approximations to the full solution of the Schrödinger equation is known as the coupled cluster (CC) theory [4]. The CC theory is generally regarded as the most accurate QM method available. It provides accurate results, often rivaling experiment, for a wide variety of molecular properties including structures, energies, and vibrational frequencies. Unlike the popular density functional theory [5], it is a systematically improvable series of approximations to the exact solution of the Schrödinger equation. Unfortunately, full inclusion of all the requisite terms is a veritable impossibility for all but the smallest molecules; therefore, approximate solutions are typically determined. However, within the current levels of approximation, the exact total energy can be recovered to less than 1 percent, but this accuracy is not without cost. The computational scaling of the CC method can depend on the 6^{th}, 8^{th}, and sometimes 10^{th} power of the system size depending on the chosen level of approximation [6]. Further, since solution of the CC equations requires an iterative approach, 10 to 15 evaluations of these expensive terms are often necessary, which further increases the computational expense. Due to this steep scaling, application of the method has been limited to small molecules with a small number of electrons. However, recent parallel implementations of the CC method in programs such as ACESIII [7] have extended the range of applicability of the CC method to a wide variety

of systems previously beyond the reach of quantum mechanics. Within ARL, the ACESIII program system is being used in research of energetic materials and to provide accurate reaction rate constants for toxic industrial chemicals using transition states determined at the CC level.

Application

A critical component of the energetic materials research effort at ARL is simulation of currently fielded and notional materials using MD. The accuracy of any MD simulation is directly contingent upon the quality of the potential used to "move the atoms." Ideally, QM should be used directly to compute interatomic forces "on the fly"; however, even with highly scalable quantum MD codes such as CP2K [8] and high performance computing (HPC) resources, the use of QM potentials directly in MD places limits on the maximum time scale/length scale that can be simulated. Due to this limitation, it is still common practice to fit a parametric function to QM data computed on a molecule whose structure/properties are representative of the target simulation material. There is a wide array of classical functionals available for MD ranging from simple Lennard-Jones, exponential-6, and Morse formulations to the complex ReaxFF [9] functional form that has seen a rise in popularity in recent years. Regardless of the complexity of the potential, determination of the parameters requires the highest quality reference data; and in the absence of experiment, this must be obtained quantum mechanically using HPC resources such as those available at the Arctic Region Supercomputing Center and other DSRCs.

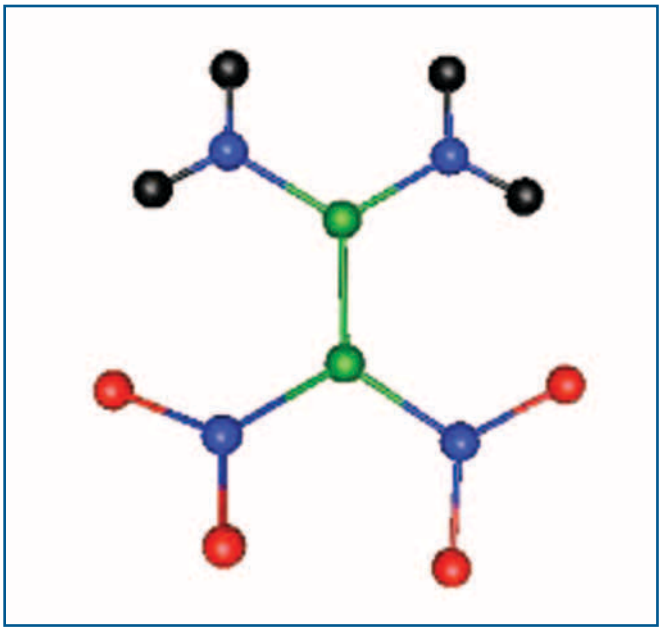

Figure 1. Fox-7 monomer. Black = Hydrogen, Green = Carbon, Red = Oxygen, Blue = Nitrogen

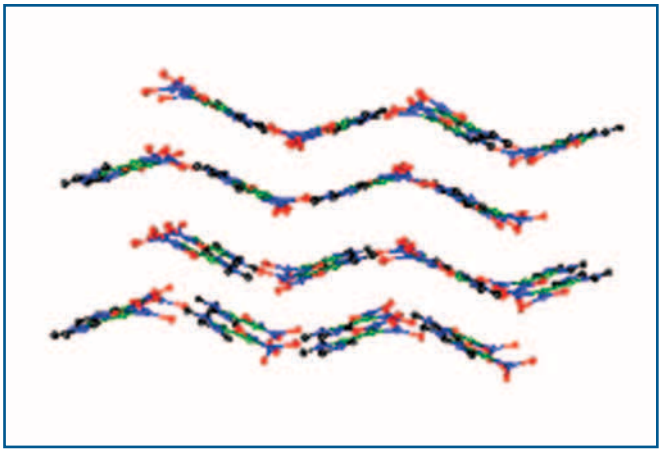

Figure 2. Wavelike layers of Fox-7 crystal

1,1-Diamino-2,2-dinitroethylene (Figure 1), known as Fox-7, has been identified as a promising energetic material based in part on its low shock sensitivity and thermal stability [10]. Fox-7 has a peculiar crystal structure consisting of wavelike layers with extensive intra- and intermolecular hydrogen bonding within the layers and weak van der Waals (vdw) forces between the molecular sheets (Figure 2). With respect to MD, in order to capture these intricate physical interactions that define the Fox-7 crystal structure, accurate QM reference data are required for potential parameterization. Using the ACESIII program system, we are computing a six-dimensional potential energy surface of Fox-7 dimers at the CC level. The Fox-7 dimer contains 28 atoms (20 heavy atoms), and within an aug-cc-pVDZ basis set, it contains 532 basis functions, a sizeable computation. In relation to the CC work, we have developed a preliminary Fox-7 potential function using the symmetry adapted perturbation theory [11] (SAPT(DFT)), which has a more affordable computational scaling than the CC method. Using SAPT(DFT), it was found that 1008 interaction energy calculations were required to obtain an accurate potential energy function; and under the Challenge Project program (with an allocation of 7 million CPU hours), we are repeating the set of dimer configurations used to determine the SAPT(DFT) potential energy function at the CC level to assess the quality of the CC method for development of potentials for use in MD simulations. We are nearing completion of the first year of a 2-year Challenge Project effort and are in preparation to do an initial parameterization to assess the quality of the CC potential and identify regions of the potential energy surface needing refinement. It is our hope that the CC potential will yield results that rival experiment, and future work will involve addition of a flexible intramonomer potential that will permit simulation of the α to β phase transition [12] that involves reorientation of

the Fox-7 molecular geometry. This is not part of our CC potential that currently assumes a rigid monomer orientation.

Summary

With the rapid increase in processor speed, computer memory, and scalable QM software packages, DoD researchers are uniquely positioned to apply high-level QM methods to molecules that were beyond the limit of computational resources in years past. Although the CC method is now a viable option for treatment of large energetics, it is still a substantial computational effort since a large number of processors are generally required for an acceptable wall clock time per computation. Without DSRC resources such as the availability of a large number of parallel processors for extended run times, work such as this would not be possible. It is our hope that with advances to hardware and more scalable QM software, high-level QM methods such as the CC theory will find increasing application in research of molecules and materials whose sizes were previously beyond the reach of these methods.

References

[1] Levine, I., *Quantum Chemistry 4th Ed,* Prentice Hall, 1991.

[2] Szabo, A. and Ostlund, N.S., *Modern Quantum Chemistry*, McGraw-Hill, 1989.

[3] Dirac, P.A.M., Proc. Roy. Soc., **123**, 714 (1929).

[4] Bartlett, R.J., in *Modern Electronic Structure Theory – Part II*, World Scientific, 1995.

[5] Hohenberg, P. and Kohn, W., Phys. Rev. B, **140,** 864 (1964).

[6] Bartlett, R.J, and Stanton, J.F., in *Reviews in Computational Chemistry, Vol. 5*, VCH Publishers, 1994.

[7] See *http://www.qtp.ufl.edu/ACES/index.shtml*

[8] See *http://cp2k.berlios.de/*

[9] Goddard III, W.A, Zhang, Q., Uludogan, M., Strachan, A., and T. Cagin, in *Fundamental Physics of Ferroelectrics - AIP Conference Proceedings No. 535*, 2000.

[10] Latypov, N.V., Bergman, J., Langlet, A., Wellmar, U., Bemm, U., Tetrahedron, **54**, 11525 (1998).

[11] Podeszwa, R., Bukowski, R., Szalewicz, K. J., Chem. Theory Comput.**2**, 400 (2006).

[12] Evers, J., Klapotke, T., Mayer, P., Oehlinger, G., Welch, J., Inorg. Chem., **45**, 4996 (2006).

Atomic "Triple-Dog" Dare: Using HPC Resources to Help Create Dipolar Matter

By Dr. James Freericks, Department of Physics, Georgetown University, Washington, DC

In 1983, the movie *A Christmas Story* [1] was released and has since become a holiday classic. One scene in the movie involves the main character, Ralphie, witnessing two of his friends debating over an old urban legend. Schwartz says his tongue will stick to an ice cold metal pole, while Flick insists it will not. Once Flick is "triple-dog" dared, there is no turning back. The urban legend is tried, and Flick's tongue is stuck to the pole, requiring a call to the fire department to release him (see Figure 1). We might ask, "How does this relate to science, high performance computing, and atomic physics?" The Department of Defense (DoD) is sponsoring work to help create quantum degenerate matter. This requires making molecules from two different elements so that the electrons of the molecule like to preferentially lie near one of the elements, creating a net positive charge on one side and a net negative charge on the other side of the molecule. This electrical charge arrangement is called a dipole, and the relative orientation of the positive and negative sides of the dipoles can be oriented by placing them in an external electric field, which makes the dipoles line up (similar to what happens to the hairs on our arms when we bring a balloon close by that has been charged

Figure 1. *Flick sticks his tongue onto a cold metal pole as Ralphie watches in* A Christmas Story *(1983) [1]*

by rubbing it on a sweater). If we further lower the temperature enough, then we can reach the so-called quantum-degenerate limit, where interesting effects can occur in the material governed by the wavelike properties of the ultracold dipoles.

So what does this have to do with the "dog dare" from *A Christmas Story*? When we want to reach the quantum limit at low temperatures, the most direct way to

make the dipole is to first cool the mixture of atomic elements, and then perform an ultracold chemical reaction that binds the two different atoms together and removes all of the exothermic energy via a coherent photon (this miraculous technique is similar to exploding a bomb next to a butterfly and having the butterfly emerge undisturbed). Such a process is called stimulated Raman adiabatic passage, which is used in these experiments after we sweep a magnetic field that weakly binds the two atoms together into something called a halo molecule. It is the formation of the halo molecule that is the "Achilles' heel" of this experiment, as we need the two atoms to be close enough to each other for them to find each other and bind together during the magnetic field sweep [2]. Here is where the "dog-dare" enters. By making the system cold enough and placing it in an optical lattice (which is similar to an egg carton, with the depressions made by interfering laser beams), we can pre-form the molecule by having exactly one atom of each species lying in each depression, then the halo molecule formation will be essentially perfect. So we need to make each atom "stick" to the other one and sit in one of the depressions of the optical lattice. This is where we connect to Flick sticking his tongue on the metal pole, and where the concept of making the atoms cold enough to stick together is so critical. In the experiments [2], mixtures of potassium (K) and rubidium (Rb) are used.

High performance computing is employed to simulate how efficient this process can be at a given temperature. Put in other words, we are trying to determine how low the experimentalists need to go in temperature in order to be able to improve their yield for making dipolar molecules (the current yield is around 20 percent of all available pairs in the initial atomic mixture). Our simulations show that this rate can increase to almost 100 percent if the system is properly prepared. We use a sophisticated algorithm called inhomogeneous dynamical mean-field theory to carry out the calculations. This approach, which becomes exact in the limit where the spatial dimensions go to infinity, has been benchmarked against exact Monte Carlo calculations in two dimensions and shown to agree to within a few percent (see Figure 2, where the efficiency for molecule formation is plotted for the two approaches). The algorithm requires the inversion of a large number of general complex matrices (of a high dimension) at each step of the iterative process. This is where the code easily parallelizes in the master-slave format by allowing each matrix to be sent to a slave node for the inversion, which is done with BLAS and LAPACK routines and hence runs efficiently (recently, we have begun to generalize the approach to a Lanczos-based procedure

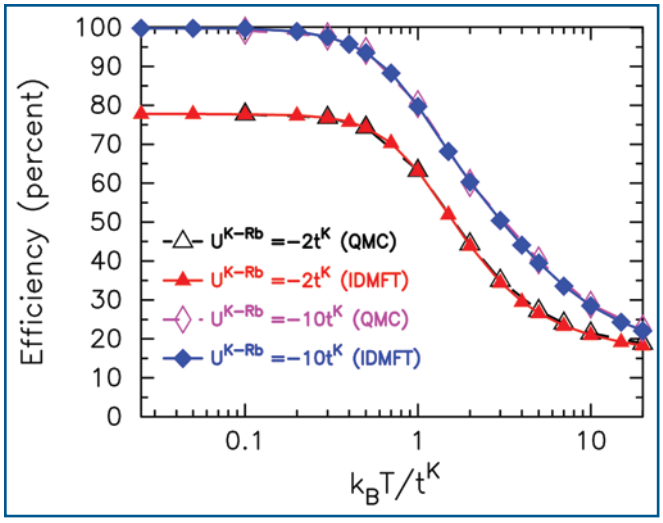

Figure 2. Comparison of the inhomogeneous dynamical mean field theory (IDMFT) solution to the exact quantum Monte Carlo (QMC) solution for the efficiency of forming molecules versus temperature. Note how accurate the approximate IDMFT approach is

since the matrices are sparse, with a factor of 2 reduction in the computation time and no loss of precision).

In the experiments, it is difficult to measure temperature directly. These experiments run at some of the coldest temperatures in the universe, just a few billionths of a degree above absolute zero. At such a cold temperature, there are no universal thermometers to use. But because there is a direct thermodynamic relation between the entropy of a system and the temperature, we can use the entropy of the system as the thermometer. The entropy, which is in some sense a measure of how disordered the mixture of atoms is, can be measured in the mixture before the optical lattice is turned on. Next, the optical lattice is turned on slow enough that the entropy remains constant. We then report our results for the efficiency as a function of the entropy per particle. Calculating the entropy of this mixture is only possible with our inhomogeneous dynamical mean field theory approach. Some sample results are given in Figure 3.

The high efficiency at low temperature/entropy results from the two species of atoms binding together and sitting on top of each other. Our results indicate that we can achieve high efficiency for forming dipolar matter if the system is cooled to a low enough temperature that we pre-form molecules in the optical lattice prior to performing the ultracold chemistry to form the dipolar molecule. We are currently working with an experimental team to implement these ideas within the context of the DoD efforts in this area. High performance computing was an invaluable tool to allow us to carry out this work. Our research has already been published in Refs. [3, 4].

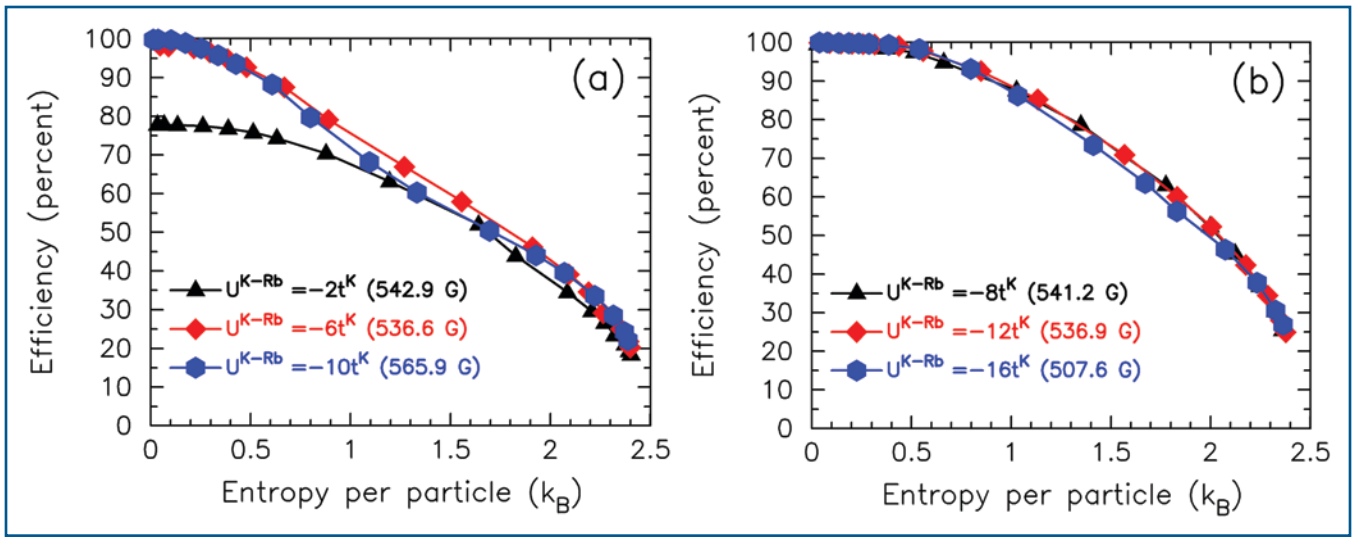

Figure 3. Efficiency for molecule formation versus entropy per particle for two cases: (a) a moderately deep optical lattice and (b) a deeper optical lattice. Note how the efficiency becomes large and assumes a nearly universal curve once the attraction between the atoms is large enough

The DoD is interested in this work for two reasons. On the one hand, they want to use the precise control of atomic systems in optical lattices to build a real materials simulator, where designer materials could be developed first in the simulator, shown to have the desired properties, and be able to be grown in a lab. On the other hand, it is believed that quantum degenerate dipolar matter will have unusual quantum properties that might allow it to be used as ultrasensitive sensors of things like electric fields.

References

[1] *A Christmas Story*, screenplay by Jean Shepherd, directed by Bob Clark, Christmas Tree Films, 1983.

[2] Ni, K.-K., S. Ospelkaus, M. H. G. de Miranda, A. Pe'er, B. Neyenhuis, J. J. Zirbel, S. Kotochigova, P. S. Julienne, D. S. Jin, and J. Ye, "A high phase-space density gas of polar molecules." *Science,* 322, pp. 231-235, 2008.

[3] Freericks, J. K., "Optimizing pre-formed molecules in mixtures of ultracold ^{40}K and ^{87}Rb on an optical lattice: A challenge grant and capabilities application project." *2009 DoD High Performance Computing Modernization Program Users Group Conference, San Diego, CA,* edited by D. E. Post, pp. 205-212, Los Alamitos, IEEE, 2009.

[4] Freericks, J. K., M. M. Maska, Anzi Hu, T. M. Hanna, C. J. Williams, P. S. Julienne, and R. Lemanski, "Improving the efficiency of ultracold dipolar molecule formation by first loading onto an optical lattice*." Phys. Rev. A*, 81, pp. 011605-1-4(R), 2010.

Dr. James Freericks received his bachelor's degree in physics from Princeton University in 1985 and his Ph.D. in physics from the University of California, Berkeley in 1991. After spending time as a postdoctoral fellow at UC's Santa Barbara and Davis campuses, he went to Georgetown University, where he has been a professor of physics since 1994. He has been involved in high performance computing on DoD machines since 2001 and has worked on numerous CAP projects and two Challenge grants. He is a Fellow of the American Physical Society.

Parametric Studies for Superconducting Quantum Interference Filters

By Dr. Juan Carlos Chaves, Army Research Laboratory, User Productivity, Enhancement, Technology Transfer and Training (PETTT) Program

Problem

Superconducting quantum interference devices (SQUIDs) are the world's most sensitive detectors of magnetic signals (sensitivity ~ femto-Teslas) for the detection and characterization of signals so small as to be virtually immeasurable by any other known sensor technology. There is mounting evidence that superconducting quantum interference filters (SQIFs), arrays of SQUIDs, having intentionally dissimilar loop areas can act as wideband (1MHz - 1GHz) transimpedance amplifiers having large (>100 dB) linear dynamic range and can provide the basis of centimeter-scale antennas for these frequencies. Recently, different types of SQIFs built by one-dimensional (1-D) or 2-D arrays of SQUIDs have been studied numerically and experimentally. SPAWAR Systems Center Pacific is developing fully phenomenological models to simulate the 1-D parallel (and 1-D serial) multiloop array architecture (with unconventional grating configurations). These models are used to perform complex parametric studies for the full characterization of SQIFs under a variety of conditions. The serial SQIF algorithms, developed in MATLAB, are too computationally intensive to allow the modeling of large enough systems or to achieve enough throughput for significant parametric studies in conventional PC/workstation architectures. Through the use of algorithmic improvements and parallel MATLAB technologies, the User Productivity, Enhancement, Technology Transfer and Training (PETTT) program is providing optimized codes that empower SPAWAR users to run much larger systems and to perform parametric studies by exploiting the DoD High Performance Computing Modernization Program (HPCMP) HPC platforms.

Methodology and Results

The serial SQIF codes are profiled to identify optimization opportunities. For example, as the SQIF codes are written in MATLAB, vectorization offers many opportunities for performance enhancement. Vectorized code takes advantage, wherever possible, of operations involving data stored as vectors. This extends to matrices and general arrays since a MATLAB matrix is stored (by columns) in contiguous locations in the computer's RAM. The speed of a numerical algorithm in MATLAB is sensitive to whether or not vectorized operations are used. Additional effort is focused in parallelization opportunities. The SQIF codes offer task and data-based parallelization opportunities that are exploited through embarrassingly and fine-grain parallelization techniques, respectively. The optimized codes are extensively tested and validated before using them for parametric studies. PETTT contributions are facilitating the completion of several parametric studies for the full characterization of large SQIF systems under a variety of conditions. These studies would not be possible without the availability of optimized/parallelized SQIF codes and the exploitation of HPCMP HPC platforms provided by PETTT.

DoD Impact

The ability to run with larger SQIF systems and investigate larger parameter space empowers the SPAWAR team to explore novel SQIF

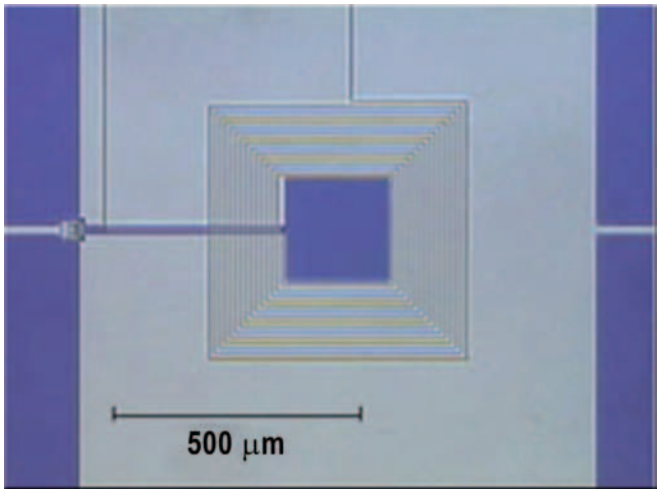

SQUID: a tiny loop of superconducting material interrupted by narrow gaps/Josephson junctions

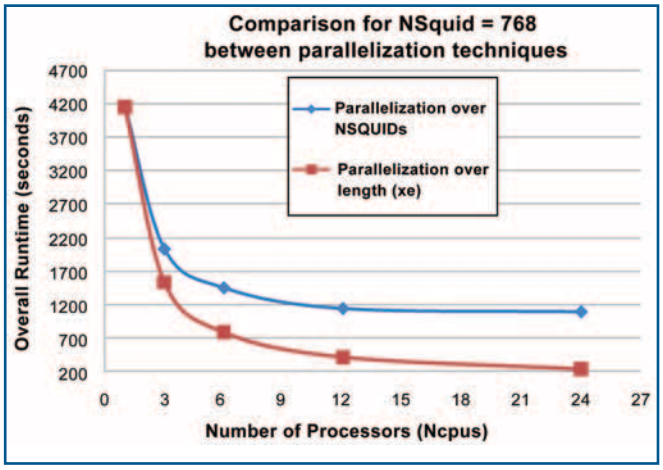

Comparison between two techniques of parallelization for SQIF codes

aspects, otherwise hidden, with a dramatic reduction in time-to-solution and significant enhancement in productivity. This directly impacts the understanding and development of extremely small superconducting sensors with large dynamic range and high sensitivity. SQIFs offer an opportunity for extremely small and extremely wideband antennas that would allow for consolidation and size reduction of detection systems while achieving environmentally noise-limited performance. In addition to the logistic benefits, this type of antenna performance will allow increased range against conventional radio frequency signals and will allow detecting a whole new class of signals previously undetectable by any technology. This is of great importance for a variety of DoD and homeland security applications.

State-of-the-Art Computational Fluid Dynamics Flow Solver Aids X-51A Research

By Dr. Susan Cox-Stouffer, Aerospace Propulsion Division, Air Force Research Laboratory

A scramjet is a supersonic combustion ramjet. It compresses, combusts, and exhausts air without moving parts. The shock that forms at the leading edge of a supersonic vehicle compresses the air, and the air moves through the engine to the combustor, where fuel must inject, mix, and burn in milliseconds. Facilitating that process is a major challenge in scramjet research, as are thermal management and structural integrity. Despite the challenges, various Department of Defense (DoD) agencies, and others, have long been interested in scramjet-powered vehicles. Current Air Force interest focuses on the possibility of a scramjet-powered missile for hardened and time-critical targets, though scramjets could also play a role in affordable, reusable space-access vehicles and perhaps even in human hypersonic flight.

The X-51A Scramjet Engine Demonstrator-Waverider program is heir to a long line of scramjet development programs supported by various National Aeronautics and Space Administration (NASA) and DoD interests, including the Air Force, Navy, and Defense Advanced Research Projects Administration (DARPA). More specifically, it is the result of a collaboration between the Air Force and DARPA, which aims to test a HyTech scramjet engine flowpath in the Affordable Rapid Response Missile Demonstrator (ARRMD) waverider

Artist concept rendering of X-51A

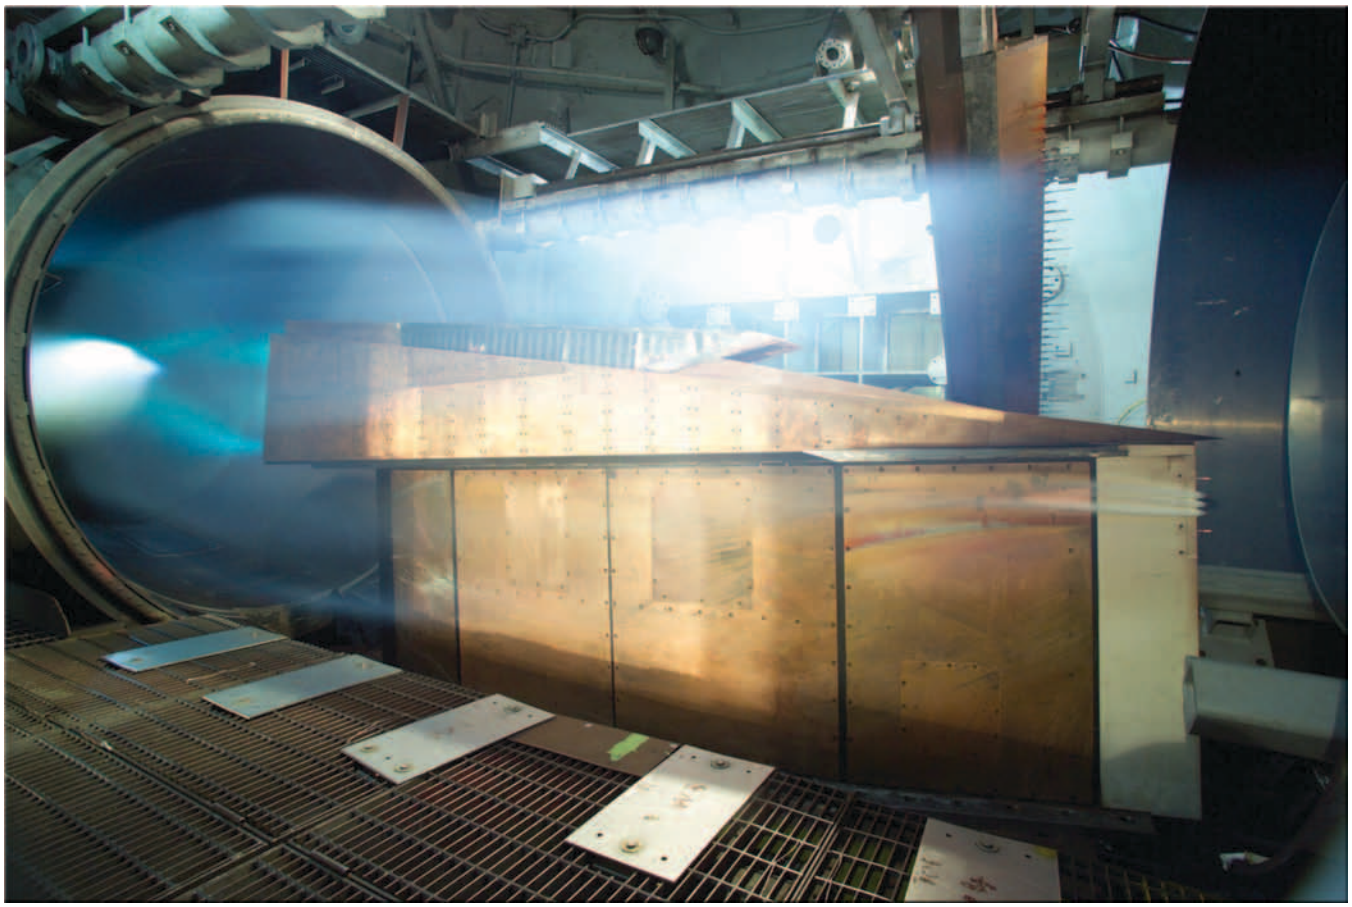

Engine test – 2007-02-21 – SJX61-1 transition to JP7

airframe design. The X-51A is a technology demonstration program, not a weapon development program, though it could be adapted for use as a hypersonic cruise missile. The main program objective is the flight test of the HyTech engine, using endothermic hydrocarbon fuel, including acceleration from boost, approximately Mach 4.5 to Mach 6.5.

The X-51A, like most aircraft, was subjected to laboratory tests long before its first flight. One of the challenges of any aircraft engine development program is the realistic simulation of flight conditions in ground-test facilities. For scramjets the challenge is particularly daunting, because the test fluid (air) must be heated to extreme conditions to maintain proper temperature at supersonic speeds. Most heating systems contaminate the test fluid, which complicates testing. X-51A tests were performed in the NASA-Langley Research Center 8-Ft. High Temperature Tunnel (8-Ft. HTT). Air in that test facility is heated by in-stream combustion of methane. Though oxygen is added to compensate for that lost to combustion, the resultant mixture of air and combustion products contains relatively high concentrations of water vapor. Since the test fluid cools as it accelerates, the water vapor can undergo a phase change during testing, vastly complicating analysis.

Most computational fluid dynamics (CFD) analysis of scramjet systems is performed using flow solvers valid only for gaseous flows. A collaboration among the Air Force Research Laboratory's Aerospace Propulsion Division (AFRL/RZA), Taitech, Inc., and researchers at North Carolina State University (NCSU) has resulted in the development of REACTMB-MP, a state-of-the-art CFD flow solver capable of simulating the formation and evolution of droplets of condensed fluid, such as water, in an otherwise gaseous flow at high-speed conditions. The flow solver is also capable of simulating supercritical and purely liquid flows and contains advanced large-eddy simulation capabilities as well. Working with the lead code developer, Dr. Jack Edwards of NCSU, Dr. Susan Cox-Stouffer, a Taitech employee long associated with the AFRL Aerospace Propulsion Division and an experienced DoD Supercomputing Resource Center (DSRC) customer (more than 14 years for both), used REACTMB-MP to predict the condensation and consequences of water droplets during X-51A ground tests. Her work included simulations of the 8-Ft. High Temperature Tunnel itself and of the various aerodynamic probes used to gather information about the test fluid. Since both the test facility and the aerodynamic probes are axisymmetric, the simulations

could be performed in two dimensions with source terms for axisymmetry, allowing quick and efficient computation and an extensive test matrix. Typical simulations were performed on 10 processors on AFRL and U.S. Army Engineer Research and Development Center (ERDC) computers (particularly, *Eagle, Hawk, Sapphire,* and *Jade*) and required approximately 2 days to reach convergence. Three-dimensional simulations of the X-51A vehicle components have also been performed on those and other DSRC computers and require vastly more resources. These simulations typically ran on 50 to 300 processors and were completed in 3 to 14 days.

The results of the test-facility and probe simulations proved particularly useful during X-51A ground testing. As Figure 1 indicates, the presence of condensation has a marked impact on test conditions. Mach number, which is a simple function of geometry for ideal gases, is in fact a function of the degree of condensation and can differ by as much as 10 percent from the expected value. Other variables are even more strongly impacted. In particular, static pressure measurements can differ by as much as 40 percent from expected values. The availability of two-phase simulations of the test facility and static-pressure probes provided the test team with advanced warning of otherwise inexplicable results, expediting tunnel calibration and saving program resources.

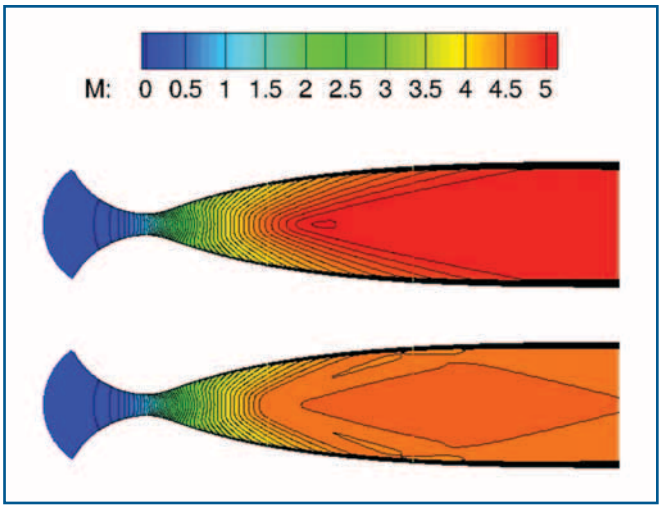

Figure 1. CFD image of Mach number. Purely gaseous simulation is on top and two-phase simulation on bottom, with lower Mach numbers and greater nonuniformity

Dr. Susan Cox-Stouffer is employed by Taitech, Inc. as an onsite contractor to the Air Force Research Laboratory's Aerospace Propulsion Division at Wright Patterson Air Force Base in Dayton, Ohio. She performs computational fluid dynamics analysis of scramjet engines and test facilities, and has provided support to the GDE-1, GDE-2, and X-51A programs. Her areas of expertise include fuel injection, mixing, and combustion and multiphase flow such as that caused by supercritical fuel. She received her bachelors and doctorate degrees from Virginia Tech (VPI&SU). Dr. Cox-Stouffer's e-mail address is Susan.Cox-Stouffer@wpafb.af.mil.

Geo-Temporal Visualization of Information Collected from Large Databases at MHPCC

By Kevin P. Roe and Maria Murphy, Senior Application Engineers, Maui High Performance Computing Center DoD Supercomputing Resource Center

The Time-Based COCOM Operation Picture (TIMECOP) system was developed at the Maui High Performance Computing Center (MHPCC) DoD Supercomputing Resource Center (DSRC) to provide the ability to visualize events over geographic locations as a function of time for large databases. It allows the user to discover geo-temporal trends, patterns, and behavior over large regions, specific sites, and user-defined time periods. This data mining capability has the capacity to make trend analysis possible to users who would have found it otherwise impossible to manually extract the necessary information from a large database and through the merging of multiple databases. Although it has been initially tuned to operate on a single large database, it is actually data agnostic and has since been fit to serve the needs of other users operating on data from different sources by merging the relevant information from multiple databases. TIMECOP can also transmit data to remote users in the kilobyte range allowing for mobile and limited bandwidth system access.

The key elements from the TIMECOP Web server can be seen in Figure 1. The pull down menus on the Web

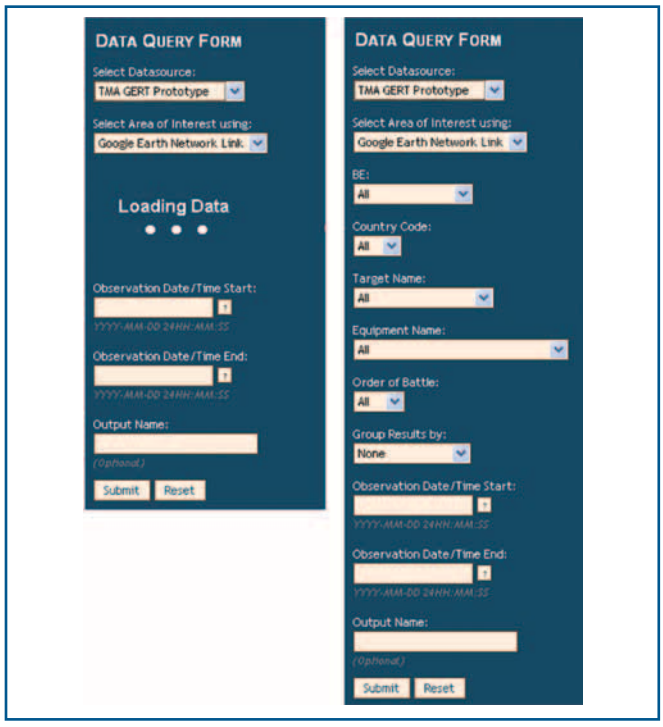

Figure 1. Key TIMECOP Elements

TIMECOP

EQUIPMENTNAME	COUNT	TIMESTAMP	LATITUDE	LONGITUDE	BE	COUNTRYCODE	TARGETNAME	TYPEORDEROFBATTLE
BELIZE29	3	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
BELIZE29	3	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
BELIZE29	6	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
BELIZE29	4	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
BELIZE29	2	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
BELIZE29	6	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
BELIZE29	3	2008-08-25T00:00:00Z	61.2125	166.0207	C92	LV	Dothan	TUR
FALKLAND ISLANDS (MALVINAS)2	5	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
FALKLAND ISLANDS (MALVINAS)2	3	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
FALKLAND ISLANDS (MALVINAS)2	3	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
FALKLAND ISLANDS (MALVINAS)2	5	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
FALKLAND ISLANDS (MALVINAS)2	5	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
FALKLAND ISLANDS (MALVINAS)2	3	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
FALKLAND ISLANDS (MALVINAS)2	3	2008-09-21T00:00:00Z	80.0191	-101.7861	B5466	VN	West Chicago	ANT
HAITI97	5	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
HAITI97	5	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
HAITI97	2	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
HAITI97	5	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
HAITI97	4	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
HAITI97	2	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
HAITI97	4	2008-06-08T00:00:00Z	77.9896	-13.3334	V81	CX	Jasper	ANT
LESOTHO65	4	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
LESOTHO65	6	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
LESOTHO65	4	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
LESOTHO65	3	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
LESOTHO65	4	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
LESOTHO65	5	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
LESOTHO65	1	2008-09-23T00:00:00Z	59.1481	-129.9505	K17	FI	Bucks	TUR
W	1	2008-06-03T00:00:00Z	81.5047	-37.6719	V66	ST	Peoria	MAL

Figure 2. TIMECOP Tabular Output

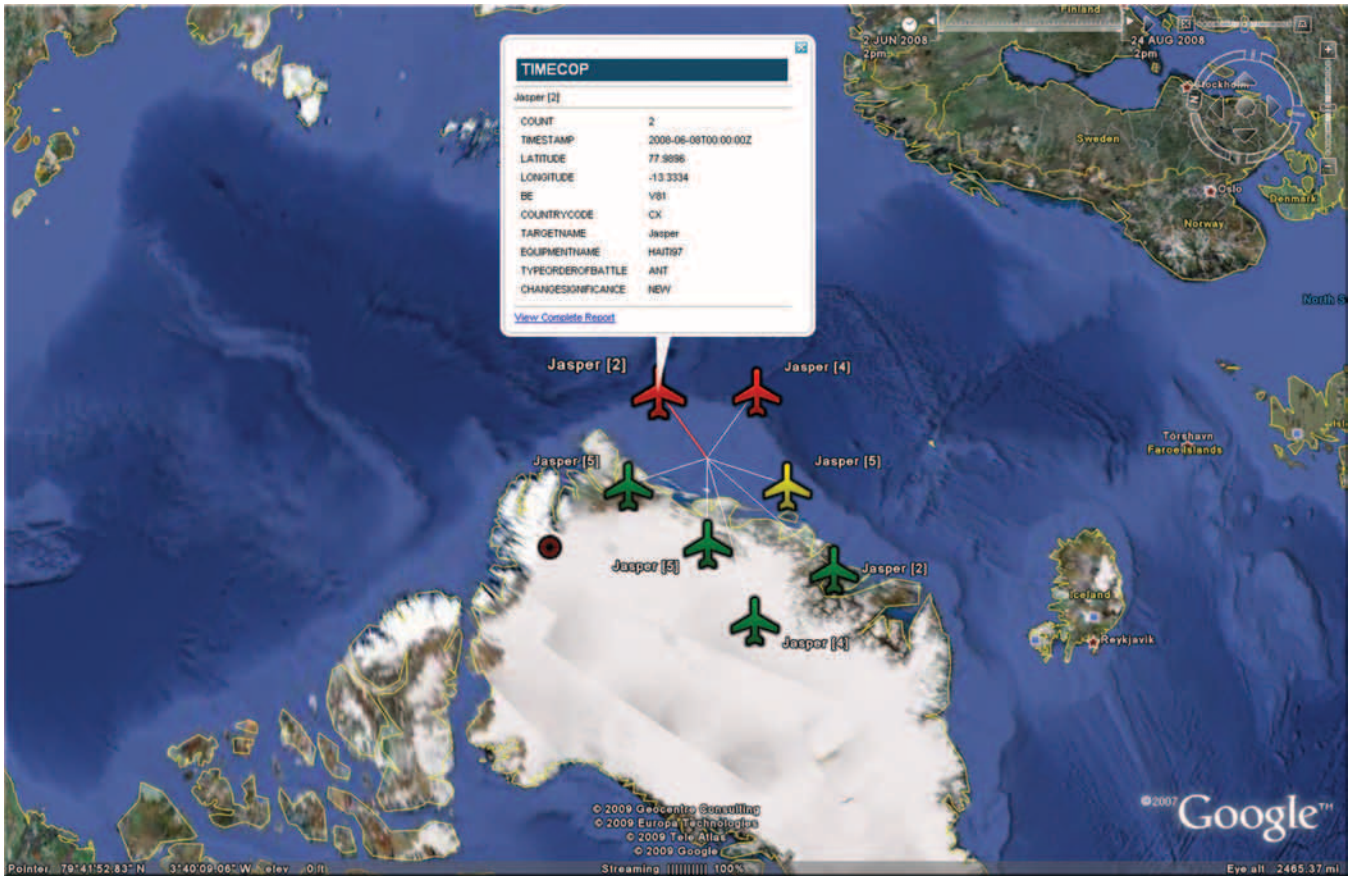

Figure 3. TIMECOP KML Output

page are dynamic and load possible query options based on the choices made in prior menus.

The TIMECOP server can output the results of the query in tabular format as well as in KML output that can be viewed with Google Earth. This is illustrated in the Figures 2 and 3.

The TIMECOP system has been implemented to an operational level over the course of a 2-year MERIT cycle. It is currently being utilized by both the National Geospatial-Intelligence Agency and Special Operations Command for their operational needs. Additional agencies are currently examining this system for application to their needs. Future design improvements include the capability to handle additional types of databases and data sources. Humanitarian applications have been found for the TIMECOP system; specifically, the Pacific Disaster Center has found it to be useful in aiding with HADR (Human Aid and Disaster Relief) efforts.

Kevin P. Roe graduated from the State University of New York, *magna cum laude,* with a B.S. degree in aerospace engineering in 1994. He obtained a master's degree from Syracuse University in aerospace engineering in 1996 and earned another master's degree from the College of William and Mary in computer science in 1998. For the past 10 years, Roe has been a senior application engineer at the MHPCC DSRC. He has been widely published in scientific journals.

Investigation of Three-Dimensional Internal and External Flow Separation

By Drs. Robert D. Jacobi, Wolfgang Balzer, Andreas Gross, and Herman F. Fasel, Aerospace and Mechanical Engineering, University of Arizona, Tucson, Arizona

Research and development of advanced aerospace and naval systems increasingly rely on computational fluid dynamics (CFD). The sophistication and reliability of CFD have grown continuously because of significant progress in the development of numerical methods and advanced turbulence models but also as a result of breathtaking progress in supercomputer performance and parallelization. At the same time, the demands posed on CFD by the applications keep ever increasing as more and more complex flow problems are being addressed. One example is the prediction of laminar and turbulent flow separation and reattachment. For most technical applications, separation has to be avoided, as it negatively affects performance by causing a dramatic drag increase and a lift reduction. Other undesirable consequences of separation are reduced maneuverability and stability and unsteady fluid dynamic loads resulting from the inherent instability of the separated flow region. For example, when considering submarines or similarly shaped underwater vehicles, high performance maneuvers may lead to flow separation from the control surfaces and other parts of the vehicle. While the majority of aeronautical applications has lifting surfaces with a rather large aspect ratio, aspect ratios seen in typical naval applications are generally moderate. As a result, separation for naval applications can be considered to be rather three-dimensional (3-D).

Separation alone is one of the least well-understood areas in fluid dynamics. Separation is typically associated with considerable unsteadiness since hydrodynamic instabilities of the mean flow give rise to coherent structures that originate in the separated boundary layer. These structures influence separation and reattachment, thereby affecting the aerodynamic characteristics such as lift and drag. While a relatively large body of research exists for 2-D separation, much less is known for 3-D separation, which is more intricate and far less understood. A better understanding of the physical mechanisms governing 3-D flow separation in general and the dynamics of the concomitant flow structures in particular, however, is essential for improving the performance of existing applications and for the development of revolutionary new concepts. In addition, better insight into the flow physics may lead to novel passive and active strategies for controlling or preventing separation.

We are investigating 3-D external and internal flow separation for prototypical low-aspect ratio geometries that are of relevance for naval applications. For the experimental part of our research, we are placing an axisymmetric displacement body with boundary layer suction over a flat plate. The pressure gradient induced by the displacement body results in the generation of a 3-D separation bubble on the flat plate (Figure 1). Because the bubble is generated on a flat plate, complicating effects of flow curvature are successfully eliminated. Without going into more detail of the water tunnel experiments, we now turn our attention to the central part of our research, which is based on direct numerical simulations (DNS) of the same geometry so that the experimental measurements can be utilized for validation and verification of codes and simulation data. For these simulations, we are employing a highly accurate and efficient in-house developed incompressible

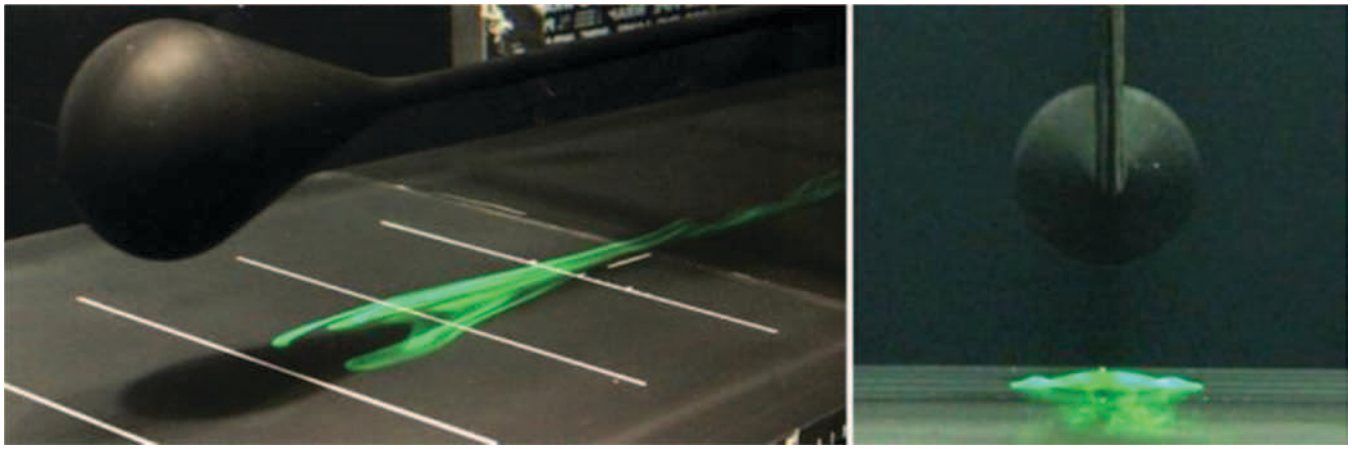

Figure 1. Perspective view (left: flow direction from left to right) and view from the back (right: flow direction towards the observer) for Re = 7500. Separation bubble was visualized using fluorescent dye

finite-difference CFD code. The acronym DNS refers to simulations that capture all scales of fluid motion down to the smallest dissipative scales.

We are investigating steady and unsteady 3-D separation bubbles on a flat plate for the same Reynolds number range, 5000-30,000, as in the water tunnel experiments. The Reynolds number is based on the approach flow velocity and the displacement body diameter. The computational domain was placed entirely under the displacement body (Figure 2). The number of grid points for our simulation ranged from 25 million to 0.7 billion. Simulations of that scope would not be possible without the support of the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP). Typical simulations were run on up to 1281 processors and required run times of up to 4 weeks. In addition to the number of available processors, memory, disk space, and turnaround times for the success of this research, data handling and processing capabilities are critical as well. Fortunately, whenever needed, the Consolidated Customer Assistance Center (CCAC) provided prompt assistance. Another asset that is crucial for simulations of such scale are in-house developed research codes tailored to the specific application and computer platform. We are currently working with Dr. Gabriele Jost from the User Productivity, Enhancement, Technology Transfer and Training (PETTT) program to improve the performance of our incompressible code for the Cray XT5.

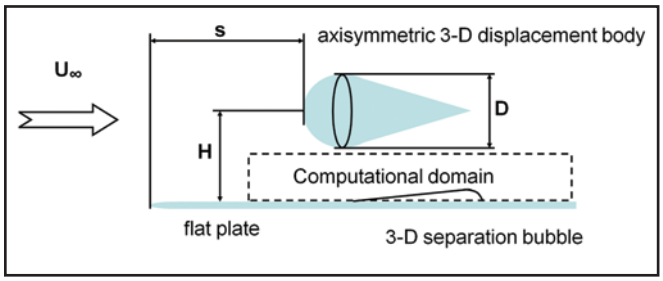

Figure 2. Setup of direct numerical simulations

The elliptic nature of the incompressible Navier-Stokes equations makes parallel solution strategies difficult. Classical domain decomposition approaches as typically used in compressible codes are not directly applicable. To make matters worse, the equations are highly nonlinear requiring massive processor communication and therefore a fast processor interconnect. The spectral discretization that we employ in the spanwise direction allows for a message-passing interface (MPI) parallelization. However, the number of processors is limited by the number of spectral modes. We therefore implemented a second layer of parallel computing by introducing shared-memory parallel programming (OpenMP) directives. This approach, which combines MPI and OpenMP, is commonly referred to as "hybrid" approach. Different modes of hybrid operation for the Cray XT5 are shown in Figure 3. Each Cray XT5 node consists of two quad-core central processing units (CPUs) that share the same memory. The maximum number of threads for a thread-based parallelization is therefore eight. However, since an MPI task can be assigned to the master thread, parallel work can be performed across several nodes.

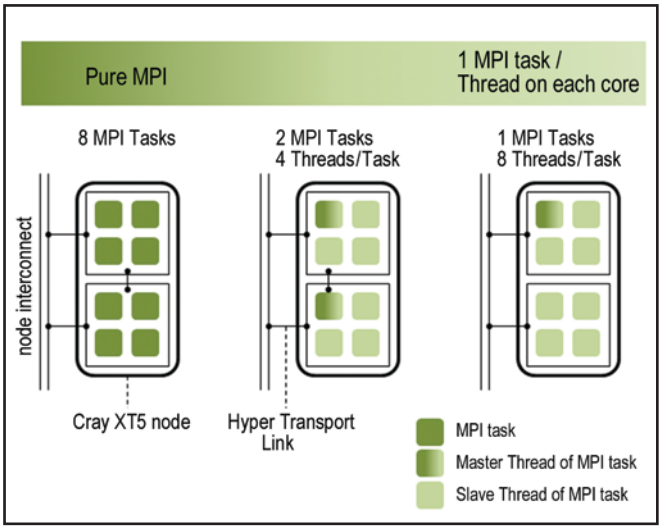

Figure 3. Modes of Hybrid operation for Cray XT5

We tested this "hybrid" code version for a problem size of 1729×256 grid points and 129 spectral modes. Compared with an ideal speedup of one for 129 processors, the hybrid approach results in a speedup of 1.91 for 258 processors (2 threads/MPI task) and 2.59 for 516 processors (4 threads/MPI task). We are confident that with further improvements of the code, it will become feasible to run simulations with four times the number of processors (4 threads) compared with the pure MPI version of the solver.

Separation bubbles for low-aspect ratio geometries have a complex flow topology that is a subject of ongoing research. Our simulations have pointed out many of the characteristic salient features of the flow field (Figure 4). The bubble topology was found to depend on the pressure gradient and the Reynolds number. Unsteady vortex shedding, which is of concern to the designer because it causes unsteady loads and noise, was observed for all but the smallest Reynolds numbers.

An instantaneous flow visualization for Re = 15,000 is shown in Figure 5. Here the Q vortex identification criterion was employed for visualizing vortical structures. Large-scale spanwise coherent structures with strong streamwise periodicity can be seen in the wake downstream of the bubble. The fact that these structures are symmetric in the spanwise direction supports our hypothesis that the dominant instability mechanisms

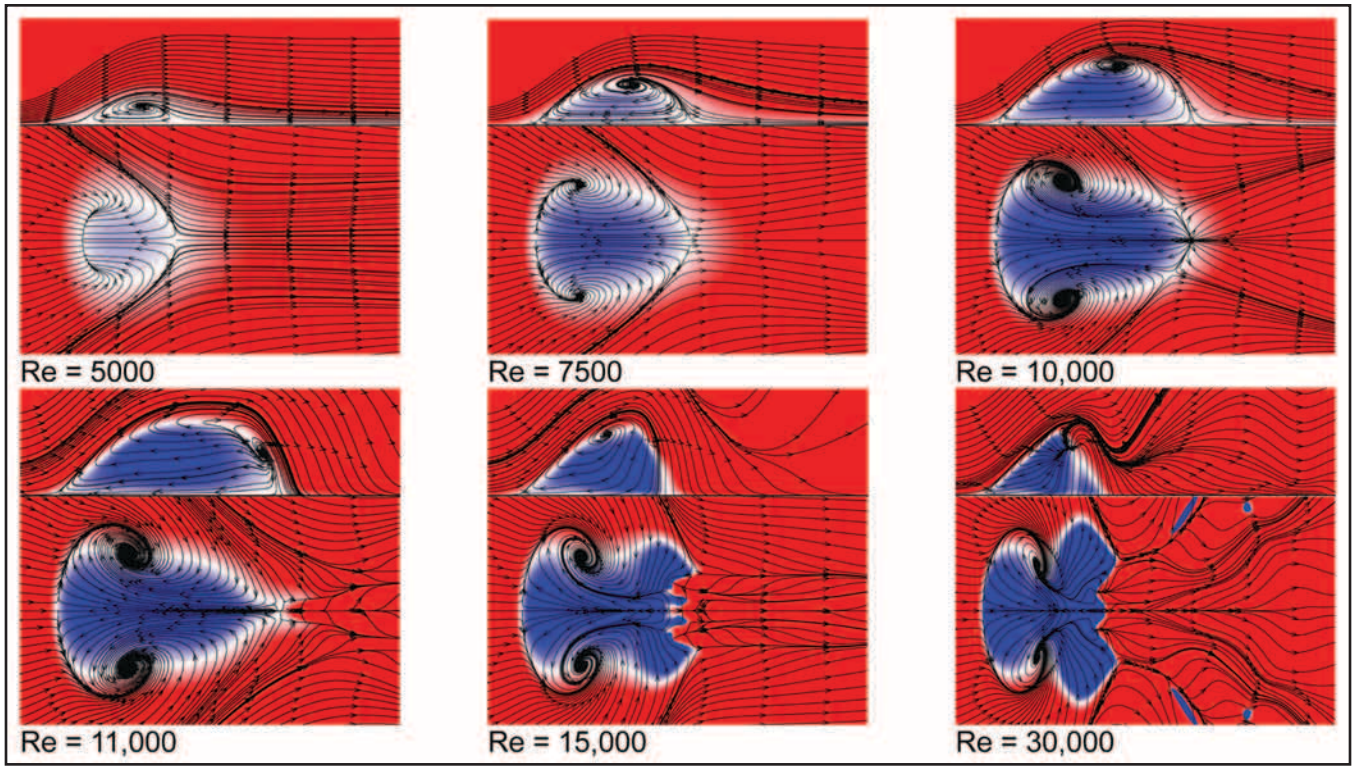

Figure 4. Streamline patterns and color contours of streamwise velocity in the symmetry plane (z = 0) and limiting streamlines and color contours of spanwise vorticity on the plate (y = 0). Blue indicates reverse flow

are symmetric with respect to the spanwise direction. Results for Re = 30,000 are provided in Figure 6. As the Reynolds number is increased, the spread in length scales between the energy bearing and dissipating structures increases. The wake is populated by an increasing amount of small-scale structures, and the streamwise spacing of the spanwise coherent structures is reduced.

We also investigate turbulent separation for the low aspect ratio asymmetric 3-D diffuser geometry of J. Eaton at Stanford University. The Reynolds numbers associated with this experiment are too high for DNS. Reynolds-averaged Navier-Stokes (RANS) that models all unsteady fluid motion is a low cost alternative but does not provide insight into the unsteady flow physics. Large-eddy simulation (LES) where the more isotropic small-scale turbulence is modeled and the large-scale resolved motion is captured and computed time-accurately is a viable alternative but suffers from a restrictive near-wall resolution requirement. A fairly new class of models that has evolved recently is the family of hybrid RANS/LES models that blend continuously between RANS near the wall and LES/DNS away from the wall. Such models are advantageous, as modern RANS models are calibrated to perform well near the wall, while LES models have shown good performance for unbounded flows. We are testing and improving hybrid RANS/LES models for turbulent

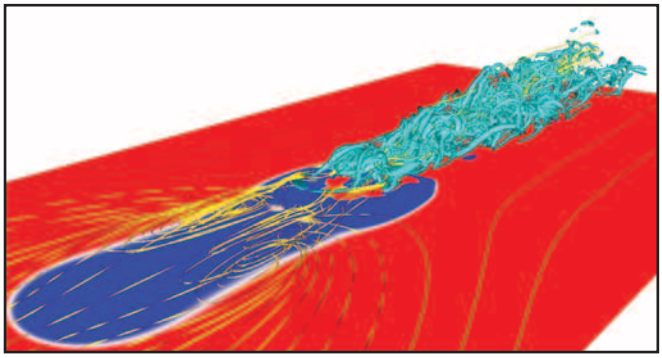

Figure 5. Perspective view (Q = 6000) of unsteady separation bubble at Re = 15,000

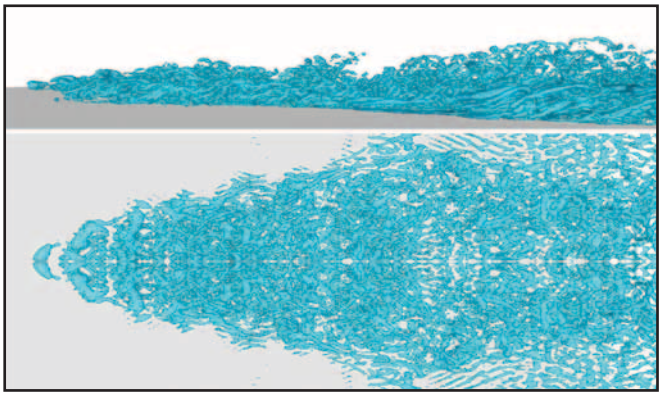

Figure 6. Instantaneous flow visualizations of vortex identification criterion (Q = 6000) for Re = 30,000

square-duct flows at a Reynolds number based on bulk velocity and hydraulic diameter of Re = 10,000. Promising candidates that perform well for the square-duct flow are then applied for simulations of the asymmetric diffuser geometry. For these simulations, we are employing an in-house developed higher-order-accurate compressible finite-volume code.

The computational grids for our hybrid RANS/LES simulations of the turbulent square-duct flow had 192×40×40 (coarse), 192×54×54 (medium), and 384×80×80 (fine grid) cells. In addition, as a reference, we also carried out a DNS on a super fine grid with 384×160×160 cells (Figure 7). Instantaneous flow visualizations from hybrid RANS/LES simulations based on the filter-based RANS hybrid RANS/LES model in combination with the 1998 Wilcox k-w model and the Rumsey and Gatski explicit algebraic stress model are provided in Figure 8. The most small-scale structures are, as expected, observed for the DNS on the super fine grid. As the grid resolution is reduced, more and more of the resolved turbulence kinetic energy is absorbed into the unresolved turbulence kinetic energy, and less and less unsteady flow structures are captured.

The computed bulk velocities were 1.03 for the DNS and 1.09 (coarse), 1.05 (medium), and 1.02 (fine grid) for the hybrid RANS/LES model. Interestingly for the hybrid RANS/LES simulations, the RANS bulk velocity is approached in the coarse grid limit, and the DNS bulk velocity is approached in the fine grid limit. The testing of hybrid RANS/LES models in duct flow simulations turned out to be quite compute time intensive. Grid resolution studies are required for each model for determining if the respective model shows the correct asymptotic behavior towards the RANS and DNS limit.

We then employed the same models for simulating the flow through the Stanford asymmetric diffuser for which measurements indicate turbulent separation. The diffuser inflow Reynolds number based on bulk velocity and hydraulic diameter is Re = 15,380. The

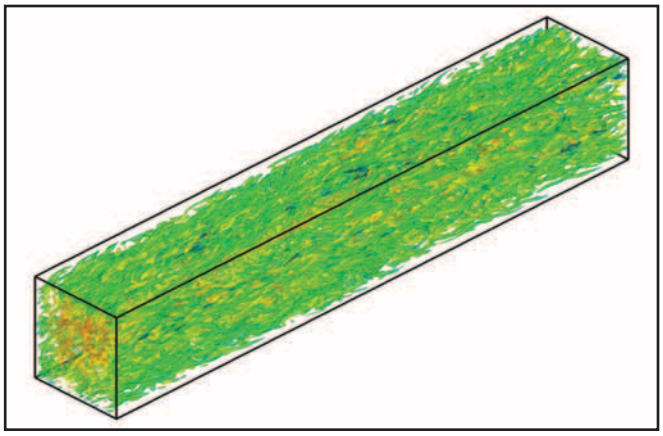

Figure 7. DNS of square-duct flow. Isosurfaces of Q = 5

computational domain was broken up into two subdomains, one for the approach channel flow and one for the diffuser. The approach channel flow subdomain had a length×height×width ratio of 6.66×1×3.33 with 92×76×96 cells. The diffuser grid had 268×76×96 cells. The total number of cells was 2.7 million.

Instantaneous flow visualizations obtained from a hybrid RANS/LES simulation are shown in Figure 9. The approach channel flow subdomain simulation is seen to provide sufficiently de-correlated turbulent inflow data for the diffuser flow subdomain simulation. Inside the diffuser, the flow was observed to separate near the lower corner where the diffuser begins to open up. Different results were obtained depending on which hybrid RANS/LES model with what underlying RANS model was employed. Turbulence models that performed well for the square-duct flow did not necessarily also work well for the separating turbulent flow inside the asymmetric diffuser. A problem encountered during the square-duct flow and diffuser flow simulations is that the time-dependent data can only be averaged in time and not in a homogeneous direction (such as the z-direction for 2-D geometries). Therefore, long time-averaging intervals are required for good time averages, which are a prerequisite for obtaining turbulence statistical data. Because we are testing sev-

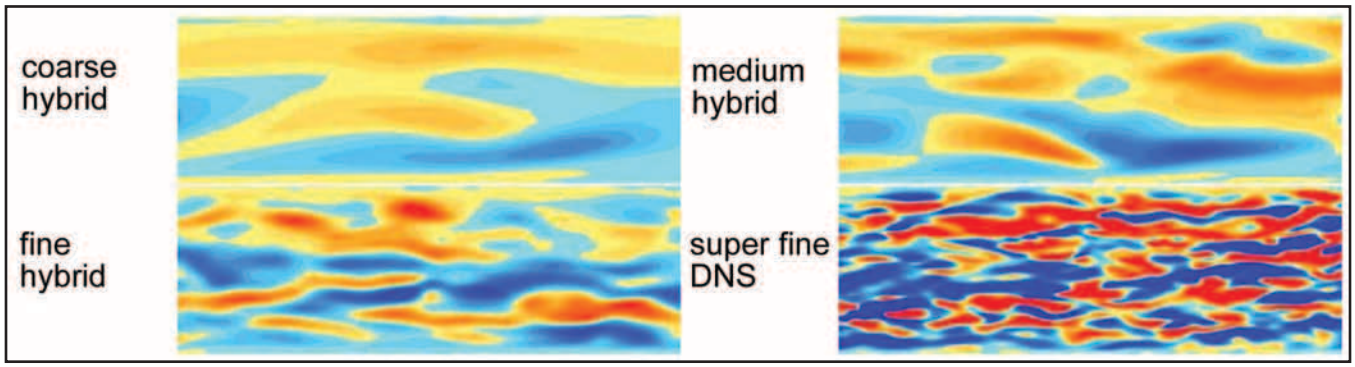

Figure 8. Square-duct flow at Re = 10,000. Isocontours of streamwise wall vorticity, $-10<\omega_x<10$

eral models, considerable computational resources are required for this research to be feasible.

Flow separation deteriorates aerodynamic performance, increases structural loads, and may result in an undesirable acoustic signature. For Navy applications, separation is mostly 3-D and unsteady. Our simulations of generic 3-D separation bubbles advance the understanding of 3-D separation, which paves the way for separation control devices that will ultimately lead to new strategies for reducing drag and acoustic signature. The accurate and reliable prediction of turbulent flows with adverse pressure gradients, separation, and non-equilibrium turbulence requires the development of new turbulence modeling strategies. Our square-duct and diffuser flow simulations allow us to validate and advance hybrid turbulence modeling approaches. New modeling approaches will benefit engineers and scientists involved in the design of high Reynolds number devices.

Typical simulations were carried out on up to 512 processors on the IBM P5 at the Navy DoD Supercomputing Resource Center (DSRC), the Cray XD1 at the Naval Research Laboratory, and on the Cray XT5 at the Arctic Region Supercomputing Center DSRC. This

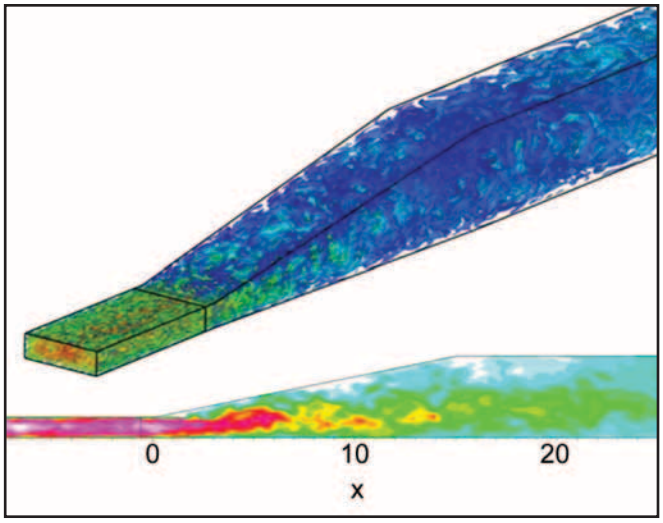

Figure 9. Asymmetric Stanford diffuser. Isosurfaces of Q = 0.01 colored by u-velocity and isocontours of u-velocity (0…1.5) at midspan plane

research was funded by the Office of Naval Research under grant number N00014-07-1-0401, with Dr. Ronald Joslin serving as program manager, and received additional support through the Department of Defense High Performance Computing Modernization Program challenge grant C3B.

Optimization of Finite Element Codes for Stimulated Brillouin Scattering in Photonic Crystal Fibers

Dr. Juan Carlos Chaves, Army Research Laboratory, User Productivity, Enhancement, Technology Transfer and Training (PETTT); and Drs. John Nehrbass and Bracy Elton, Air Force Research Laboratory, PETTT

Problem

The onset of stimulated Brillouin scattering (SBS) is the primary obstacle limiting power scaling in photonic crystal fiber lasers. Such lasers have many important applications including directed energy weapon (DEW) systems and atmospheric light detection and ranging (LIDAR). Researchers at the Laser and Optics Research Center (LORC) at the Department of Physics in the United States Air Force Academy (USAFA) are investigating the SBS process with the aim of increasing the delivered power as components of directed energy weapon systems. Although fiber lasers have many excellent characteristics for tactical applications, they present relatively low optical nonlinear thresholds where the onset of SBS is the most limiting phenomenon. For example, acoustic vibration in the optical fiber media due to stimulated Brillouin scattering is critical to understanding frequency shifts due to electrostriction phenomena (See figure below).

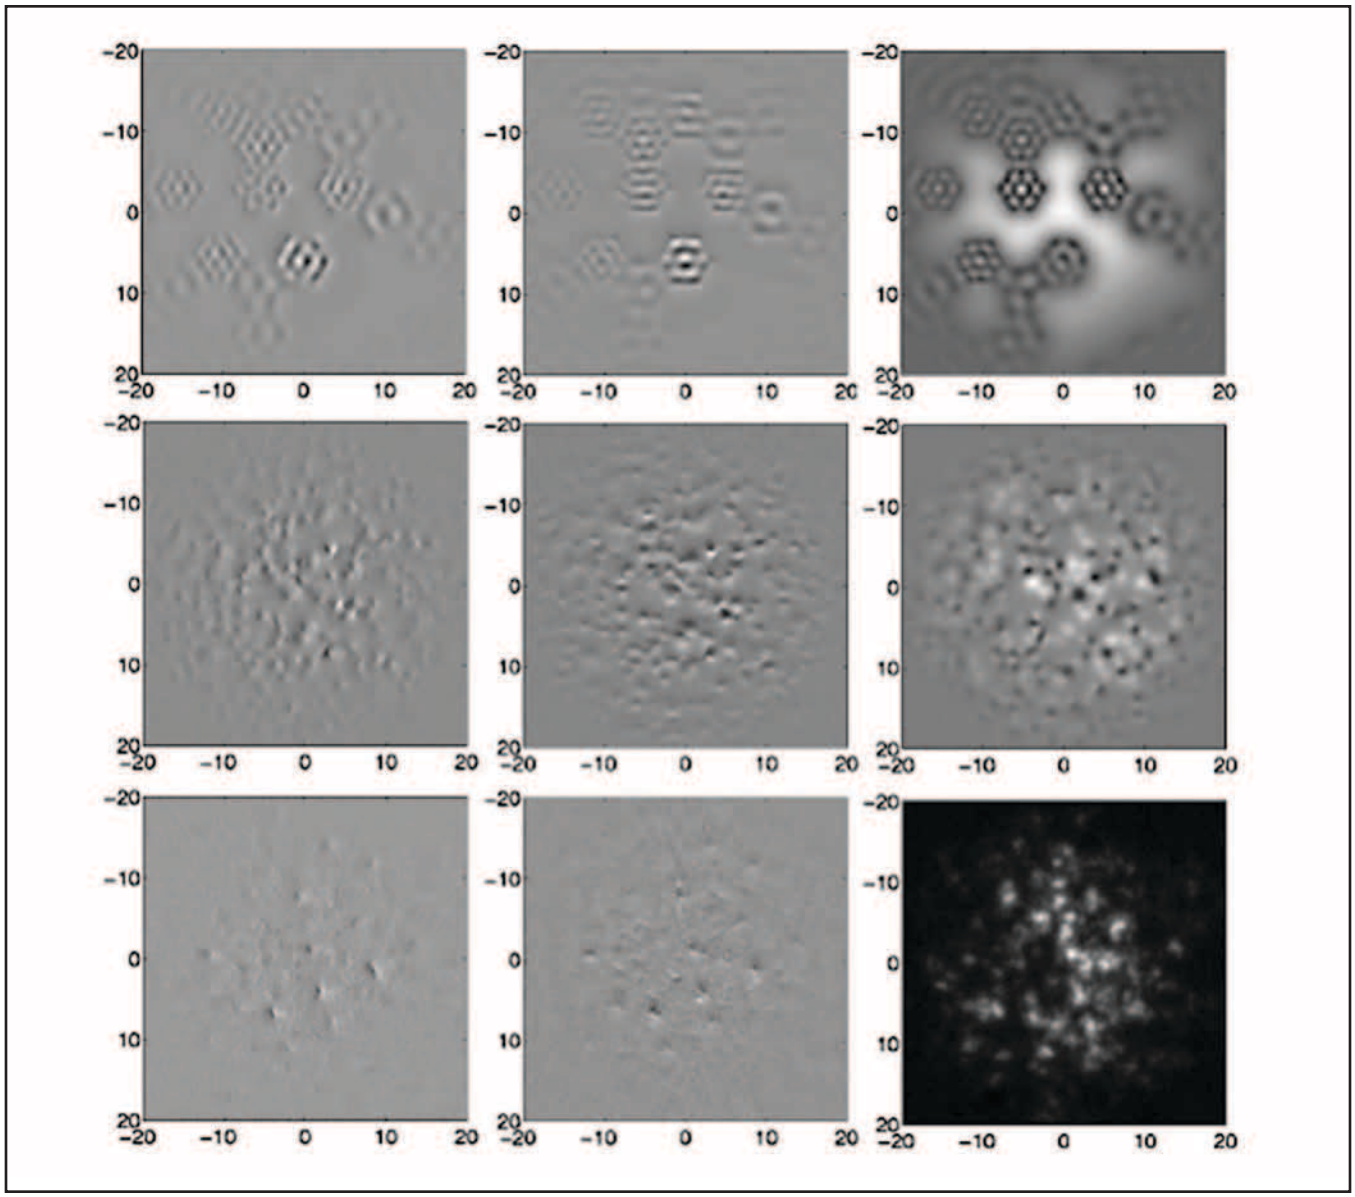

Acoustic displacement field components for largest, optimal, and smallest domain sizes at the peak Brillouin gain frequency

The User Productivity, Enhancement, Technology and Training (PETTT) user and LORC Principal Investigator have applied a hybrid C-MPI/Parallel MATLAB strategy to investigate these complex optical systems on the Air Force Research Laboratory DoD Supercomputing Resource Center (AFRL DSRC) high performance computing (HPC) platforms. A model based on numerically intensive finite element methods and MatlabMPI has been developed to investigate various scenarios for increasing the output power of photonic crystal fiber lasers. Through the use of innovative parallel MatlabMPI compilation technology, PETTT has provided optimized codes that perform 33 percent faster than previous versions. Additionally, these codes can be scaled to large numbers of computational nodes since only a single MATLAB/MATLAB compiler license is initially required. This effort allows USAFA users to model significantly larger laser systems by exploiting innovative parallel MATLAB technologies and enabling important new discoveries in this research area. Additionally, this has resulted in significant reduction of MATLAB license fees for the DoD High Performance Computing Modernization Program (HPCMP) HPC platforms.

Methodology and Results

MatlabMPI compilation technology dramatically extends the applicability of MatlabMPI, as it allows creating parallel MATLAB applications with a single MATLAB/MATLAB compiler license and running the compiled parallel applications without requiring any MATLAB licenses. This becomes especially important on distributed-memory architectures such as most platforms of the HPCMP in which one MATLAB license per node is required to run the unmodified version of a MatlabMPI code. For example, on most DSRC systems, there is an insufficient number of available licenses to run on the desired number of nodes. In order to address this problem, MatlabMPI programs are compiled into stand-alone executables through the MATLAB compiler. Although the approach is straightforward, the implementation and transfer of a workable/robust solution to a DSRC system is a complex task. Dr. John Nerhbass accomplished most of the MatlabMPI porting to a compiler friendly version. Dr. Juan Carlos Chaves adapted and transferred the solution to the AFRL *Falcon* cluster and provided extensive documentation, LSF/Perl specially tailored scripts, and coordination with the user. Dr. Bracy Elton resolved several SSH connectivity problems.

The PETTT user provided the following feedback:

> "Dear Dr. Chaves: Your support has been excellent. You provided a working solution in a very short amount of time. This solution has saved us 1/3 of our computational costs. Furthermore, you documented the solution so well that a noncomputer scientist was able to get it working on production code (not test code) within a day without any additional assistance. We have discovered using our model, and the MatlabMPI solution you developed, that by engineering acoustic domain sizes, we can increase the output power of these lasers by about 8x from about 500 watts to 4 kilowatts. We currently have an article describing these findings in preparation for submission to Physical Review B. As soon as I have a complete draft, I will forward it along. I'd also like to say that as an academic institution, we compete for our research funding with civilian institutions. I'm convinced that our HPC work, aided by PETTT support, has set us above our competitors and allowed us to win our grants and hence continue to operate."

DoD Impact

Photonic crystal fiber lasers have many important applications including directed energy weapon systems and atmospheric LIDAR systems. LIDAR systems exploit the scattering characteristics, shorter wavelength, and high directivity of laser energy to image objects at much finer resolution than is readily attainable with longer wave millimeter and microwave systems. Civil applications of LIDAR systems include detection of weather phenomena and pollutants in the atmosphere. The military applications of LIDAR systems include range-finding for targeting and missile defense systems. Directed energy laser weapons and LIDAR systems are an increasingly important technology for a variety of DoD and homeland security applications.

Air Force Research Laboratory DoD Supercomputing Resource Center

From the Director's Desk – Frank Witzeman

Frank Witzeman
Director, AFRL DSRC

Since the last issue of *HPC Insights*, the AFRL DSRC completed the decommissioning process for both of our systems acquired in the DoD HPCMP Technology Insertion 2005 (TI-05) activity. *Eagle,* a 2048-processor SGI Altix 3700 system, was split into two systems and delivered to new owners in the DoD intelligence and rapid fielding communities. *Falcon,* a 2048-processor HP XC Opteron system, is being redeployed to a unit in the Naval Air Warfare Center. Our last standing major HPC system is *Hawk,* a 9216-core SGI Altix 4700 acquired under TI-07. *Hawk* currently represents the only shared-memory architecture among the 16 HPC systems in the DoD HPCMP enterprise. It's a unique configuration of 18 nodes (or partitions) consisting of 512 cores for each node and 1 to 2 TB of memory per node, 20 TB of total memory.

As we implement our new 44,000-core Cray XE6 system, recently obtained through the TI-10 process, it becomes increasingly important to facilitate any necessary transitions of the *Hawk* shared-memory applications onto the new platform. Another consideration is our ability to effectively manage the anticipated workload on such a large core-count system. User applications that are able to scale to tens of thousands of cores must share priorities with thousands of user jobs on the order of 4 to 16 cores. The nature of high performance computing within the DoD HPCMP has become, and will remain for some time, a struggle between capability and capacity computing.

So, what's the next step in determining how we meet user requirements in 2011 and beyond? There is some emphasis on application software development, to enable users to scale up codes to investigate massive problems on existing and emerging HPC systems. Another focus is related to hybrid supercomputing architectures consisting of combinations of standard chips and graphics processing units (GPUs) programmed for fast, general-purpose computations. Yet another area of growth is the "HPC Cloud," where users simply buy core hours, memory, and storage as services accessible from the Internet. The benefits of the above alternatives appear to be attractive from a number of perspectives, and they need to be investigated from the DoD's position on security and protection of scientific and technical information.

Perhaps the DoD HPCMP can leverage the Defense Advanced Research Projects Agency (DARPA) initiatives in Ubiquitous and Omnipresent High Performance Computing. Under these projects, DARPA is attempting to stimulate industry's development of new systems/technologies that address lower power consumption, higher user productivity, and resiliency to cyber attack (Reference: DARPA Transformational Convergence Technology Office Broad Agency Announcement DARPA-BAA-10-37, March 9, 2010). Although DARPA is seeking revolutionary advances in computer science and engineering, we can take notice and perhaps advantage of the resulting ideas in power reduction, performance dependability, and cyber security. In fact, two of the five "Challenge" problems in the Ubiquitous computing area propose the use of DoD HPCMP applications from the performance benchmark suite or from the Computational Research and Engineering Acquisition Tools and Environments (CREATE) effort.

Whether we continue to pursue commodity supercomputers in the future DoD HPCMP Technology Insertion process or chase more innovative solutions, we must remain committed to the DoD scientific and engineering user community. Our objectives are consistent with providing the most secure and dependable high performance computing capabilities while reducing operating costs and exceeding user productivity expectations. As an integral member of the larger DoD enterprise, we are naturally partnered with our users who are solving the most complex problems and delivering superior technologies in defense of our Nation. Supposedly HPC can be bought affordably from a cloud, grid, or even a mobile container. However, synergy among DoD HPC suppliers and customers working on the same team in the best environments is priceless!

AFRL DSRC Prepares for Cray XE6 Supercomputer

In mid-2007, the AFRL DSRC began preparing its facility for the Technology Insertion 2010 (TI-10) supercomputer. Our facility team knew there would be challenges to overcome. To help forecast the potential hurdles, it brought in an outside company specializing in IT complex integration and current technology. Although there were many unknown aspects of the 2010 system, including heating, ventilation, and air conditioning (HVAC) and electrical requirements, brand, and size, it was decided to ensure that the facility should be able to handle up to a 4-MW system. Additional HVAC and space requirements would wait until the system was closer to conception.

Icon designed to represent the Cray XE6

Pictured above, workers install two 600-ton cooling towers with pipe stanchions that will support the interconnecting feed and return pipes that will provide a balanced environment for the XE6 to operate efficiently

In 2009, the bids were in and the choice was made. The new AFRL DSRC high performance computer (HPC) was a 43,712 compute core Cray Inc. XE6. The system was named *Raptor* after the Air Force fifth-generation fighter aircraft, the F-22.

The total weight of the system, approximately 100,000 lb, compares with about thirty-three 2010 Toyota Prius automobiles, and each of the 15-plus cabinets requires the cooling equivalent of four average homes.

The system is considered energy efficient and is air cooled with the option of ECOphlex™, an inert coolant that converts liquid to gas, which helps reduce the need for additional lesser efficient computer room air conditioners (CRACs). Our newly installed HVAC system should help the AFRL DSRC greatly reduce energy costs and consumption, especially during the winter months, by using outside air when the external temperature is less than 50 °F (the average annual temperature in Dayton, Ohio, is 51.7 °F).

The preparation of our facility for this immense system has every one excited. From the network administrators, computer engineers, and the support staff, the challenges are ready to be met. With its multicore architecture, distributed memory, AMD powerful multicore processors, and the revolutionary Gemini™ interconnect, *Raptor* will allow us to continue to deliver premier and innovative services to our user community.

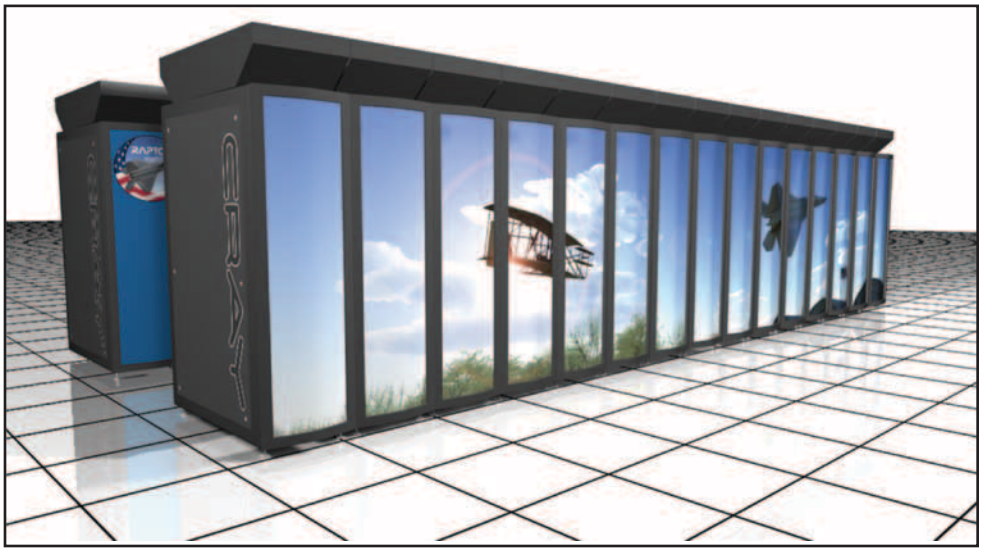

Conceptual computer rendering of Cray XE6 with front mural graphics inspired by our blending of history and aeronautical innovations

Army Research Laboratory
DoD Supercomputing Resource Center

From the Director's Desk – Charles J. Nietubicz

Charles J. Nietubicz
Director, ARL DSRC

I write my last ARL DSRC Director's column with mixed emotions. As I take a break from the many exciting things that are happening within the HPCMP and the ARL DSRC, I realize that I have over 39 years of Government service and the time has come to retire. This will happen on October 1, 2010. While the excitement of entering a new phase of my life is before me, the end of being your Director of the ARL DSRC is coming to a close. I have been fortunate since the beginning of the HPCMP to have the great opportunity to serve as Director and see our computational capability increase exponentially, our computational environment become more robust, and our collective energy focused on providing the best HPC for our scientists and engineers. Most importantly, we can all look with pride to being an integral part of helping to provide the men and women of our armed forces with the finest weapon systems in the world.

Before I reflect briefly on how we started and where we are, I would like to point to the ARL DSRC articles included in this edition of *HPC Insights.* First, we were pleased to host the 2010 HPCMP Joint Educational Opportunities for Minorities (JEOM) Research Workshop at ARL. This excellent workshop was pulled together by Bob Sheroke and his team who worked hard to provide the attendees with a learning experience on HPC as well as ongoing research at ARL that utilizes the HPCMP capabilities. Please see this article for additional details. Also, the 29th of June was a great day here at the ARL DSRC as we had our formal dedication of the new SGI Altix ICE system. While the last edition of *HPC Insights* identified *Harold* as the new machine name and gave some information, the current article describes the activities that took place on the dedication day. I would also point your attention to the "HPC at Work" section where our onsite PETTT staff has articles on "Optimization of Finite Element Codes" and "Parametric Studies of Quantum Interference Filters."

I could also spend some time here talking about the technology, the advances at the Center, our current projects, and more, but I would rather reflect, in my last column, on what has made the HPCMP outstanding, what will continue to take it to the next level, and what makes it so hard to leave – it is the "People of the HPCMP" – all those who have contributed, some who went on to other areas, those who have provided guidance, insight, and leadership, and those who continue to work every day to make this Program better.

I still vividly remember the exciting times and flurry of activity at the Army Research Laboratory in 1994 as we were developing a response to a DDR&E initiative requesting submission of proposals to become a DoD Major Shared Resource Center. Similarly this same activity was ongoing at other Service Laboratories. We all had questions: What is an MSRC? What does early access mean? What should DREN look like? Parallel computing vs. vector – will it ever work? What about integration contractors? Where is Fern St? What is PET, etc.? If we look back to those questions and see where we are today – there is one common theme. That is, we have been fortunate to have smart, hard-working, and dedicated people who provided the answers to these questions and more. I look with pride to so many people who have given unselfishly of their time and talent to build and sustain a Program that is such a benefit to others. I would like to take this opportunity to thank all of you who I have had the pleasure of working with, for letting me be a small part of such a great vision and success, and allowing me to help build the HPCMP into what it is today. While I will miss being a part of the advancement in technology and the environments that make HPC an integral part of our science and engineering research, I know I will miss even more my colleagues and friends here at the ARL DSRC and throughout the HPCMP. Thank you again for letting me be a part of this outstanding evolution and transformation of HPC and for all of your help and support. As I leave at the end of this FY, I will take with me all the great memories that you the "People of the HPCMP" have given me.

Thank you, God Bless, and 'Sto Lat.'
Charlie

ARL TI-09 System *Harold* Dedicated in June 29 Ceremony

An official dedication and ribbon cutting ceremony was held on June 29 to celebrate the contributions made to the Army and the DoD by Harold Breaux. A capacity audience filled the BRL Hall, an iconic meeting room that has been part of the research laboratory's culture for many years, to share in the event. On hand were numerous family members and ARL, Army, and DoD officials who were all pleased to share the day with Harold.

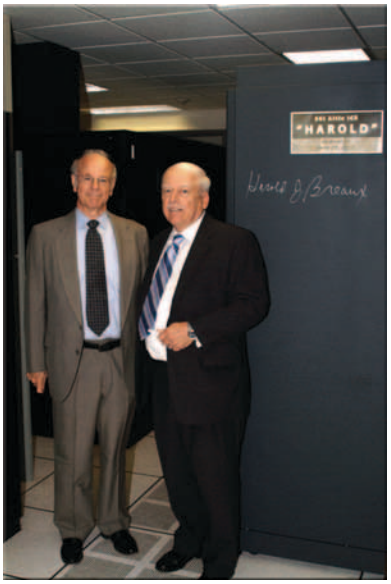

ARL DSRC Director Charlie Nietubicz with Harold Breaux and the Harold *supercomputer*

The system that was dedicated is an SGI Altix ICE 8200 Linux cluster and the largest of its kind ever deployed at the ARL DSRC. It offers 10,752 Intel Xeon 5500 series processor cores and higher memory bandwidth than other high performing computers here and across the DoD HPCMP. The computer arrived in June 2009 and has been in production use since October 2009.

Harold has spent over 33 years in Government service serving as a research mathematician, Chief of the High Performance Computing Division at ARL, and as manager of ARL HPC Systems. One of his most significant contributions to supercomputing came through his lead role as an Army representative on a DoD working group in 1992, which founded the multibillion dollar HPC Modernization Program. Breaux is also recognized for writing the proposal that resulted in ARL's designation as one of the first three DoD laboratories chosen as a Major Shared Resource Center (MSRC), later renamed as a DoD Supercomputing Resource Center (DSRC). His achievements include the Army R&D Achievement Award, Superior Civilian Service Award, and Meritorious Civilian Service Award. In 2005 he was given a "HERO" award by the Defense Department for his long-term contributions to the HPCMP.

After the ceremony, the group moved from BRL Hall to the computer room housing the *Harold* system to sign the side panel of the supercomputer, leaving a record of the day and creating a tribute to the man and his legacy. Harold remarked that he was humbled that the system shares his name; to him, it is a permanent and public reminder of the "good ride" he has experienced in his career with the Army.

HPCMP Director Cray Henry signs the Harold *supercomputer*

(Left to right) Dr. Paul Dietz, Director, Army Materiel Systems Analysis Activity; Harry Reed, retired Division Chief, BRL/ ARL; Dr. Reed Skaggs, Associate Director for Plans and Programs, ARL; Harold Breaux; Charlie Nietubicz, ARL DSRC Director; Cray Henry, Director, DoD HPCMP; Jeff Gosciniak, Lockheed Martin, Program Manager, NGTS Contract

2010 HPCMP JEOM Faculty Workshop

By Eldred Lopez, Lockheed Martin, ARL DSRC

The ARL DSRC hosted the HPCMP Joint Educational Opportunities for Minorities (JEOM) Faculty Workshop, which provided Historically Black Colleges and Universities/Minority Institutions (HBCU/MI) Program faculty a firsthand perspective of Army research programs. These programs are impacted by the use of advanced computing and supercomputing resources and will provide faculty with tools and resources to introduce and implement a computational science program within their existing science and engineering academic curriculums.

The workshop focused on developing and sharing information, resources, and strategies to assist HBCU/MI faculty by establishing awareness of DoD/Army research programs and supporting the professional development of undergraduate students in science, technology, engineering, and mathematics (STEM) disciplines at HBCU/MI institutions. The key goal of the workshop was to introduce HBCU/MI faculty to established methods and practices of computational science programs in academic curricula currently in place at nationally recognized academic institutions.

To establish awareness, the workshop brought in speakers from different universities such as keynote speaker Dr. Patricia Teller from the University of Texas of El Paso. Her topics touched upon how universities could prepare students to meet educational challenges of the future and how to have these students prepared to meet the expectations of employers and graduate schools.

Invited guests were also encouraged to have roundtable discussions on developing interests of students in computational science, and to share the experiences of their respective programs. They were also encouraged to create new contacts and develop a flow of information.

The workshop was held June 29 through July 1 with 26 invited speakers and guests, representing nine universities.

Attendees and staff of the HPCMP Joint Educational Opportunities for Minorities (JEOM) Faculty Workshop

Arctic Region Supercomputing Center
DoD Supercomputing Resource Center

From the Director's Desk – Frank Williams

Frank Williams
Director, ARSC DSRC

ARSC has a 15-year proven record of success in running large HPC systems. The Center's achievements are directly attributable to a culture our staff establish by providing unique connections among academic and defense computing communities with the computers, data storage systems, high-speed networks, and next-generation experimental systems needed by DoD HPC users to solve big problems.

Sustaining this level of commitment to excellence extends to the ARSC DSRC new 11,648-core Cray XE6 obtained under the HPCMP Technology Insertion acquisition process for 2010. It will be part of a full complement of HPC hardware and software services that will be physically located at the U.S. Army Engineer Research and Development Center DSRC. The ARSC DSRC will remotely operate the new HPC resources from Fairbanks as an open research system.

This new supercomputer has been dubbed *Chugach* (pronounced choo-gatch) and is named after a mountain range in Alaska. *Chugach* is a little over three times the size of *Pingo*, the ARSC 3456-processor Cray XE5, but with a faster interconnect and the new Cluster Compatibility Mode environment, which makes it easier to run applications needed by our users.

The ARSC DSRC is scheduled to be the first Center in the Program to provide user access to the HPCMP Enhanced User Environment (EUE), early in the first quarter of 2011. Under the EUE umbrella are (1) CUES-Common Utility Enhancement Services, (2) CWFS-Center Wide File System, and (3) SLM-Storage Lifecycle Management.

The timeline for bringing *Chugach* and EUE services online for general user access is anticipated in January 2011. The ARSC DSRC will continue to operate *Pingo* into FY2011 as an unallocated system to provide uninterrupted service to the open research community.

Being one of two EUE test and development sites for the HPCMP, ARSC will have multiple test systems in service to support all of the Centers as they support users. In that role, ARSC will test EUE software and hardware prior to release to the other DSRCs. Following a similar, parallel path, the ARSC DSRC will lead the Storage Test and Sustainment activities for the HPCMP and will be the storage systems test site for the Program.

The ARSC DSRC HPC Mass Storage Specialist Gene McGill is coleading the implementation of the EUE storage initiative project.

Gene McGill in front of the StorageTek SL8500 at ARSC DSRC

ARSC DSRC
DoD Supercomputing Resource Center

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

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

Dr. Robert S. Maier
Director, ERDC DSRC

HPCMP Needs an Exascale Computing Initiative Now

John Levesque remembers the transition from vector to parallel computing. Levesque, a long-time member of the Cray staff, is responsible for helping his customers prepare for exascale computing. He recently spoke on this topic to the Workshop on Application Drivers for Exascale Computing at the National Institute for Computational Science. Levesque said that in order to move existing all-MPI applications to exascale architectures, applications must be refactored, or analyzed to identify high-level, shared-memory parallelism. Levesque made a clear distinction between refactoring and simply supporting low-level OpenMP, which is easily achieved by automatic compiler parallelization, but not sufficient to effectively use a shared-memory, many-core node. He described the ORNL Application Readiness Effort, in which teams will ready DOE applications for an MPP system equipped with accelerators on each node. Teams will first identify the most important kernels and use hand-coded CUDA

Garnet, *the new Cray XE6*

(as well as directive-based OpenMP) to prepare a hybrid application, comparing performance to that of a present-generation multicore system. Once the kernels are analyzed, the communication and memory management of the hybrid app is examined to identify kernels that can be run asynchronously with other kernels and communication. One of his recommendations is that teams should strive to develop a performance-portable application that runs on existing multicore as well as target hybrid systems. It's important that we launch a similar effort in the HPCMP now. Recalling our past experience in moving codes from vector to parallel computing, we should allow 3 years for an individual code, and at least two or three times longer for phased project awards. Establishing this type of initiative will take some planning, too. It's time to start.

So how can we afford to undertake new software initiatives in a time of shrinking defense budgets? Tightening belts, while providing users the latest technology, is now the operating mode. We should always practice such thinking. Just as the economy affects our personal lives, so it affects the workplace.

The ERDC DSRC has recently taken delivery of a new Cray XE6 supercomputer with 20,224 compute processors and 194-teraflop peak performance. Keeping with the theme of precious minerals, we named this system *Garnet.* It joins *Diamond,* an SGI Altix ICE 8200, and *Jade,* a Cray XT4, and brings the ERDC DSRC computing capacity to 437 teraflops.

In the following article by Greg Rottman, you'll read about the infrastructure and financial planning behind the installation and operation of our TI-09 and TI-10 systems, as well as future systems. This planning has allowed us to save Program funds and to provide space for the T1-10 open system *Chugach*, while still leaving room for a TI-11 system.

ERDC DSRC
DoD Supercomputing Resource Center

Practically Green

By Greg K. Rottman, ERDC DSRC Assistant Director

When you are a Government facility that must get reimbursed for all costs associated with the services provided to a customer, it is important to be cost conscious. There are many ways to reduce cost. This article will describe how we at the ERDC DSRC addressed reducing construction costs of our central energy plant.

With limited funding, we set out to design a power delivery infrastructure that was scalable and efficient. To that end, we developed plans for an uninterruptible power supply (UPS) system that was scalable in units of 2 MW with a maximum capacity of 8 MW. Our initial assumption was that we would have to build a shelter for the UPS, switchgear, and batteries. But as soon as we had funding for the equipment, we discovered we did not have authority to construct the shelter—so we went back to the drawing board.

After a few weeks of looking for alternatives, we came upon a company, S&C Electric, that manufactures outdoor UPS systems. This is exactly what we needed to keep our project on schedule. We could design and build an electrical delivery system that did not require construction of a shelter. An additional benefit of the system was that it could deliver UPS-protected power at utility line voltages—in our case 13,800 volts. This system would operate at better than 98 percent efficiency.

Being able to deliver protected power at that voltage level would reduce the conduit and cabling work by a factor of at least 25 times. This means there would be 25 times less conduit to run and 25 times less copper cable to buy. Also, there would be an efficiency increase due to less loss of energy because of less cable through which to transmit.

Once we had started down the efficiency trail, we made a conscious decision to examine other aspects of our design for possible ways to be more efficient. One thing we recognized was that anything that produced heat was doubly inefficient. First, it was losing energy to heat, and second, the heat was in a building and added to the air conditioning load.

To reduce loss to heat, we worked to increase the efficiencies of the transformers, which resulted in a reduced air conditioning load. We asked a company called Power Quality International (PQI) to analyze our application. They delivered step-down transformers and a connection scheme that increase the efficiency of the system while reducing the heat released into the conditioned space. These transformers have a return on investment of less than 3 years.

Our investigations turned over a revealing (to us) piece of information. We found upon examination that computer power supplies appeared to run more efficiently at higher voltages. That set us to thinking, "How high could these voltages go?" We found that the maximum voltage of most computer vendors was 240 volts. The highest common U.S. voltage is 208 volts. So if we could set all the computer system equipment to 208 volts, we would only need to step down the 13,800-volt feed twice, once from 13,800 to 480 volts and then 480 to 208 volts. We could also place the 13,800- to 480-volt transformer just outside the building, leaving the heat outside to dissipate and using the PQI-designed transformer to step down from 480 to 208 volts.

Since installing this system, we have been able to acquire a computer that accepts 480-volt electrical feed, and all the peripheral equipment operates at 240 volts, providing that much more efficiency gain.

On the cooling side, we built an infrastructure that was capable of scaling in nominally 500-ton increments up to 2000 tons. Each of the computer systems use chilled water to remove heat from the room, while computer room air handlers (CRAH) provide the heat removal for the peripheral equipment such as data storage disks and network gear. These CRAH units work together to effectively deliver the cooling and airflow necessary to maintain the proper computer room environment, while incorporating variable speed fans that minimize the current draw required to produce the necessary airflow.

Remember my saying earlier in this article that we did not have authority to construct a shelter? Well that posed another interesting challenge. How could we design, contract, and build a 10,000 sq ft computer room in less than a year without authority to build it? We worked closely with the U.S. Army Corps of Engineers, Vicksburg District, and designed the computer room while simultaneously working to get authority from Congress to construct it. Once we received authority, we contracted to have the shelter built with a contractual requirement to be complete in 110 days. The contractor, Brassfield & Gorrie, out of Birmingham, Alabama, was up to the challenge and delivered the project on time.

With advanced planning, adaptable thinking, and a great deal of effort by many people, we now have a 10,000 sq ft computer shelter that is calculated to require only 30 percent overhead on the power delivered to the computers, as compared with nearly 100 percent overhead on our legacy systems of just a few years ago.

In a 1-MW data center with power cost at $0.10/kWh, every 10 percent increase in efficiency equates to almost $88,000 in savings in operating costs. In a 4-MW center, this value rises to in excess of $350,000. Being able to reduce your data centers overhead by 60 to 70 percent can result in a huge favorable impact to your operations budget.

This is what I call practically green.

ERDC DSRC new 10,000 sq ft computer room

Maui High Performance Computing Center DoD Supercomputing Resource Center

From the Director's Desk – David Morton

David Morton
Director, MHPCC DSRC

FY10 continues to be a year of significant growth and success for the MHPCC DSRC that fits in well with the SC 2010 theme of "The Future of Discovery."

It has been becoming increasingly obvious to HPC professionals in our DoD community that the hardware capabilities are increasing much more rapidly than the software used to harness that power. The HPCMP recognized this trend several years ago and has invested significant resources in the development of applications that are better suited to scale and deliver the next generation of capabilities to our DoD researchers and warfighters. Two of those HPCMP applicatons that MHPCC participates in are the HPC Software Applications Institute (HSAI) program (we host the Space Situational Awareness HSAI) and the Computational Research and Engineering Acquisition Tools and Environments (CREATE) program.

Operationally, MHPCC's role in the CREATE program has increased significantly. In addition to providing administration and management services for the infrastructure servers and virtual machines, MHPCC embarked on a precedent-setting approach of providing dedicated HPC resources for the CREATE AV software development releases. Through negotiations with CREATE team members, MHPCC dedicated 2,000 cores of the *Mana* system to the CREATE AV developers in support of release testing for the Kestrel and Helios applications. Beginning July 15, an independent queue with a compute pool of 2080 cores was dedicated to the developers of the Kestrel code for a period of 6 weeks. Following those 6 weeks, independent validation tests were conducted on the code release for another 4 weeks. Following the dedicated Kestrel testing, identical 6- and 4-week periods will be dedicated for the new Helios release. Helios testing will begin October 1 and run through December 15. This dedicated partition will utilize *Mana* resources under the Advanced Reservation Service.

MHPCC intends to continue to be an advocate for using our supercomputing resources in this manner. The development and release of these next-generation HPC codes require these kinds of dedicated resources to test and ensure the stability and performance requirements. Already, these codes are discovering key phenomena that have not been modeled in the past due to the ability of the code to scale to higher processor counts and the resulting increased physical resolution.

Continuous facility upgrades are ongoing to meet the growing demands of the MHPCC DSRC that include the expansion of the infrastructure of the Center and implementation of alternate green technologies (an R&D effort in advanced photovoltaic technologies). MHPCC recently did an extensive energy audit in collaboration with the DOE Federal Energy Management Program. As a result of this audit and implementation of many of the recommendations, the energy efficiency of MHPCC improved significantly. The Power Utilization Effectiveness (PUE) went from a baseline of 1.64 to 1.44. This was a decrease of approximately 12 percent of the total power used at the Center. This increase was done with a number of small incremental efficiency steps with no major facility renovations. The long-term goal is to increase energy efficiency with a resulting PUE of 1.35 or better.

In June, the Defense Security Service (DSS) recognized MHPCC as one of the nine recipients of the prestigious James S. Cogswell Outstanding Industrial Security Achievement Award. The award criteria focus on principles of industrial security excellence. The selection process for the Cogswell Award is rigorous. A DSS Industrial Security Representative may only nominate facilities that have at a minimum two consecutive superior industrial security review ratings and show a sustained degree of excellence and innovation in their overall security program management, implementation, and oversight.

At the forefront of HPC, the MHPCC DSRC has established itself as a leader in the DoD research and development community. Chartered to support a diverse base of DoD and other Government users, the MHPCC DSRC is facilitating the collaborations needed to solve tomorrow's complex computational problems today and dedicated to accelerate the development and transition of the ever-changing advanced defense technologies into superior warfighting capabilities.

Interns from Universities and Military Academies Receive Research Experience at the MHPCC DSRC

By Marie Greene, MHPCC DSRC Deputy Director

The MHPCC DSRC greatly expanded its summer intern program in 2010, held from May 16-August 25. More than 20 undergraduate students, graduate students, and interns were awarded research internships. Programs supported by the MHPCC DSRC included HPCMP Military Academies, the Air Force Research Laboratory Directed Energy Directorate (AFRL/DE) Scholars Program, AF TRIP (ROTC), AFRL, University of Hawaii, HPCMP Joint Equal Opportunities for Minorities (JEOM), Akamai-Center for Adaptive Optics (CfAO), and the Maui Economic Development Board. Universities represented were the U.S. Air Force Academy, U.S. Military Academy, Princeton University, University of Colorado, University of Kentucky, University of Wisconsin, University of Hawaii, and the Honolulu Business College.

Research projects were individually designed for each intern and their major area of study. HPC was employed in each research effort. Project examples included modeling and simulation, image processing, algorithm development, computational fluid dynamics, space weather, visualization tools, etc. Examples of some of the titles of the individual research projects included the following:

1. Interactive Live Event Toolkit (ILET) Image Vis Tool
2. Plasma for Use on Aerodynamic Surfaces
3. Algorithm Development for Satellite Collision Probability Analysis
4. CFD Analysis of F-15 and F-16 Grids

(Front row,left to right): Viola Silva (University of Hawaii), Skip Williams (AFRL/RDSMA), Michael McMinn (University of Kentucky), and Tim Fram (University of Wisconsin). (Middle row): Anna Cruickshank (USMA), Spenser McIntyre (USAFA), Marie Greene (AFRL/MHPCC), Daan Stevenson (University of Colorado), Herbert Harms (Princeton University), Tim Sinoga (University of Hawaii), and Gabriel Font-Rodriquez (USAFA Faculty). (Back row): Jonathan Plyler (USAFA), Gene Bal (MHPCC/UH), Joseph Dratz (AFRL/MHPCC), Tom Ainscough (USAFA), and Neil Grigsby (USAFA)

Maximizing Performance of the WRF Model on *Mana*

By Kevin P. Roe, MHPCC DSRC Senior Application Engineer

The Hawaiian Islands consist of dramatic terrain changes over short distances, resulting in a variety of microclimates in close proximity. To handle these challenging conditions, weather models must be run at fine vertical and horizontal resolutions to produce accurate forecasts. Computational demands require the WRF (Weather Research and Forecasting) model to be executed in parallel on MHPCC's *Mana* system, a PowerEdge M610 Linux cluster. This machine has 1152 compute nodes, each with two 2.8 GHz quad-core Intel® Nehalem processors and 24 GB RAM. Realizing maximum performance on *Mana* relied on the following:

- Determination of an optimal number of cores to use per socket
- Efficiency of an MPI-only implementation (versus hybrid MPI/OpenMP)
- Determination of an optimal set of parameters for adaptive time-stepping
- Maintaining the strict stability requirements necessary for Hawaii
- Effective choices for processor and memory affinity
- Parallel automation techniques for producing forecast imagery in a timely manner

The Haleakala Weather Center Web page can be seen at http://weather.mhpcc.edu.

All island domains down to individual counties can be selected and contain data for a 48-hour simulation in 1-hour increments for temperature, wind speed and direction, relative humidity, and rainfall.

After a detailed examination of the multiple techniques to maximize performance, the Haleakala Weather Center produces reliable high-resolution forecasts every night for the Hawaiian Islands in a timely fashion while minimizing its use of computational resources.

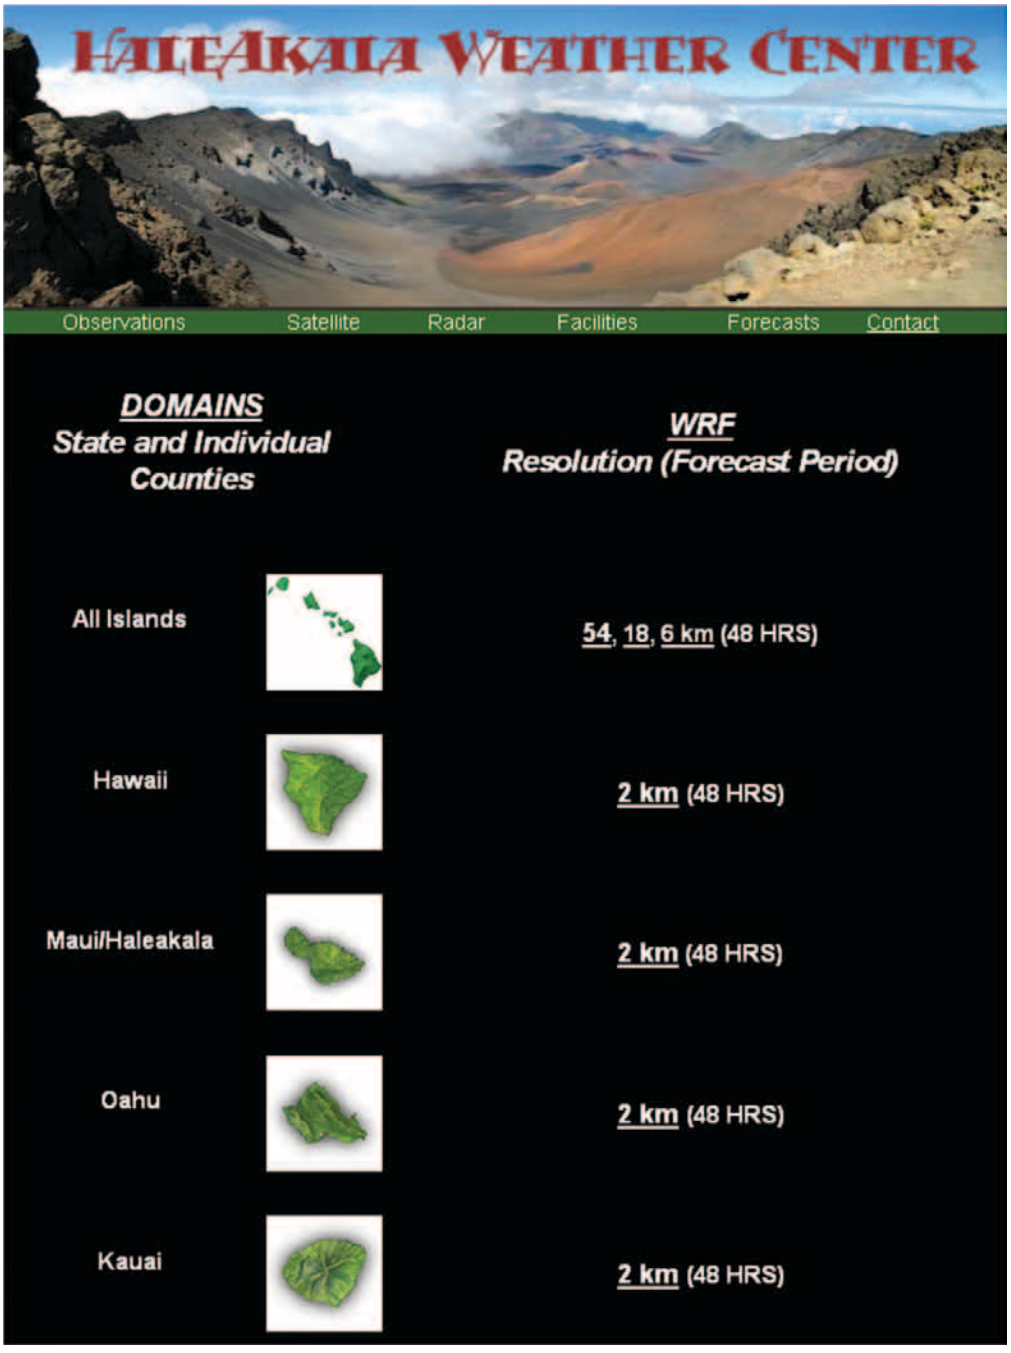

Figure 1. Weather Center Web Page

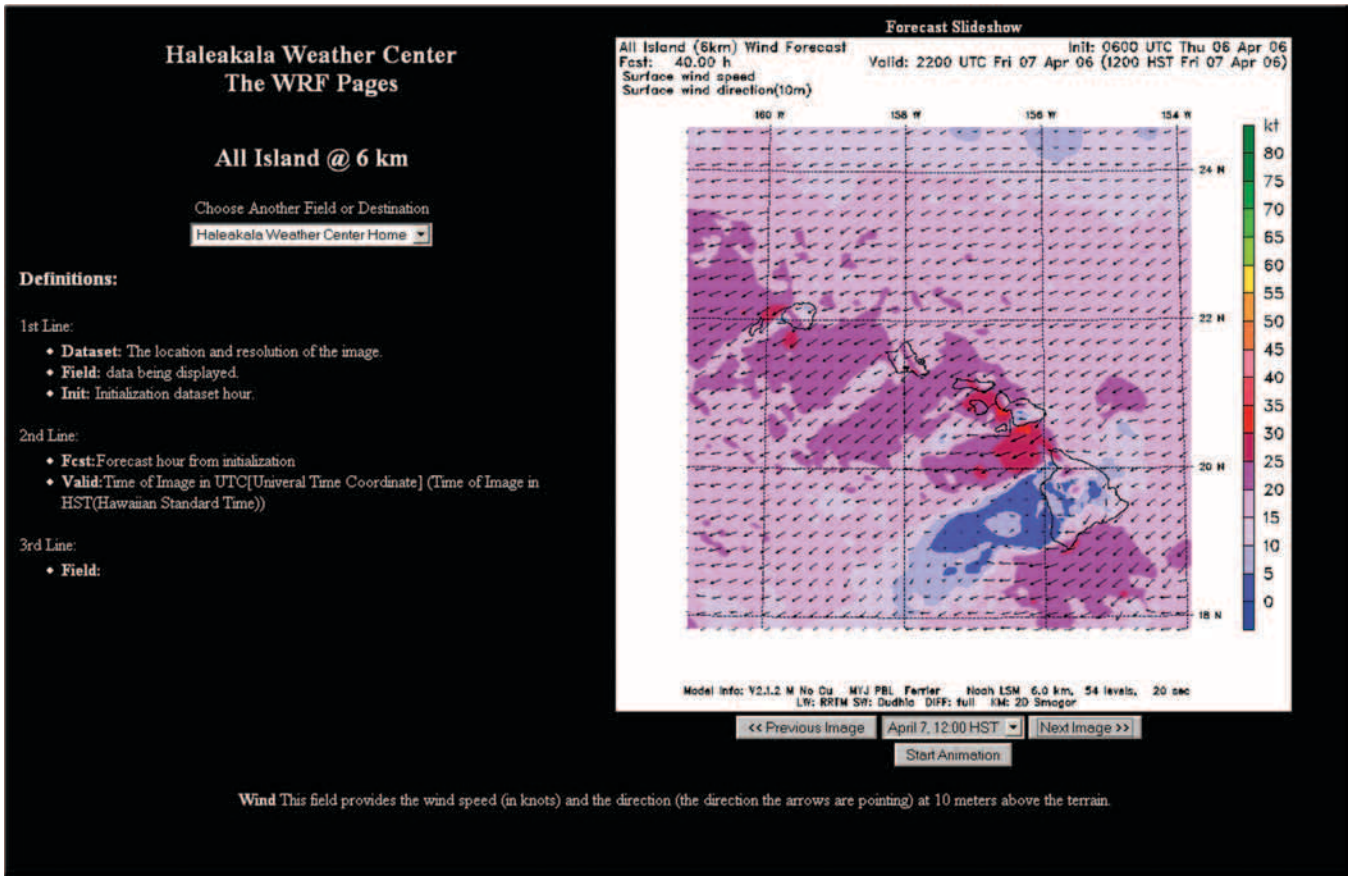

Figure 2. Wind Modeling

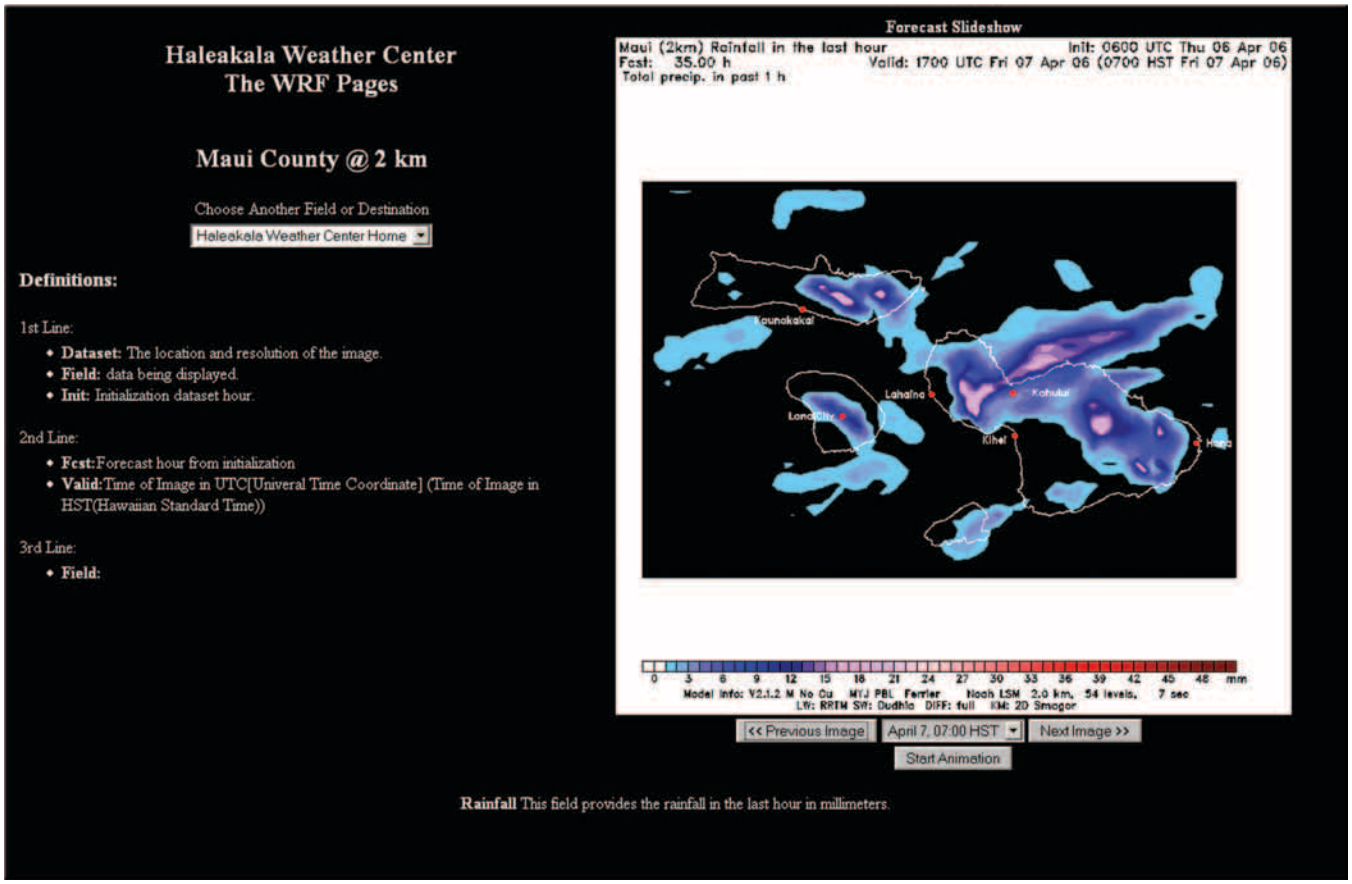

Figure 3. Precipitation Modeling

Navy DoD Supercomputing Resource Center

From the Director's Desk — Tom Dunn

Tom Dunn
Director, Navy DSRC

This edition of *HPC Insights* once again provides the opportunity to communicate recent accomplishments, upcoming events, and how the Navy DSRC is helping the HPCMP user community support the warfighter through supercomputing. We describe the implementation of an Automated Compliance Checking capability for the Baseline Configuration Web site, provide an update on the Navy DSRC high performance computing (HPC) systems, and offer a success story on how HPC has enabled the ocean modeling community to support a global community.

In her article "Automated Compliance Checking," Sylvia Seal summarizes the implementation of a capability that will provide a consistent, predictable way for identifying and managing baseline compliance across the entire Program. This project is a process improvement activity to ensure baseline compliance accuracy and performance.

We bid a fond farewell to the venerable POWER5+ HPC system *Babbage* and describe the current status of the Advance Reservation System in the article "What's New with Navy DSRC Systems" by Bryan Comstock.

On page 2 of this publication, Christine Cuicchi provides a success story detailing the chronology of how HPC has enabled the Navy numerical modeling community to provide value-added forecasts that enable the warfighter to characterize the environmental battlespace. She also discusses how HPC is currently being used to support the National Oceanic and Atmospheric Administration's (NOAA) Office of Response and Restoration (OR&R) with the Gulf of Mexico oil spill trajectory forecasts.

What's New with Navy DSRC Systems

By Bryan Comstock, Navy DSRC HPC Systems Analyst

Saying Goodbye to *Babbage*

Looking to the future, the Navy DSRC will soon be a part of the Technology Insertion (TI) acquisition process. As that process begins, we must also look to the past and see that one of our TI-06 systems, *Babbage*, will have reached the end of its life cycle at the close of this fiscal year.

Babbage has been a great asset to both the Navy DSRC and the HPCMP community as a whole. It became a part of the Navy DSRC, then the Naval Oceanographic Office Major Shared Resource Center, through the TI-06 acquisition process. Early in the existence of *Babbage*, it exhibited great stability, especially during the Capability Applications Project (CAP) periods. The Phase I CAP users on *Babbage* were able to utilize one million hours within a month, and the CAP II users clocked two million plus hours over a similar time span.

That performance simply set the stage for the remainder of *Babbage's* career. It routinely provided consistent, stable results over its lifetime as demonstrated by the sustained system performance (SSP) tests run each quarter by the HPCMP benchmarking team.

Babbage was also a leader in another area for the Program—Advance Reservation Service (ARS). Starting in Fiscal Year (FY) 2009, the ARS capability was transitioned from our IBM Power4+ system, *Kraken*, to *Babbage*. And just recently, we implemented an updated version of the ARS on *Babbage* and our other systems.

Although *Babbage* will be leaving us, its smaller sibling, *Pascal*, will continue to provide the same features and reliability to the HPCMP user community until it is replaced during the TI-11/12 acquisition process.

With that, we'd like to bid *Babbage* a fond farewell.

And Re-welcoming the Advance Reservation Service

As part of providing additional methods of computing to our HPCMP user community, the Navy DSRC, in conjunction with the Workload Management Team (WMT), has implemented the ARS on all of our unclassified HPC systems, including *Babbage*, *Davinci*, and *Einstein*.

Through the ARS, a user will be able to reserve a specific number of nodes for a specific block of time. This can be useful for time-sensitive computational work. For instance, if users need to interactively debug code and know they can be in front of a terminal for a given time range, the users could submit a reservation to the system via ARS and have sole access to the required nodes for that time period.

The table below shows the relevant ARS configurations for each of our systems as of Summer 2010:

Table of Advance Reservation System Availability per HPC Resource at Navy DSRC

System Name	Nodes	Cores
Babbage	40	640
Davinci	32	1024
Einstein	336	2688

The updated ARS has the potential to provide a great service to the user community, especially with 25 percent of each HPC system being available for reservation creations. We hope that the user community can take advantage of the ARS on our systems and look forward to providing any assistance necessary.

Learn About the Automated Compliance Checking Project

By Sylvia Seal, Information Assurance Lead, Lockheed Martin, Navy DSRC

The Automated Compliance Checking (ACC) project was a Special Emphasis project for Period 4 for the Next Generation Technical Services (NGTS) project. The ACC project delivered a tool to automate the reporting status maintained in the Baseline Compliance (BC) matrix.

The BC matrix reports policy status, in red or green, across the High Performance Computing Modernization Program (HPCMP) Centers. This project provided a process improvement activity to increase baseline compliance accuracy and performance.

The ACC project began several years ago by a development team from Instrumental, Inc. It was guided to leverage off the scripts, tests, and configurations developed by Instrumental to build the ACC utility. In addition to utilizing the Instrumental software, the ACC Charter stated two clear directives:

- Current list of Baseline Compliance Policies will be used
- Not all of policies will lend themselves to automated checking

At the time of development, the BC matrix supported 26 policies. The matrix was updated using a manual operation involving DoD Supercomputing Resource Center (DSRC) personnel. In determining the policies to automate, requirements were reviewed, including which policies could actually be automated. The ACC team worked closely with the BC team to get consensus and approval of policy automation.

The ACC project progressed using a two-phased approach. During Phase I, the team implemented the seven policies included with the initial Instrumental delivery; and Phase II implemented six of the remaining policies. Phase III followed successful implementation of Phases I and II with a bonus of two automated checking policies. The ACC utility also provides configuration files for those policies not a part of the automated checking scripts. These configuration files are transferred along with the compliance files for updates to the matrix.

One of the challenges confronted by the ACC team was the lack of an approved cross-center data transfer method. Several options were available, but the team was limited on providing an HPCMP solution that would meet the individual security requirements of each Center. After several weeks of discussions, the best solution of data transfer was determined to be e-mail. The solution is simple, easily understood, and is a common file transfer method. Even though e-mail as a requirement was a basic concept, the implementation was complex for the BC Web development team.

The ACC team was successful mainly due to the experience brought to the project by the team members. The Government Executive Agent was Jeff Graham, and the team included Perl developers, system administrators, BC team members, and a BC Web developer. Lockheed Martin leveraged additional Perl experience from a member of its Civil Programs Navy Engineering Support team at Stennis Space Center, Mississippi.

The ACC team consisted of the following:

Brent Siglar	AFRL, BC team
Greg Brewer	AFRL, BC Web team
Jonathan O'Reilly	ARL
Alec Bennett	ARSC
David Sanders	ERDC, BC Team
Lance Terada	MHPCC
Lee Whatley	Navy
Sylvia Seal	Navy
Bill Owens	Lockheed Martin, Civil Programs, Navy

The ACC utility now provides a means to alleviate the burden of manual compliance checking from DSRC personnel. The utility also provides for more accurate information, thereby creating meaningful information for the Centers as they manage compliance issues.

As of this writing, the following HPCMP systems are reporting compliance using the ACC utility:

AFRL:	*Hawk*
ARL:	*Mjm*
ARSC:	*Pingo, Midnight*
ERDC:	*Jade, Diamond, Sapphire*
MHPCC:	*Mana*
Navy:	*Einstein, Babbage, Davinci*

You can check their status by clicking the "Compliance Matrix" link at URL http://www.ccac.hpc.mil/consolidated/bc/policy.php.

Data Analysis and Assessment Center

DAAC Wins Department of Energy OASCR Award

By Randall Hand, Data Analysis and Assessment Center (DAAC) Scientific Visualization Lead

The Department of Energy (DOE) held its annual Scientific Discovery Through Advanced Computing (SciDAC) Conference on July 14. As part of the SciDAC Conference, a visualization night was held where participants could highlight their work. At the end of the event, the Office of Advanced Scientific Computing Research (OASCR) handed out 10 awards. The High Performance Computing Modernization Program (HPCMP) Data Analysis and Assessment Center (DAAC) won an award for work done with Dr. Eric Fahrenthold and sponsored by the Office of Naval Research and the National Aeronautics and Space Administration.

When asked about the impact of the HPCMP Office on his research, Dr. Fahrenthold had this to say:

> "The supercomputing support provided by Arctic Region Supercomputing Center (ARSC) and the high performance visualization support provided by DAAC were outstanding. User support is obviously a top priority for both organizations. ARSC provides highly reliable high performance computing systems in a very user friendly environment. DAAC converts a mountain of numerical data into a visualization with clear physical meaning, for critical comparison with experiment. Systems and graphics support of this kind has been of great benefit to my research program."

The winning video showcased the hybrid particle-finite element simulation code EXOS following physical impact experiments conducted at the University of Dayton and Los Alamos National Laboratory. The two cases show the impact of a 10-mm aluminum sphere on a 2.5-mm aluminum plate at 6.7 km/s, and 7.67 mm uranium alloy rod impacting a 6.4-mm plate at a 73.5-degree angle at 1.2 km/s. DAAC was able to process the raw simulation outputs and then create some detailed visuals complete with fracture effects, shadows, reflections, and pressure-visuals, all in a 4-minute video.

You can view the video on YouTube at *http://www.youtube.com/watch?v=XPEey2QGcgw*.

Impact of a uranium alloy rod on an aluminum plate, colored by pressure

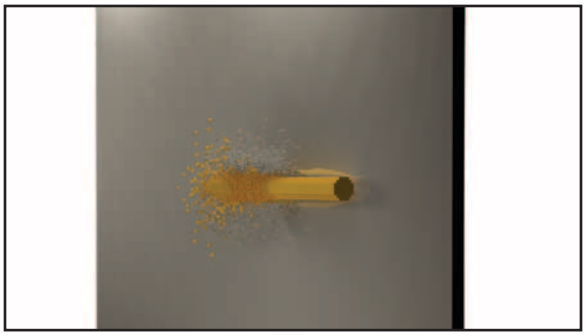

Impact of a uranium alloy rod on an aluminum plate

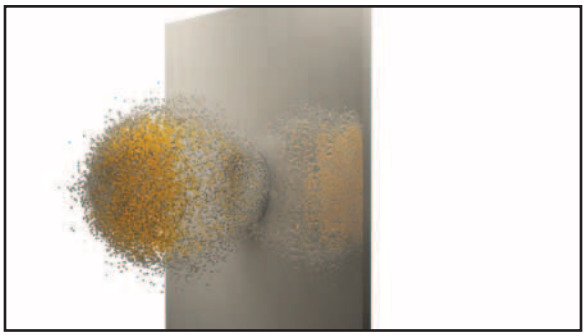

Impact of an aluminum sphere on a plate at 6.7km/s

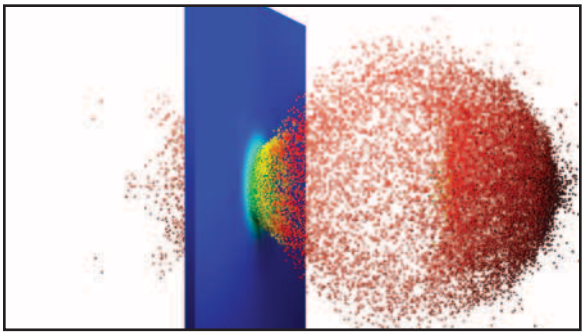

Impact of an aluminum sphere on a plate at 6.7km/s, colored by pressure

Announcements

DoD HPCMP Hero Awards for 2010

Presented by Cray Henry at the DoD HPCMP 2010 Users Group Conference
Chicago, Illinois, June 16, 2010

Technical Excellence
Kathy Borelli
MHPCC DSRC

Innovative Management
John Carter
AFRL DSRC

Long Term Sustained
Jeanie Osburn
NRL

Up & Coming Within the HPCMP
Peter Fulton
AFRL DSRC

pictures courtesy of Leah Glick, HPCMP

DoD High Performance Computing Modernization Program Users Group Conference 2011

Portland, Oregon, June 20-23, 2011, Hilton Portland & Executive Towers

DoD High Performance Computing Modernization Program

SOLVING THE HARD PROBLEMS

SUPERCOMPUTING FOR THE WARFIGHTER

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

2010 SC10 Edition

Evolving Technology