DOD HPC INSIGHTS

Fall 2012

A publication of the Department of Defense High Performance Computing Modernization Program

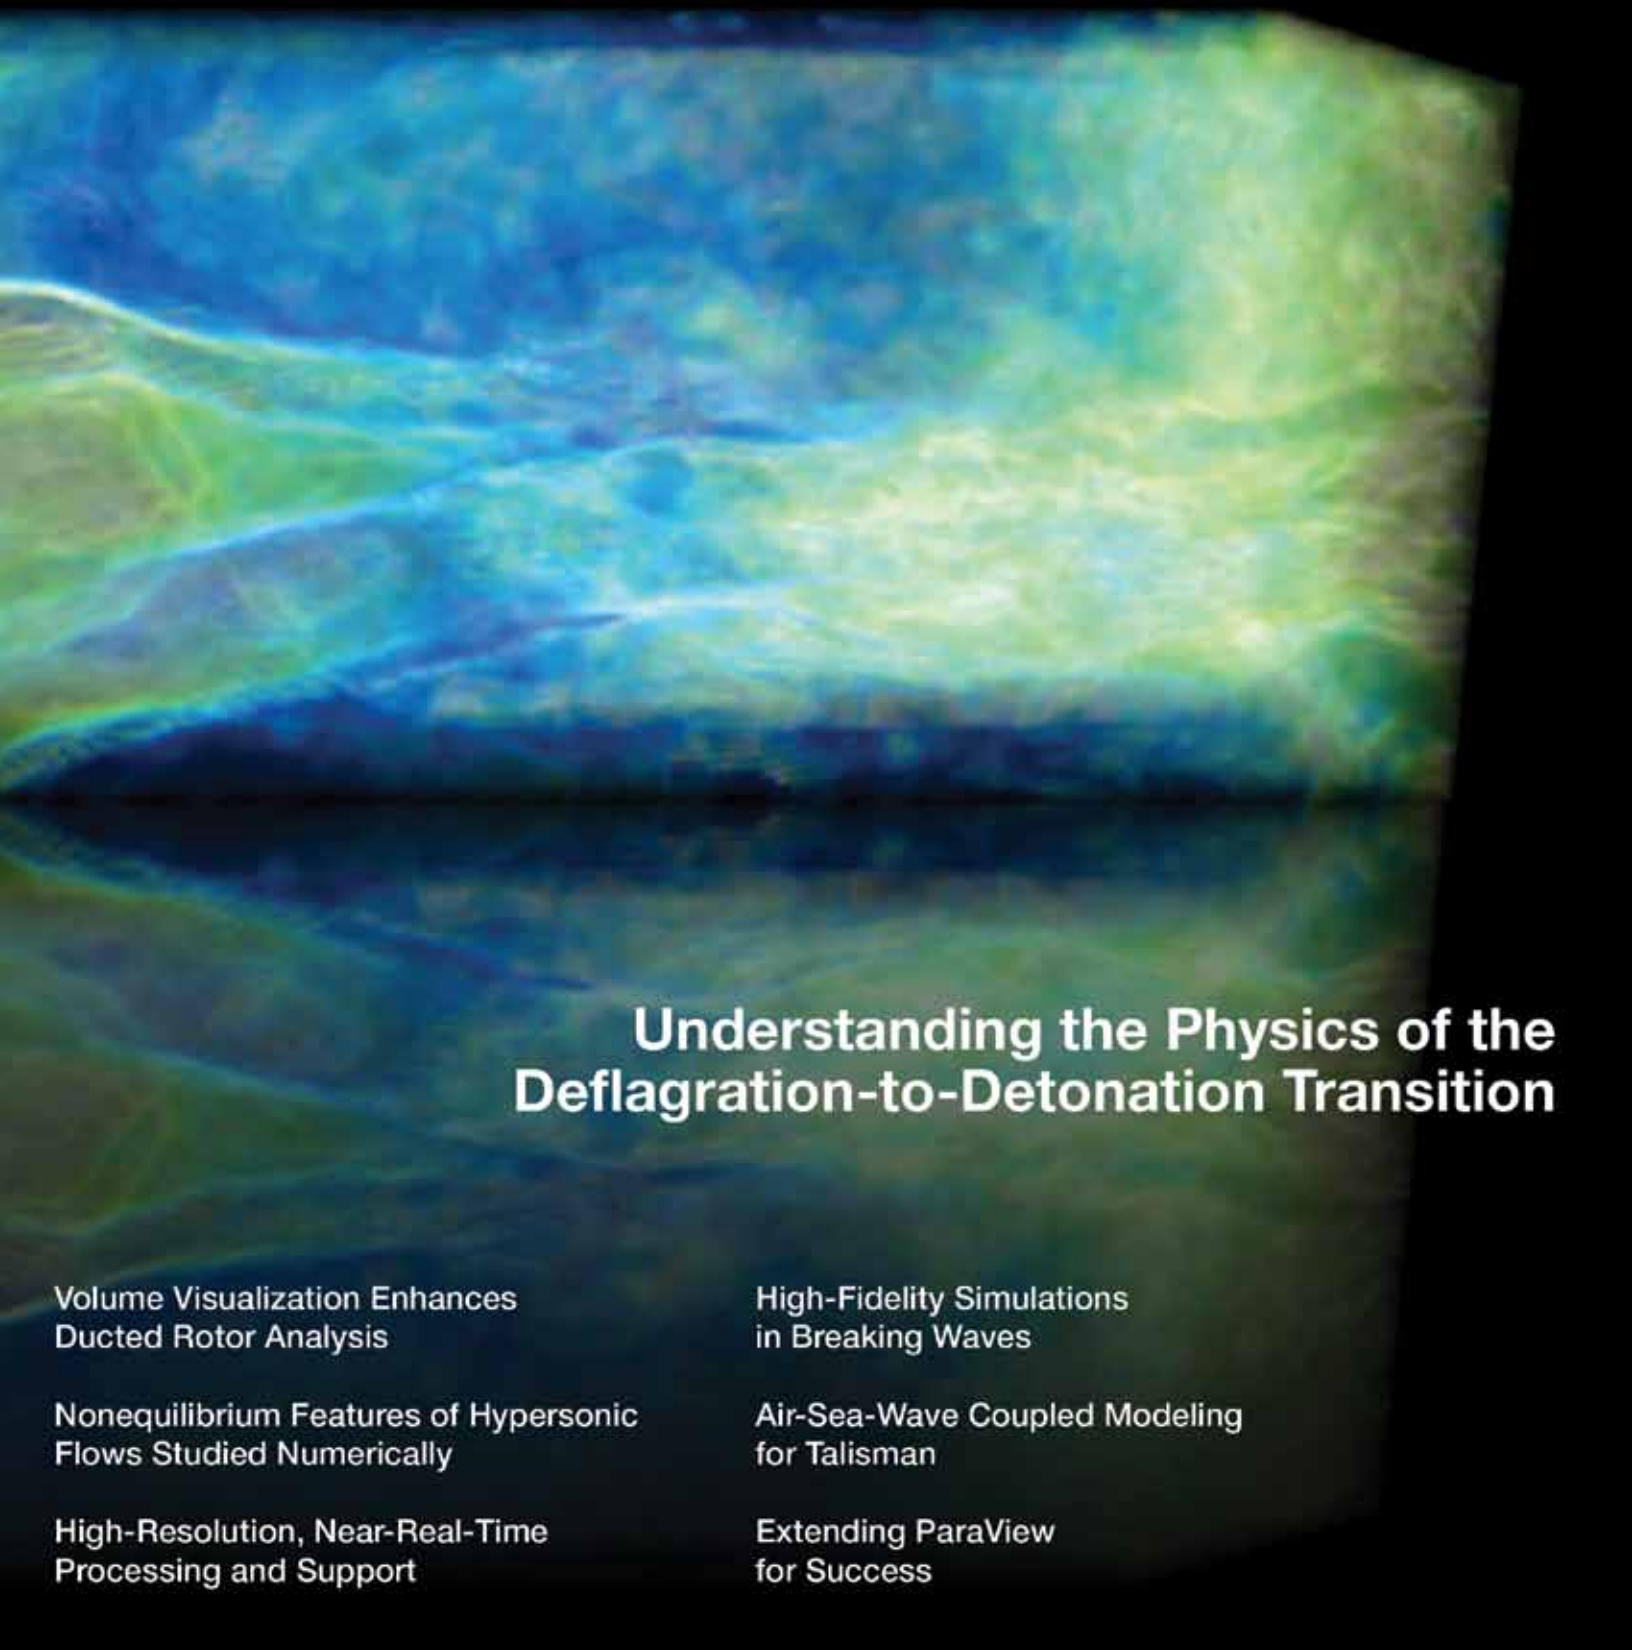

Understanding the Physics of the Deflagration-to-Detonation Transition

Volume Visualization Enhances Ducted Rotor Analysis

Nonequilibrium Features of Hypersonic Flows Studied Numerically

High-Resolution, Near-Real-Time Processing and Support

High-Fidelity Simulations in Breaking Waves

Air-Sea-Wave Coupled Modeling for Talisman

Extending ParaView for Success

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
Gregg Anderson
Chuck Abruzzino

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

ERDC DSRC, Vicksburg, MS
Rose J. Dykes

MHPCC DSRC, Maui, HI
Betty Duncan

Navy DSRC, Stennis Space Center, MS
Bryan Comstock
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
Data Analysis and Assessment Center

Contents

About the Cover: Volume rendering of a pressure wave reacting to a detonation of fuel. This is part of a more detailed study of the deflagration-to-detonation transition, which could lead to an increase in fuel efficiency.

The data were provided by Dr. Alexei Y. Poludnenko, Naval Research Laboratory, Washington, D.C., and the visualization done by the HPCMP Data Analysis and Assessment Center.

First Word

By Robert M. Hunter, Associate Director for HPC Centers, Department of Defense High Performance Computing Modernization Program Office, Vicksburg, Mississippi

Welcome to the fall 2012 edition of HPC Insights, and a special welcome to readers who may be visiting us for the first time at our High Performance Computing Modernization Program (HPCMP) booth at SC12. This issue of HPC Insights focuses on several success stories that highlight the return on HPC resource investment garnered by the Department of Defense (DoD) and its mission. We're proud of the systems and services we make available to the DoD Research, Development, Test, and Engineering (RDT&E) community and are always honored to share our successes and lessons learned with the larger HPC community.

The HPCMP is strengthening its strategic approach to providing the RDT&E and other defense communities with the most effective means of accessing and utilizing HPC resources for their increasingly varied requirements. While the Unix/Linux command line remains a formidable tool for the scientific computing professional, we recognize that domain scientists shouldn't always need to know all the command line switches for, say, the tar command. Toward that end, we're investing resources in new ways to enable scientists and engineers to interact with our HPC resources. We have a vibrant HPCMP portal effort in development that aims to provide our users with streamlined methods for conducting their scientific inquiry, while maintaining our systems' security posture in an unobtrusive manner and still offering sufficient flexibility and power to enable users to get the job done.

We've also made significant investment into Big Data and are currently in the testing phase of a new Storage Lifecycle Management system that provides, among other features, metadata tagging to enable efficient postprocessing data manipulation. Moreover, we realize that our users need on-the-fly information about their usage of our resources, and toward that end, we're developing a customizable user dashboard to allow deep insight into the status and statistics of user projects.

The HPCMP's strength comes from the talented staff in our Program. The Program's "Green Team," comprised of staff from our five DoD Supercomputing Resource Centers (DSRCs), has been recognized by the U.S. Army Corps of Engineers with the 2012 Green Dream Team Award, under the Chief of Engineers Awards of Excellence Program, for "recognizing an innovation or idea with clear potential to transform the Federal community's overall energy and environmental performance." Experienced in implementing energy conservation measures such as flywheel energy storage devices, uninterruptable power supplies, and step-down transformers, the Green Team makes recommendations for improvements at the DSRCs. Some of these improvements include the use of magnetic-levitation chiller compressors with economizers, automated control of chilled water temperature, and flow rate based on ambient temperature. The team's initial efforts improved one Center's overall power usage effectiveness rating from 2.0 to 1.4 with the expectation that the rating will continue to improve to below 1.3 during fiscal year 2013.

As 2013 nears, we are bringing approximately four petabytes of new HPC resources online at our DSRCs; details of these systems are highlighted by the Center Directors in this publication. These additional resources herald the largest single system the HPCMP has fielded: a 162,796-core Cray XE6 created from three upgraded systems that previously served the Program as separate systems in different Centers.

These new systems, in addition to our current systems, will provide 2765 billion core hours of computing in FY13. This incredible number of hours will provide DoD researchers, scientists, and engineers with the resources and opportunities to explore new and exciting science. I look forward to the new discoveries to come and the positive impact they will have in support of the DoD mission and its warfighter.

DOD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM

Understanding the Physics of the Deflagration-to-Detonation Transition

By Alexei Y. Poludnenko and Elaine S. Oran, Naval Research Laboratory, Washington, D.C.

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

On the morning of December 11, 2005, a violent explosion shook the London suburb of Hemel Hempstead. The incident occurred at the Hertfordshire Oil Storage Terminal, also known as the Buncefield complex, when the spill from one of the storage tanks led to the formation of a large fuel-vapor cloud. Subsequent ignition resulted in an exceptionally powerful explosion considered by many accounts to be the largest in Europe since the end of World War II. Thirty-five years earlier, almost to the day, on December 9, 1970, a similar explosion of an open-air vapor cloud, which formed as a result of the propane pipeline break, occurred in Port Hudson, Missouri. In both cases, the magnitude of the devastation was unprecedented for this type of industrial accidents.

Unlike these two examples, which are fortunately relatively rare, a different and incredibly more powerful kind of explosion is observed on a daily basis. These explosions, however, occur on cosmological distances and are an astronomical phenomenon known as the Type Ia supernovae (SN Ia). They result from the thermonuclear incineration of compact, degenerate, white dwarf stars. These stars are the end product of the evolution of normal stars such as the Sun. When formed in isolation, they lead a fairly peaceful existence slowly cooling after their nuclear fuel has been exhausted. However, in binary systems when a close-by stellar companion is present, white dwarfs can end their life in an extremely powerful and bright explosion. This fully disrupts the star in a matter of seconds and is capable of outshining the entire galaxy for the period of several days. Such brightness, along with the remarkable similarity of the majority of the exolosions, has turned in the last 15 years into an indispensible tool for measuring cosmological distances. This led to the discovery of the accelerating expansion of the universe and of the existence of dark energy.

These are the examples of some of the most powerful explosions on Earth (aside from nuclear ones) and in the universe. While they might seem different from each other, there is, in fact, a fundamental similarity of the key physical processes that power these events. In both cases, the explosion starts when a subsonic flame is ignited in the system. In an industrial explosion, this could occur due to some external factor, e.g., a spark. In the SN Ia, ignition would be caused by the increase in density and temperature in the stellar core as the white dwarf siphons matter from its stellar companion and grows in mass approaching the Chandrasekhar mass limit. Resulting energy release and the associated fluid expansion produce turbulent motions, which wrinkle and fold the flame, thus, accelerating burning significantly. An example of a turbulent thermonuclear flame inside the white dwarf during the later stages of the SN Ia explosion is shown in Figure 1. Figure 2 gives an example of a turbulent chemical flame formed in a methane-air mixture initially under atmospheric conditions.

Turbulent flame acceleration alone, however, is often not sufficient to explain the observed properties of these explosions, primarily their power. The possible missing piece of the puzzle is the deflagration-to-detonation transition (or DDT), in which a subsonic flame develops a detonation, a supersonic shock-driven reaction wave. The rate of energy release in a detonation is significantly higher resulting in a much more powerful explosion. It must be emphasized, however, that it is far from certain whether DDT can indeed occur in the interior of a star or in an open-air vapor cloud. The reason for that is our lack of understanding of the DDT process in such unconfined systems. In fact, it is not even clear whether truly unconfined DDT is possible at all. For SN Ia, for instance, this prompted a

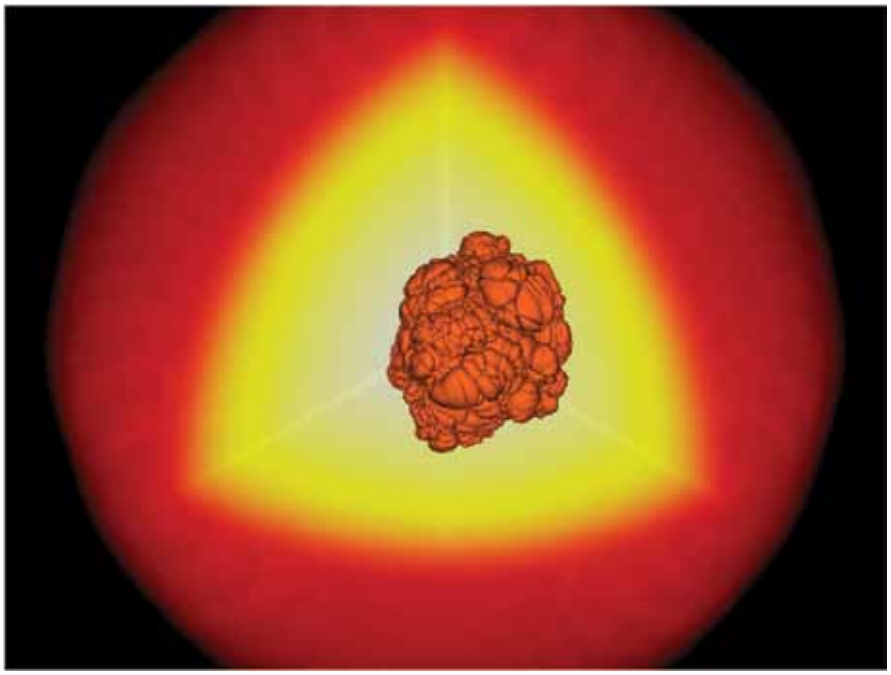

Figure 1. Complex structure of the turbulent thermonuclear flame propagating through the interior of a white dwarf star during a SN Ia explosion prior to the transition to a detonation. Computation is performed on an adaptive mesh with the highest resolution of ~600 m. Diameter of the star is ~4000 km. The turbulent flame dynamics is captured using a subgrid-scale model intended to reproduce the correct burning rate on the small unresolved scales.

search for alternative explosion models, which, however, are currently not as successful in explaining the observations as the model based on DDT. Therefore, elucidating the physics of DDT, as well as the conditions that can lead to it, is important for problems ranging from the safety of fuel storage and chemical processing facilities to the nature of the SN Ia phenomenon and of the enigmatic dark energy.

Furthermore, better insight into the process of detonation formation is crucial for the development of the next generation of propulsion systems. This primarily concerns detonation-based engines, which hold promise to provide up to a 25 percent increase in fuel efficiency.

Over the past 50 years, significant experimental and theoretical progress has been made in understanding the mechanism of DDT in confined systems, such as closed channels. In this case, however, there is a significant simplification. Burning in a closed space naturally leads to pressure increase and the formation of shock waves, which are driven by hot expanding burning products. These shock waves can eventually become strong enough to ignite a detonation. In contrast, in an unconfined system, there is no obvious way to form shock waves of sufficient strength. Furthermore, turbulence that can develop in the course of the explosion in the interior of a white dwarf or in the fuel-vapor cloud is subsonic and, thus, it cannot itself form shocks or accelerate the flame to supersonic speeds. The question then arises: can a highly subsonic flame interacting with the highly subsonic turbulence spontaneously develop a supersonic detonation without any assistance from the confining effect of external walls, boundaries, or obstacles?

Answering this question using numerical modeling presents a number of challenges. The main difficulty is associated with a broad dynamical range of scales involved in the problem. For instance, an open-air vapor cloud can reach hundreds of meters in size, while the characteristic burning scale (flame width) can be less than a millimeter (Figure 2). The dissipative (Kolmogorov) scale of the fast turbulence that develops in the explosion can be even smaller. This disparity of scales is further exacerbated in SN Ia. While the thermonuclear burning scale is similar to that in chemical flames (fractions of a millimeter to centimeters), the overall size of a star is thousands of kilometers (Figure 1). Thus, modeling the full system from the first principles, while resolving all relevant scales, is not feasible.

In addition to the multiscale nature of the problem, there is also a wide range of physical processes involved. These include complex nuclear or chemical reactions, thermal conduction and species diffusion, complex equation of state, radiation transport, etc. An attempt to include a detailed description of all these processes would typically make the cost of any three-dimensional (3D) computation prohibitive.

Therefore, in our approach, we attempted to find the simplest yet realistic setting that would exhibit a spontaneous transition to a detonation. Most importantly, this requires fully resolving the flame width in a 3D computation in order to avoid using any model descriptions of burning that could introduce significant uncertainties into the end result. The resolution requirement associated with this constrains the maximum practical physical dimensions of the computational domain. Since the acceleration of the flame through its interaction with turbulence is an essential part of the overall process, accurate generation of the turbulent flow field is also essential. Turbulence is typically stirred on large scales, which cannot be captured in the simulation. Therefore, in our calculations, homogeneous, isotropic, Kolmogorov-type turbulence is generated on the computational grid using a spectral method. Finally, we use a simplified, single-step, Arrhenius-type reaction kinetics calibrated to represent stoichiometric hydrogen-air and methane-air mixtures. In particular, this reaction mechanism produces realistic speeds and widths of both flames and detonations and is much more efficient computationally than more complex reaction networks.

The primary numerical tool for this work was the code *Athena-RFX*, developed in collaboration with Dr. T. Gardiner (Sandia National Laboratories). This is the reactive flow extension of the astrophysical magnetohydrodynamic code *Athena*, initially developed by Drs. T. Gardiner and J. Stone (Princeton University). The code is fixed grid, finite volume, higher order, Godunov-type and is massively parallel with excellent scalability demonstrated up to 50,000 CPU cores. It was successfully deployed both at the Air Force Research Laboratory DoD

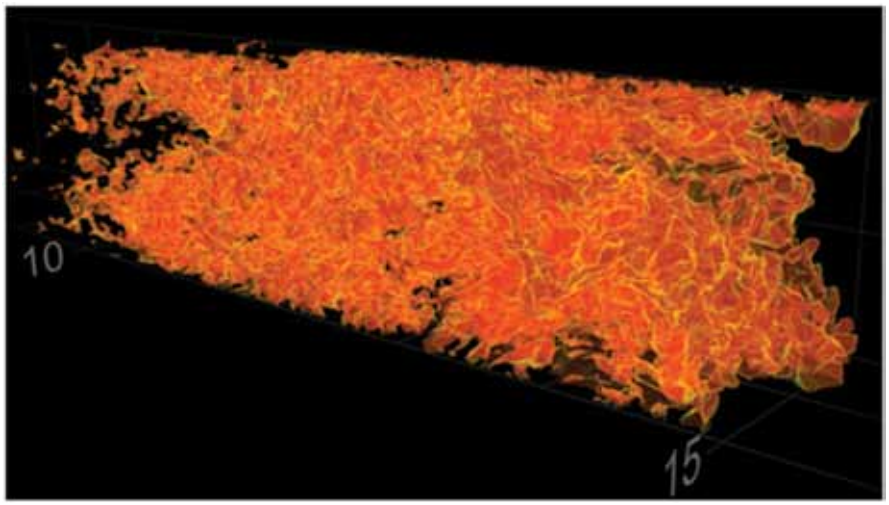

Figure 2. Complex structure of the turbulent chemical flame in a stoichiometric methane-air mixture. Uniform computational grid has dimensions 1.33 × 1.33 × 42.5 cm and resolution of ~50 μm. Shown is the isosurface of the fuel mass fraction corresponding to peak reaction rate. In this simulation, burning is fully resolved on all scales. (Image rendering by the HPCMP Data Analysis and Assessment Center)

Supercomputing Resource Center (DSRC) on *Raptor* and at the U.S. Army Engineer Research and Development Center (ERDC) DSRC on *Garnet*, where most of the results discussed here were obtained.

Since the conditions required for the onset of DDT were not known *a priori*, a survey of a large parameter space was required. This involved varying the type of the reactive mixture, which primarily affected the characteristic burning speed, as well as the system size and turbulence intensity. The largest calculations had the computational grid size in the range from $256 \times 256 \times 8096$ cells (0.5 billion cells) to $256 \times 256 \times 16384$ cells (1 billion cells). Since we use a fully compressible, explicit numerical solver and prior to the development of a detonation the flow in the system is highly subsonic, the overall number of time-steps per calculation was substantial, exceeding 100,000 or, equivalently, ~ 10^{14} cell-steps. The total CPU cost of each calculation ranged from 100,000 to 500,000 CPU hours. Overall, several dozen calculations were required to fully sample the parameter space. Thus, the availability of the 10,000-core class machines, such as *Raptor* and *Garnet*, was crucial for this type of study.

The key conclusion that emerged from these models was that subsonic turbulent flames are inherently susceptible to the formation of a detonation even in the absence of any confining factors, such as walls or boundaries.[1] Upon reaching a critical burning velocity, the flame develops a catastrophic runaway process illustrated in Figure 3, which shows DDT in a stoichiometric methane-air mixture interacting with fast turbulence. The characteristic turbulent velocity at the scale of the domain width is ~ 36 m/s, or ~ 10 percent of the sound speed. Shortly after ignition, once the turbulent flame becomes fully developed, pressure begins to rise throughout the volume of the flame accelerating it and forming the leading planar global shock. Shock waves, repeatedly generated within the flame, coalesce at the leading shock front amplifying it until the detonation is ignited (panels e-i).

[1] A.Y. Poludnenko, T.A. Gardiner, E.S. Oran, *Phys. Rev. Lett.* 107 (2011) 054501.

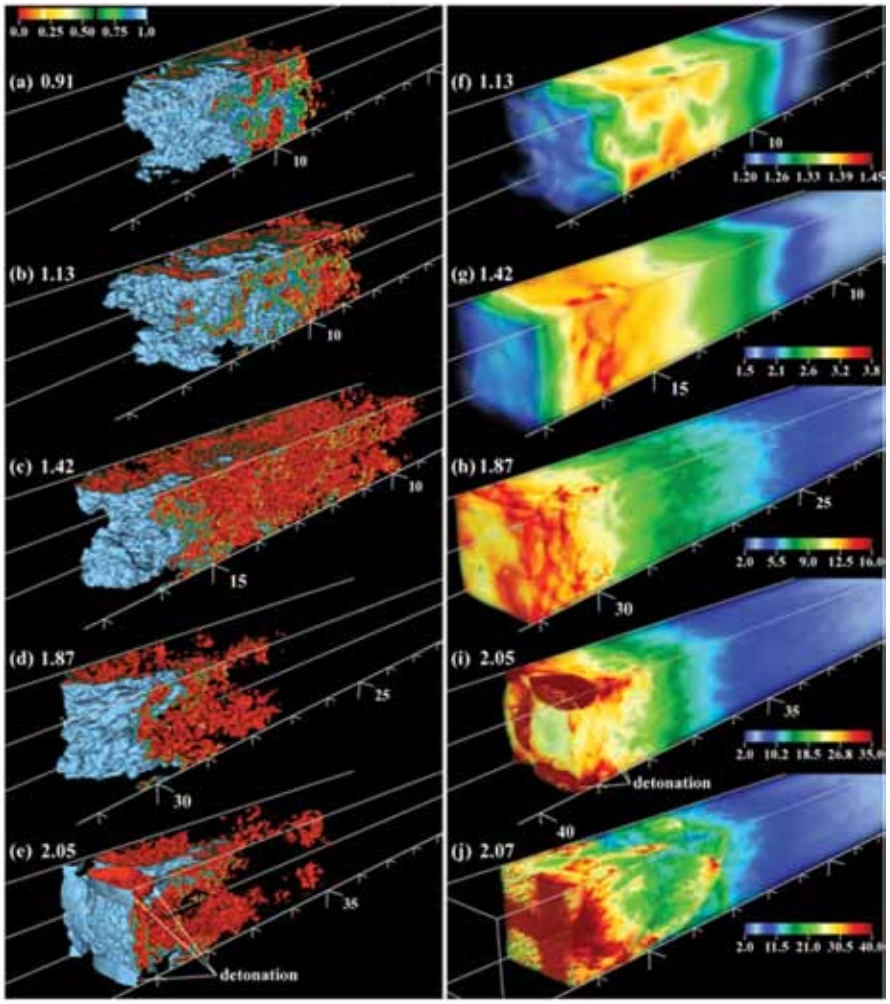

Figure 3. Structure of the turbulent flame and the corresponding pressure distribution during DDT in a stoichiometric methane-air mixture. (a) – (e): isovolume of the fuel mass fraction. (f) – (j): volume rendering of pressure normalized by the initial pressure in the domain (note a different colormap range in each panel). Horizontal axis scale gives the distance from the right boundary of the domain in cm. The time from the start of the simulation is indicated in each panel in units of the large-scale turbulent eddy turnover time. Note that panels (b)-(f), (c)-(g), (d)-(h), (e)-(i) correspond to the same time instants.

The underlying physical cause of the spontaneous pressure increase is the development of the supersonic flow of burning products downstream of the flame. In the reference frame comoving with the flame, fuel enters the flame with the speed equal to the flame-burning velocity. Products leave the flame with a much higher velocity due to the overall fluid expansion caused by heating. This means that at a certain subsonic flame speed, the product velocity will become equal to the speed of sound. At this point, any pressure increase as a result of burning cannot be propagated upstream by pressure waves, which will cause an overpressure to form within the flame volume. Such overpressure compresses and heats up the fuel, which, in turn, accelerates burning and further increases the outflow velocity of the burning products. This promotes pressure confinement and sets off the runaway process, which ultimately leads to a detonation.

The critical threshold, at which this process begins, is known as the Chapman-Jouguet (CJ) deflagration speed, the theoretical maximum speed for the steady flame propagation. Laminar flames, both chemical and thermonuclear, never reach such high speeds. Turbulent flames, in contrast, can become sufficiently fast. When they exceed this threshold, their steady-state propagation is indeed no longer possible.

These results are, by no means, the end of the story but rather a promising starting point. They show that spontaneous DDT is indeed possible in unconfined systems, such as open-air fuel-vapor clouds or the interior of a star. Furthermore, they specify a precise condition for the flame speed required for DDT to occur.

One of the major outstanding problems is that it is not known how to predict reliably the speed of a turbulent flame formed by the turbulent flow field, which may exist in a given practical situation. This is particularly true for highly unsteady flows encountered in the course of DDT. Turbulent flame models can be developed on the basis of *ab initio* 3D simulations, such as the ones discussed here. However, their validity has to be verified through calculations that use much larger ratios of the domain size to the characteristic burning scale along with more realistic turbulent flow fields. In particular, calculations must allow for the turbulence (possibly, inhomogeneous and anisotropic) to form self-consistently, rather than through the artificially imposed mechanisms. It is also crucial to relax some of the simplifications that made this study possible. This primarily concerns more detailed descriptions of the reaction kinetics, which can properly account for the complexity of real chemical and thermonuclear flames. Such next generation of calculations will not be possible without the next generation of petascale supercomputing platforms in the 100,000-core class. Scheduled deployment of such a cluster at the ERDC DSRC in early 2013 will allow further insights to be gained into the physics of the deflagration-to-detonation transition.

Acknowledgments

Financial support for this work provided by the Air Force Office of Scientific Research, the National Aeronautics and Space Administration, and the Office of Naval Research/Naval Research Laboratory 6.1 Base Program is gratefully acknowledged. Computing resources were provided by the DoD High Performance Computing Modernization Program. We also thank Randall Hand, Christopher Lewis, and Michael Wissmann from the High Performance Computing Modernization Program Data Analysis and Assessment Center for their assistance with data visualization.

High-Fidelity Simulations of Bubble, Droplet, and Spray Formation in Breaking Waves

By Zhaoyuan Wang, Jianming Yang, Frederick Stern, IIHR-Hydroscience & Engineering, University of Iowa, Iowa City, Iowa

HPC Resources: Cray XT5 (*Einstein*), Navy DSRC; SGI Altix ICE (*Diamond*), ERDC DSRC

Background and Objective

Plunging wave breaking is of great importance to ship hydrodynamics, including strong turbulence with large amounts of air entrainment, bubbles, droplets, jets, and spray. Previous experimental fluid dynamics (EFD) and computational fluid dynamics (CFD) studies on plunging wave breaking are mainly focused on the global structures of the wave breaking, such as wave elevation, jet, air cavity, and wave breaking processes. Moreover, due to the technical difficulties, the experimental measurements usually can only be done in the water region, and detailed description of the energetic wave breaking region is not available. With the development of the CFD technology, detailed studies of the two-phase region become possible.

The experimental study (Deane and Stokes 2002) shows that two distinct flow conditions that drive bubble creation in breaking waves are jet plunging and collapse of the cavity. The bubble sizes are from 2.0 mm down to at least 0.1 mm for the jet plunging period, and larger bubbles from 2.0 mm to ≥ 10.0 mm are created due to the collapse of the air cavity. In the recent study by Koo et al. (2012), plunging wave breaking over a submerged bump was simulated using the exact experimental flow conditions provided in Kang et al. (2012). The simulations were carried out on a 2D domain and focused on the overall plunging wave breaking process. Flows around a wedge-shaped bow were simulated by Wang et al. (2010) to study the wave breaking mechanism and main structures of ship bow waves.

The objective of the present study is to investigate the small-scale features of breaking waves such as the spray formation, air bubble and water droplet size, and distribution via numerical simulations. In order to resolve the small structures of the wave breaking at the scale of several micrometers, large-scale parallel computing (several billion grid points) is needed. In the present study, wave breakings around a wedge-shaped bow and over a submerged bump are simulated. The grid sizes used in these two simulations are 1.0 billion (1536 × 768 × 848) and 2.2 billion (1920 × 896 × 1280), respectively.

Approach

The challenges and difficulties for simulations using huge grid sizes (billions of grid points) lie in the data pre- and postprocessing, data storage, and transfer and computational speed, etc. Grid generation and grid file reading is easy for Cartesian grids since only one-dimensional arrays are required for grid coordinates. It is difficult to generate billion-point grids for complex geometries using curvilinear or unstructured grids with current mainstream grid generation software.

It is too time-consuming to transfer even a single data file (size can be tens of GB for billion-point grids) between a local PC and HPC, and it is impossible to store all the data files on a local PC, which is difficult and inconvenient for data visualization, analysis, and animation. Moreover, large grid size requires more memory

and processors, and smaller grid spacing needs a smaller time-step, which leads to even longer or unacceptable total computational time.

An orthogonal curvilinear grid solver, CFDShip-Iowa Version 6.2 (Wang et al. 2012a), is used for the computations. This solver was extended from the Cartesian grid solver (Version 6.1) (Yang and Stern 2009) for two-phase incompressible viscous flows recently developed at the University of Iowa. In this solver, both gas and liquid phases are considered for the strong interactions between two phases, such as spray dispersion and bubble entrainment. The level-set based ghost fluid method is adopted for sharp interface treatment, and the volume-of-fluid (VOF) method is used for the interface tracking.

In order to speed up the computations and improve the accuracy and efficiency, some enhanced technologies have been implemented such as a new VOF method (Wang et al. 2012b), semi-Lagrangian advection schemes for both the Navier-Stokes and VOF equations (Wang et al. 2012c), parallel grid/solution files reading/writing using MPI2-I/O (Yang et al. 2008), and optimized memory usage. Due to relatively simple geometries of the bump and wedge, 2D grids were generated first using Gridgen on a local desktop and translationally extruded to 3D grids on HPC. The water/air interface is extracted as PLY polygon file format for postprocessing.

Simulation Conditions and HPC Setup

The wedge geometry is similar to the large wedge model used by Waniewski et al. (2002). The side length of the wedge is L = 0.75 m, and the height of the wedge is H = 1 m. The half wedge angle is Θ = 26° and the flare angle φ = 0°. For the case considered here, the water depth is d = 0.0745 m and the upstream velocity is U = 2.5 m/s, the corresponding Reynolds number, $\mathrm{Re} = \rho U d / \mu$ = 1.64×10^5, and the Froude number, $Fr = U/\sqrt{gd}$ = 2.93. The geometry of the bump is the same as that in the EFD and CFD studies (Kang et al. 2012; Koo et al. 2012). The constant velocity imposed at the inlet boundary is u = 0.87 m/s for water and zero for air. The initial interface elevation is 0.2286 m, and a uniform velocity field is prescribed in the water domain with the air phase at rest.

The wedge flow case with 1.0 billion grid points was performed first on the Navy DoD Supercomputing Resource Center (DSRC) Cray XT5 (*Einstein*) from September 18, 2010 – October 19, 2010, and then moved to ERDC DSRC SGI Altix ICE (*Diamond*). The simulation ran smoothly on both platforms using 1024 cores and was also able to run on 512 cores (64 nodes) on *Einstein*. This was the largest case that had ever been conducted in the Ship Hydrodynamics group at the University of Iowa. In March 2012, the bump flow case with 2.2 billion grid points was performed on *Einstein* first using 1024 cores (128 nodes), but it ended up with "out of memory" (OOM) error messages. So the number of the cores was increased to 2048 (256 nodes), which was the maximum number of cores that could be used for the standard queue on *Einstein*. The code was terminated again with the same OOM error messages reported. This is because the amount of memory required by the multigrid Poisson solver has some dependence on the global grid size.

The total processes were then reduced to 1024 on 256 nodes (2048 cores), and the simulation ran without OOM errors anymore. However, this would waste much CPU time since only half of the cores on each node were used, and the CPU-hours would be charged for the total cores requested. The code might be able to run 4096 processes on 512 nodes (4096 cores) on *Einstein*, and then the CPU time would not be wasted. A request was sent to the Consolidated Customer Assistance Center (CCAC) for 4096 cores on *Einstein* on April 12, 2012. The request was approved with a "special queue" assigned. The code ran smoothly without any errors with 4096 processes on 512 nodes (4096 cores). The simulation was started on April 18, 2012, and terminated on April 27, 2012. The total size of all the data files generated in this simulation was about 51 TB, which was approximately 10 percent of the entire scratch file system disk space on *Einstein*.

Simulation Results

The wedge geometry and the computed wave profile are shown in Figure 1. Figure 2 shows the close-up view of the computed wave profile, which is similar to the experimental observation (Waniewski et al. 2002), such as the thin liquid sheet at the leading edge of the bow, overturning jet, jet plunging onto the free surface, and splashes at the wake. As the liquid sheet overturns, the sheet is stretched and fingered up, and some ''cylindrical drops'' then pinch off from the liquid sheet. When the detached drops impact the water surface, a spray region is created. In the experimental study conducted by Deane and Stokes (2002), the diameters of most observed bubbles due to the fragmentation process are greater than 2.0 mm. Mean drop size observed in experiments by Karion et al. (2004) is 2.3 mm. For small-size bubbles/droplets, surface tension force is dominant, and further fragmentation is difficult. In the present study, the grid spacing is 0.125 mm near the wedge and 1.0 mm in the plunging region. With the current grid, the minimum drop size is 0.8 mm near the wedge. The droplets

Figure 1. Wave breaking around a wedge-shaped bow

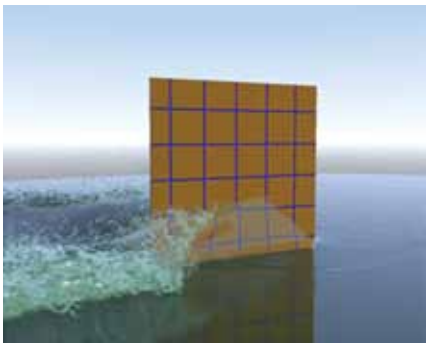

Figure 2. Wave profile of the wedge flow

and bubbles near the wedge can be effectively captured. Further grid refinement (3 to 4 billion grid points) is needed to increase the resolution in the wake region.

Figure 3 shows the overall plunging wave breaking downstream of the bump. The major events of the plunge wave breaking are demonstrated, i.e., maximum height, first plunge, oblique splash, and vertical jet. The computational results match the experiments well. Figure 4 shows the close-up of the bubbles and droplets in the breaking waves downstream of the bump, where detailed small interface structures are well demonstrated. With current grid resolution, the experimental observed mean droplet/bubble size can be effectively captured. Further grid refinement (up to 16 billion points) is needed in order to capture the minimum drop/bubble size observed in the experiments.

Conclusions and Future Work

Breaking waves around a wedge-shaped bow and over a submerged bump were numerically simulated using billions of grid points. An orthogonal curvilinear grid solver, CFDShip-Iowa Version 6.2, is used for the computations with enhanced technologies such as a new VOF method and semi-Lagrangian advection schemes for both the Navier-Stokes and VOF equations to improve the accuracy and efficiency. The computational results match the experimental observations well.

In future work, the grid resolution for the wedge flow will be increased (3 to 4 billion grid points) in order to effectively capture more droplets and bubbles away from the wedge. Further grid refinement (up to 16 billion points) is also needed for the bump flow case in order to capture the minimum drop/bubble size observed in the experiments. To accomplish these goals, much more HPC resources are needed, e.g., if 3 billion grid points are used for the wedge flow case, more than 2.8 million CPU-hours are required. This accounts for about 20 percent of our current total annual project allocation. Therefore, procurement of additional CPU-hours is required to carry out these large grid simulations.

Figure 3. Overview of the plunging wave breaking over a submerged bump

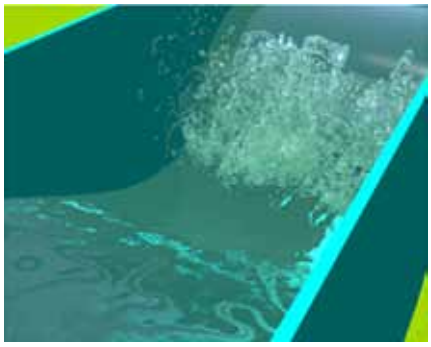

Figure 4. Close-up view of the wave breaking over a submerged bump

Acknowledgments

This research was sponsored by the Office of Naval Research under Grant N000141-01-00-1-7, under the administration of Dr. Patrick Purtell. The simulations were performed at the Department of Defense (DoD) Supercomputing Resource Centers (DSRCs) through the High Performance Computing Modernization Program (HPCMP). The authors would like to give special thanks to the Navy DSRC for the help on the "special queue," without which the simulations would not have been finished.

References

G.B. Deane, M.D. Stokes, Scale dependency of bubble creation mechanisms in breaking waves, *Nature*, Vol. 418, 2002, pp. 839–844.

D. Kang, S. Ghosh, G. Reins, B. Koo, Z. Wang, F. Stern, Impulsive plunging wave breaking downstream of a bump in a shallow water flume—Part I: Experimental observations, *J. Fluids Struct.*, Vol. 32, 2012, pp. 104-120.

B. Koo, Z. Wang, J. Yang, F. Stern, Impulsive plunging wave breaking downstream of a bump in a shallow water flume—Part II: Numerical simulations, *J. Fluids Struct*., Vol. 32, 2012, pp. 121-134.

Z. Wang, J. Yang, F. Stern, Numerical simulations of wave breakings around a wedge-shaped bow, 28th Symposium on Naval Hydrodynamics, September 12-17, 2010, Pasadena, California.

Z. Wang, J. Suh, J. Yang, F. Stern, Sharp Interface LES of Breaking Waves by an Interface Piercing Body in Orthogonal Curvilinear Coordinates, AIAA-2012-1111, 50th AIAA Aerospace Sciences Meeting and Exhibit, January 2012a, Nashville, Tennessee.

J. Yang, F. Stern, Sharp interface immersed-boundary/level-set method for wave-body interactions, *J. Comp. Phys.*, Vol. 228, 2009, pp. 6590-6616.

Z. Wang, J. Yang, F. Stern, A new volume-of-fluid method with a constructed distance function on general structured grids, *J. Comp. Phys*. Vol. 231, 2012b, pp. 3703-3722.

Z. Wang, J. Yang, F. Stern, A simple and conservative operator-splitting semi-Lagrangian volume-of-fluid advection scheme, *J. Comp. Phys.*, Vol. 231, 2012c, p. 4981-4992.

J. Yang, S. Bhushan, J. Suh, Z. Wang, B. Koo, N. Sakamoto, T. Xing, F. Stern, Large-Eddy Simulation of Ship Flows with Wall-Layer Models on Cartesian Grids, 27th Symposium on Naval Hydrodynamics, October 2008, Seoul, Korea.

T.A. Waniewski, C.E. Brennen, F. Raichlen, Bow wave dynamics, *J. Ship Res.*, Vol. 46, 2002, pp. 1–15.

A. Karion, T.C. Fu, T.W. Sur, J.R. Rice, D.C. Walker, D.A. Furey, Experiment to examine the effect of scale on a breaking bow wave, Carderock Division, Naval Surface Warfare Center, Hydromechanics Research and Development report, 2004, NSWCCD-50-TR-2004/060.

Nonequilibrium Features of Hypersonic Flows Studied Numerically

By Evgeny Titov, Pennsylvania State University, University Park, Pennsylvania; and Jonathan Burt and Eswar Josyula, Air Force Research Laboratory, Wright-Patterson Air Force Base, Ohio

Design optimization of hypersonic vehicles such as wave riders (a typical example is DARPA's HTV-2) is, in part, based on high-fidelity computational techniques including numerical solutions of the Navier-Stokes (N-S) equations. To make N-S solutions practical, two major issues need to be addressed. Firstly, physics-based, computational models for complex processes in the flow and at the boundary need to be developed; and secondly, a significant increase of computational efficiency needs to be obtained.

One of the challenges is due to a strong influence of the thermal nonequilibrium on the heat flux into the vehicles's thermal protection system at the wing leading edges. The second issue arises from typical design stage time limitations when many vehicle configurations and design materials need to be tested for an effective optimization. In this article, these issues are studied with an AFRL code [1] and with a Direct Simulation Monte Carlo (DSMC) solver HAP-DSMC [2]. We implemented and tested velocity slip boundary condition in AFRL code that accounts for local rarefaction effects. To utilize the modern GPU hardware, we attempted vectorizing parts of the code that had the largest computational load.

Maxwell formulated a velocity slip boundary condition in his paper [3] on viscous stresses arising in rarefied gases. The idea behind the model was to relate the velocity slip at the wall to the tangential shear stress and to the heat flux. The idea was based on two experimental observations of the time: a finite gas velocity at the wall in a rarefied gas and an appearance of the temperature creep force. This expression is widely used in computational fluid dynamics (CFD), where the local rarefaction effects are to be modeled in a predominantly continuum flow.

On the high performance computing front, graphics processing units (GPUs) demonstrate efficient high performance computing solutions of flow problems [4,5,6]. Typical implementations are based on the Nvidia CUDA libraries and C language extensions. More recently, PGI accelerator and PGI Cuda Fortran have gained popularity; and finally, hardware-independent standard OpenCL has emerged as a framework to unify GPU and multi-CPU parallel code development. In this article, we make use of the PGI accelerator.

Flow Solution with Velocity Slip Boundary Condition

To highlight the importance of accurate modeling of the flow at the boundary, we designed a hypersonic (Mach number equal to 7), blunted-wedge flow case of relatively high density such that the flow is not expected to be in the slip regime entirely, although local areas of slip are possible. Results for the case, whose flow solution is illustrated by the local Knudsen numbers (a relationship between the local molecular mean free path and the problem's characteristic size) contour plots in Figure 1, are further presented in Figures 2 and 3. Figure 1, in particular, illustrates a thin area of nonequilibrium in the flow along the shock wave that is formed in front of the body (as illustrated by elevated Knudsen number levels in the area.).

Figure 2 shows the local Knudsen number and mean free paths along the surface. It can be observed that the local Knudsen number has a maximum at a junction between the wedge and the leading-edge rounding not exceeding the critical value of around 0.1 above which the flow is typically considered to be in the rarefied regime where the slip boundary condition may fail.

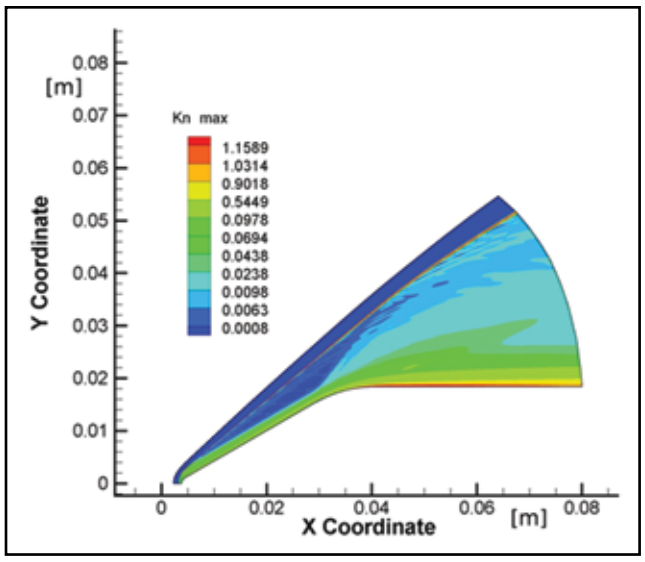

Figure 1

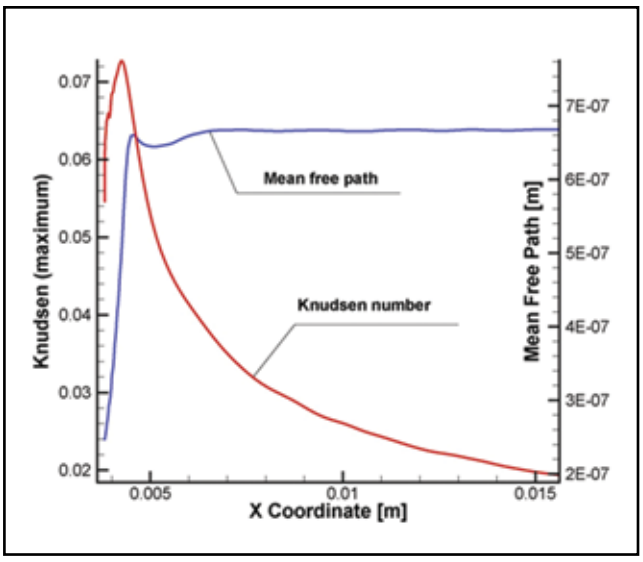

Figure 2

To verify the velocity slip values at the wall obtained with the N-S solver, a DSMC solution (a more computationally expensive but accurate solution) was obtained. The results are compared and presented in Figure 3. Although noisy, the DSMC solution shows a clear pattern, and the N-S results are close to the averaged DSMC.

Special attention, however, is required for the leading-edge area as presented in Figure 3, where two AFRL solutions are shown and compared with the DSMC results. The first AFRL solution, illustrated by a magenta line, was obtained with velocity gradients in the Maxwell's model that were numerically approximated by the first order of accuracy technique in the direction normal to the wall with the mesh step size in that direction on the order of 5.0×10^{-6}. The green line, on the other hand, shows a case in which the Knudsen layer thickness (a layer at the wall of a thickness of one molecular mean free path) that was approximately 1.0×10^{-7} was mesh-resolved. Although both of the solutions agree well with DSMC in the area along the flat surface, a significant difference is observed between these two AFRL solutions at the leading edge where the green line is much closer to DSMC results. A further study showed that velocity variation close to the wall in this area was not linear on the scale close to the molecular mean free paths, and, therefore, using derivative approximations with the mesh step larger than the mean free path would underestimate the resulting velocity slip values.

Mesh resolutions on the order of the local mean free paths are not typical for CFD analysis and may produce severe computational demands in high-density cases. It follows, however, from the provided observations, that such resolutions are needed especially in cases with local effects such as presented in Figure 3, where a large spike of velocity slip at the nose-tip area is confirmed by both DSMC and Maxwell's velocity slip model implemented in the AFRL N-S code. To meet such mesh requirements in cases of practical 3D application, a substantial increase of computational power is needed. The newly emerged alternative to the state of the practice of MPI and Open MP hardware and software capabilities also need to be explored.

AFRL Code GPU Implementation

The AFRL code profiling showed that the inviscid flux, source term, and the limiter subroutines accounted for approximately 80 percent of the computational time. These subroutines were targeted for GPU acceleration. After experimenting with various approaches, including the PGI accelerator, PGI Cuda Fortran, Nvidia CUDA framework, and OpenCL, the chosen approach was the PGI Accelerator.

After initial implementation and testing, it was found that because of external calls from the major program cycles, the implementation cannot be efficient. The original driver subroutine for the computational loop was then replaced by two subroutines, each designed to contain independent loops over I and J indexes of this 2D code. All the external function calls from the targeted subroutines were in-lined making the major loops applicable for GPU parallelization.

After testing of this alternative implementation with PGI accelerator, it appeared that more time was spent on data transfer between GPU and the main memory than on computations, which made the code inefficient. To avoid the data transfer overhead, the inviscid flux subroutines were redesigned to copy all the required data into GPU memory before the main cycles began and to transfer it back after the end of the computational loops. It was not possible to implement this approach fully because of GPU memory limitations and also because of partial interdependence of the resulting kernels on the data that are executed. The approach however worked to a degree to make the GPU implementation efficient. The PGI accelerated program that resulted from this approach was used to model a nozzle flow case presented in Figure 4 where the density gradient inside the nozzle is presented.

After the data movement overhead problem was resolved, most of the GPU time was spent on computations with

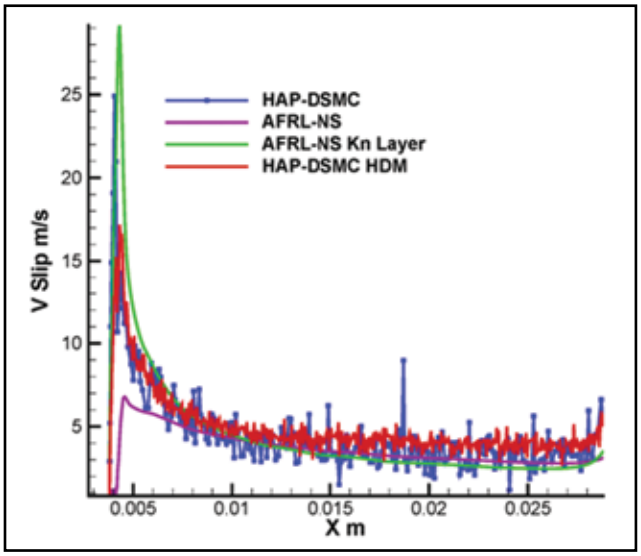

Figure 3

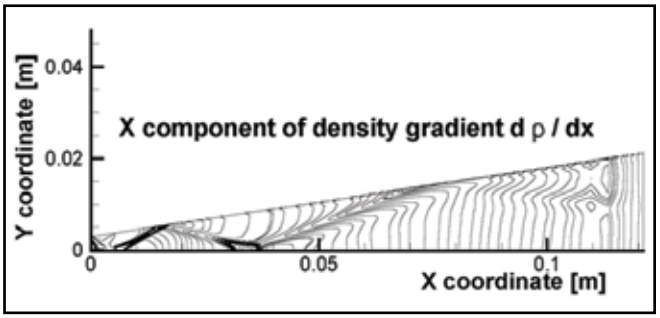

Figure 4

minimum time spent on data transfer. However, when applied to a typical problem considered here with a mesh size of up to 500 × 300 points, the implementation yielded a negative speedup (~1.5 times slower as compared with an advanced Intel Xeon CPU). The following reason was revealed: GPU implementation was correct, but because of the small size of the dataset, kernels did not have enough work to engage all the parallel processing elements. To verify this finding, a larger computational case, a long nozzle, was designed. The results are illustrated in Figure 4.

The solution presented in Figure 4 has an interesting feature that cannot be captured without high mesh resolution and solution-adaptation. A weak shock wave, which originates at the sharp corner of the nozzle throat, propagates through the nozzle and reflects several times from the nozzle walls and its centerline.

Two computational meshes were created for the case. A coarse mesh consisted of 202 × 52 points, and a larger one was 2500 × 300, where the larger mesh dimension is along the nozzle axis. The results for 10 time iterations are as follows: on a GPU Tesla C2070 at 1147 MHz accelerator, wall time was 17.2 sec.; on an Intel Xeon CPU 5160 @ 3.00 GHz, wall time was 30.4 sec. This constitutes an overall acceleration of 1.76 times. Since the accelerated programs accounted for 80 percent of the original CPU time, the actual acceleration of the subroutines that were modified was ~2.18. For the coarse mesh of 202 × 52 points, the speedup was negative.

Summary and Conclusions

A numerical study was conducted that consisted of the implementation and testing of slip boundary conditions in the AFRL perfect gas, continuum code. It was also found that for the slip velocity computed by the Maxwell model to be accurate, the cell size in the normal direction close to the wall should be close to the Knudsen layer thickness. It is important to note that smaller mesh resolutions may produce accurate results along flat surfaces, but fail to capture the physics at the strongly curved surface locations, as observed from comparisons with DSMC results. The nonlinearity of the velocity profile close to the wall results in significant differences of the numerically computed velocity gradient for different cell sizes. Velocity slip decreases with increase of flow density. However, if there is a need for accurate modeling of small-velocity slip values, the computational requirements are severe resulting from the necessity of resolving flow features with scales on the order of one mean free path.

The AFRL code was accelerated using the PGGPU technology. A PGI Fortran accelerator was used that required relatively small modifications of the code. However, to make the code efficient, the data flow in the code had to be reorganized in a way that allowed copying large portions of data into GPU memory before the main computational cycles start. A more fundamental redesign of several key subroutines using OpenCL standard was attempted and required rewriting portions of the code into C language and linking them with the rest of the Fortran code. Although promising, the technique did not show better acceleration compared with the PGI results. It was also found that GPU acceleration is efficient only for large computational meshes starting with resolutions on the order of 2500 × 250 mesh points. A good example of such cases, in this study, was a long nozzle case that required a dense mesh for the weak shock waves to be captured.

Future work consists of providing reliable heat flux predictions for design of leading edges of hypersonic vehicles such as HTV-2. This study finds that N-S solutions for such flows may not be accurate even for relatively high-pressure continuum flows. Sharp leading edges, on the other hand, provide substantial aerodynamic advantages at high supersonic and hypersonic speeds. The limitation to the leading-edge sharpness comes from maximum temperatures that thermal protection system materials can sustain (which are due to heat flux increases with reduction of leading-edge size). Since thermal conductivity of typical hypersonic leading-edge materials is low, radiative heat balance could be assumed to compute the local surface temperatures. This makes it possible to study the accuracy of N-S-based techniques in prediction of the leading-edge heat fluxes for practical applications. If it is shown that the N-S technique is not adequate for that purpose even in cases of relatively dense continuum flows, then more accurate kinetic approaches will be needed. This research has a potential of providing such an assessment.

References

[1] Josyula, E., "Computation of Hypersonic Flow Past Blunt Body for Nonequilibrium Weakly Ionized Air," *AIAA 93-2995, AIAA 24th Fluid Dynamics Conference July 6-9, Orlando, FL*, 1993.

[2] Burt, J. M., Josyula, E., and Boyd, I. D., "Novel Cartesian Implementation of the Direct Simulation Monte Carlo Method," *Journal of thermophysics and heat transfer,* Vol. 26, No. 2, 2012.

[3] Maxwell, J. C., "On Stresses in Rarefied Gases Arising from Inequalities of Temperature," *Phil. Trans.R. Soc. London,* Vol. 170, 1879.

[4] Hagen, T. R., Lie, K., and Natvig, J. R., "Solving the Euler equations on graphics processing units," In *International Conference on Computational Science, University of Reading,* 2006.

[5] Elsen, E., LeGresley, P., and Darve, E., "Large calculation of the flow over a hypersonic vehicle using a GPU," *Journal of Computational Physics,* Vol. 227, 2008.

[6] Corrigan, A., Camelli, F., Lohner, R., and Wallin, J., "Running unstructured grid-based CFD solvers on modern graphics hardware," *19th AIAA Computational Fluid Dynamics, San Antonio, Texas*, 2009.

Volume Visualization Enhances Ducted Rotor Analysis

By Rick Angelini, Army Research Laboratory (ARL), Data Analysis and Assessment Center (DAAC), Aberdeen Proving Ground, Maryland

HPC Resource: SGI Altix ICE 8200 (*Harold*), ARL DSRC

Background

Army Research Laboratory (ARL) researchers are working on understanding the thrust-generating mechanisms and characteristics of ducted rotor systems for powering small unmanned aerial vehicles (UAVs). Several characteristics make the ducted rotor system an attractive choice as a thrust-generating device for UAVs. Scientists at ARL's Vehicle Technology Directorate (VTD) performed comprehensive computational and experimental research on the ducted rotor systems for potential applications on small-scale UAVs and large-scale rotorcraft vehicle platforms. These simulations were run at the ARL DoD Supercomputing Resource Center (DSRC) on the SGI Altix system *Harold* using the computational fluid dynamics flow solver package CFD++. A series of simulations were used to study the effect of tip clearance and the shape of the duct on the performance of the rotor in hover. This work sheds light on how the augmented thrust from the duct varies with tip clearance for a rotor in hover. The researchers also looked at the thrust and torque coefficients and the figure of merit for different tip clearances.

Volume Visualization

Members of the Data Analysis and Assessment Center (DAAC) at ARL were anxious to introduce new volume-rendering techniques recently made available in EnSight version 10. Until the introduction of volume rendering, the visualization techniques in EnSight used points, lines, and surfaces to represent data, requiring volumetric data to be subsampled in some way. Traditionally, users would rely on clip planes, contour plots, isosurfaces, isovolumes, and particle traces to extract information from their computational mesh. These techniques allow the analyst to view a specific value or small range of values within their dataset, and in most cases, the results are then animated over time. Certain aspects of the calculation are visible using these techniques, and the results are effective for a concise range of values.

Volume visualization is a technique used in scientific visualization and computer-based imaging to represent an entire volume of data in such a way as to allow for complex physical phenomena to be clearly displayed. The resulting image can best be described as an opaque cloud of data that utilizes a color scale to represent values within the computational mesh. The ability to manipulate and animate this volume "cloud" over time provides insight into a large three-dimensional dataset and allows the viewer to see the results of millions (or billions) of computational values simultaneously. Volume visualization allows the analyst to view an entire range of values within a computational domain rather than calculating a specific scalar value for an isosurface or a range of values on a simple two-dimensional clip plane. By adjusting minimum and maximum data ranges and through the judicious use of color opacity, complex physical phenomena can become clearly visible within the rendered volume.

The volume-rendering algorithm implemented in EnSight uses a structured mesh defined by the user to resample the underlying computation mesh. This methodology is efficient, and EnSight is able to better leverage the computational resources with a structured volume algorithm; the volume calculated is slightly less precise than the original methodology, but significantly faster and more flexible in terms of end user interactive response. The user is able to create a bounding box around the

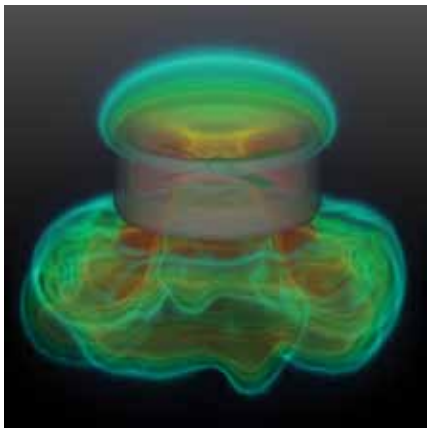

Volume rendering of the variable velocity with surface rendering of the duct and rotor visible

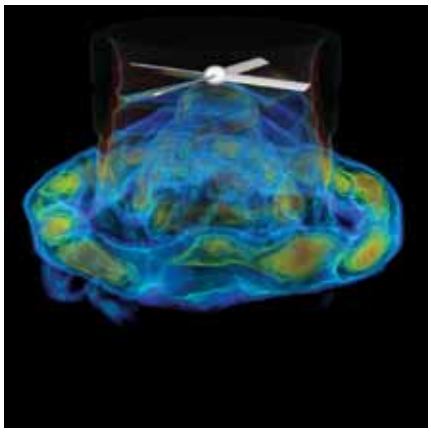

Volume rendering of the variable vorticity magnitude with a surface rendering of the rotor visible

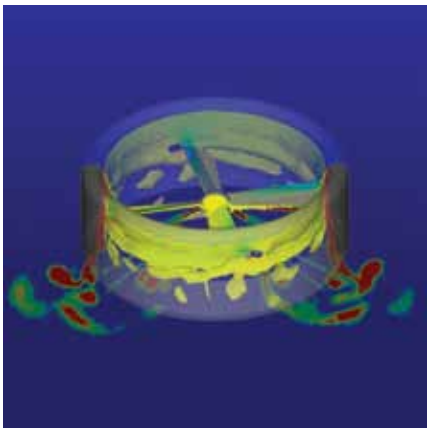

Traditional visualization of the ducted rotor dataset using an isosurface and clip plane created from vorticity magnitude variable

region of interest and create a three-dimensional structured grid within that bounding box. From this defined grid, the underlying computational data are subsampled into the structured grid, and the volume-rendering algorithm is able to efficiently calculate the volumetric image. Despite the reduced accuracy of the calculated volume resulting from the subsampling of the original computed data, there is often little to no perceivable difference in the rendering results.

Results

DAAC members were able to obtain access to this ducted rotor dataset in order to test and evaluate how volume visualization can be applied to a real-world dataset. DAAC staff successfully used both a traditional client-server connection between the desktop workstation and the ARL utility server as well as the HPCMP-provided secure remote desktop (SRD) application. When using a traditional client-server connection with EnSight, the volume renderer depends on the local graphics hardware on the desktop workstation to display the volume. When creating volume animations in batch, the EnSight volume renderer requires the use of the ARL Utility Server—specifically, the backend graphics nodes as the rendering algorithm still relies on the availability of graphics hardware. (The volume renderer cannot be used on traditional compute nodes that do not have graphics hardware).

The actual process of creating the volume rendering is beyond the scope of this article. More detail on the volume visualization process in EnSight will be presented in a forthcoming ARL technical report. However, it is worth noting that the process of extracting information from the volumetric data relies on manipulation of the color palette of the variable being represented in the volume. Refining the definition of the variable's color palette is accomplished by carefully selecting the appropriate min/max variable values and the opacity/intensity levels of the corresponding color values within that variable range. Sharp delineations in the rendered volume can be achieved by adjusting the variable palette using a "sawtooth" pattern. Creating abrupt changes in the opacity/intensity value at the correct value within the scalar variable range will result in a sharp peak in the representation of that data point in the volume image and seems to be an effective means of capturing complex dynamic changes within the rendered volume.

It was the intent of the DAAC team members to simply use this dataset for an internal training exercise to allow for understanding of the volume-rendering implementation in EnSight. However, the results that were obtained were of such interest that they were shared with the subject matter experts. The animated results of the volume visualization clearly showed the transient flow field as the rotor starts spinning and how the quiescent fluid is disturbed and moves as the rotor spins up and reaches a quasi-steady (periodic) state. Volume rendering was able to relay significantly more information than traditional surface-based visualization techniques such as an isosurface or clip plane. More than just "eye candy," the use of volume rendering allows the computational scientists to see more of the calculated result than was previously possible.

Notes

EnSight is a commercial, general-purpose scientific visualization package developed by Computational Engineering, Inc., that has been actively used on DoD high performance computing resources since the inception of the HPCMP. Numerous articles have appeared in previous editions of this publication regarding the use of EnSight, most recently the Spring 2012 Edition article "Client-Server HPC Job Launching." Additional information on volume rendering can be found in the in-line user documentation under the "Help" menu in the EnSight interface. General descriptions of EnSight's implementation of volume rendering were extracted from the in-line documentation for this article. Additionally, there are numerous on-line volume visualization tutorials on the EnSight website at http://www.ensight.com/Volume-Rendering.

The author thanks Drs. Surya Dinavahi and Rajneesh Singh, both of ARL, for providing the ducted rotor dataset and for guidance and insight of the results obtained from application of the volume-rendering technique.

Air-Sea-Wave Coupled Modeling for Talisman Sabre 2011

Dr. Travis A. Smith, Oceanographer, Naval Research Laboratory, Stennis Space Center, Mississippi

HPC Resources: Cray XT5 (*Einstein*) and IBM (*Davinci*), Navy DSRC

The Naval Research Laboratory (NRL) and the Naval Oceanographic Office (NAVOCEANO or NAVO) implemented the Coupled Ocean/Atmosphere Mesoscale Prediction System (COAMPS) version 5.0 at the Navy DoD Supercomputing Resource Center (DSRC). COAMPS, which is a joint effort between the NRL Oceanography Division at Stennis Space Center, Mississippi, and the NRL Marine Meteorology Division in Monterey, California, utilized the DSRC's high-performance computing (HPC) resources in support of the Talisman Sabre naval exercise in July 2011 (TS11). Talisman Sabre is a biennial joint Australia-United States military exercise that takes place in coastal waters along the northern and central Australia coast. The purpose of these exercises focuses on crisis-action planning and contingency response with an aim of enhancing the military capabilities of both nations.

NRL and NAVO supported TS11 by providing daily real-time COAMPS

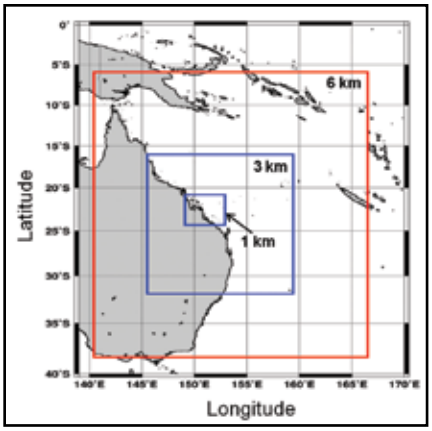

Figure 1. COAMPS nests for TS11. Inner atmospheric nest and SWAN nest (red, 6 km resolution) and two NCOM nests (blue, 3 and 1 km resolution) are shown

netCDF products and graphics through the Navy Portal, such as 48-hour forecasts of ocean currents, sea-surface temperatures (SST), salinity, and sound speed. COAMPS version 5.0 couples the COAMPS version 3.0 atmospheric model with the Navy Coastal Ocean Model (NCOM) and the Simulating Waves Nearshore (SWAN) model through the Earth System Modeling Framework (ESMF). The fully coupled air-ocean-wave COAMPS system was implemented and demonstrated for the first time on the DSRC's Cray XT5 (*Einstein*) platform in May 2011. The TS11 COAMPS simulation consisted of 18 and 6 km resolution atmospheric nests, 3 and 1 km resolution NCOM nests, and a 6 km SWAN nest centered on Shoalwater Bay that is located on the east-central Australian coast (Figure 1). All required input data (global model and observational data) were made available at the DSRC to produce initial and boundary conditions for both the atmospheric and oceanic nests. Additionally, WaveWatch III was run simultaneously on the DSRC's IBM (*Davinci*) platform to create boundary conditions for SWAN. The DSRC provided a special queue for timely model runs while maximizing

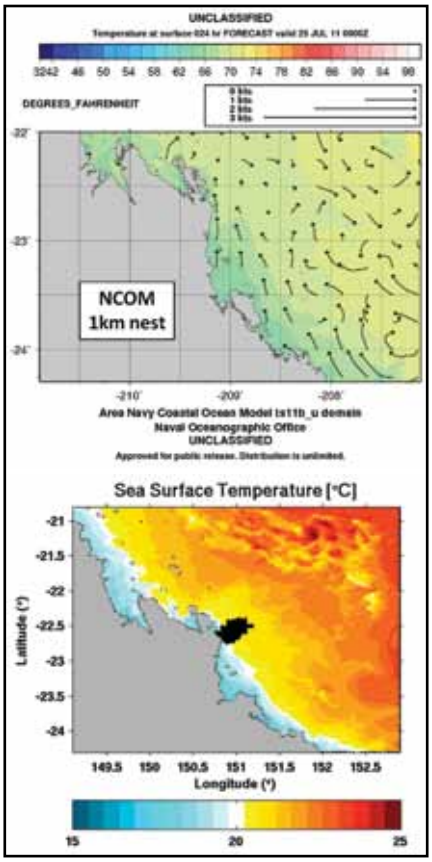

Figure 2. NCOM 1 km nest surface current (vectors) and SST plot (below) created by NAVO for TS11 guidance

resources as much as possible during the operational period. Each daily COAMPS forecast run utilized 160 processors, enabling COAMPS to produce a 48-hour forecast within 5 hours. Spin-up of the TS11 COAMPS nests commenced on June 1 and ended on June 30, 2011, while forecasts were generated for TS11 from July 1-29, 2011.

The NCOM 1 km nest output generated on *Einstein* (Figure 2) during TS11 was compared with two Australian SLOCUM shallow-water gliders that were deployed on July 11-17, 2011, off Freshwater Beach within Shoalwater Bay. The gliders collected conductivity, temperature, and depth (CTD) measurements. Errors in both temperature and salinity to 60 m depth were low

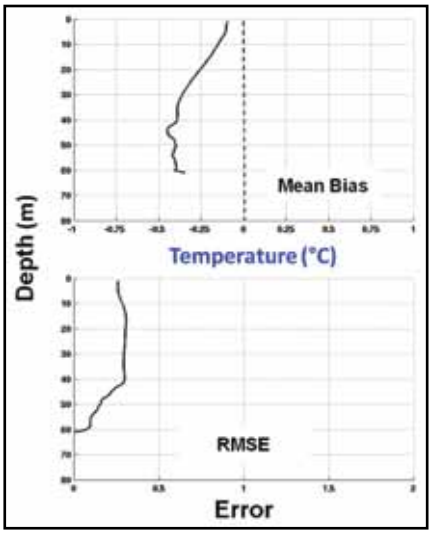

Figure 3. NCOM 1 km nest SSTs and nearshore path of Australian SLOCUM gliders (black crosses). Comparisons of ocean temperatures and glider data show that the RMSE and mean bias are less than 0.5°C

despite the gliders being in shallow water (Figure 3). The ocean temperature mean bias and RMSE were lower than 0.5°C, while the salinity mean bias and RMSE were no greater than 0.25 psu throughout the entire water column. Although sparse in coverage, several wave buoys off the eastern coast of Australia provided data, such as SST, and comparisons were similar. Additionally, NCOM, SWAN, and atmospheric outputs were used by NRL for Delft3D support during TS11.

The Navy DSRC's HPC resources are an invaluable asset in real-time military operational exercises, and efforts by both NRL and NAVO led to a successful partnership with the Australian Department of Defence. The collaboration may lead to additional NRL and NAVO support in future Talisman Sabre exercises.

The author acknowledges Scott Smith (NRL) and Pete Spence (QinetiQ North America) for glider comparisons (Figure 3) and Jim Dykes (NRL) for WaveWatch III products.

Extending ParaView for Success

By Joel P. Martin, Army Research Laboratory, Data Analysis and Assessment Center; and Brian Vonk, Army Research Laboratory

ParaView, VisIt, and EnSight are all visualization workhorses used extensively within the DoD High Performance Computing Modernization Program. Users throughout the Program are successfully using these applications to visualize and postprocess their data. Unfortunately, these packages are not always capable of doing everything researchers need to obtain their results.

Recently, Army Research Laboratory computational researcher Brian Vonk came to the Data Analysis and Assessment Center (DAAC) with such a visualization issue. His CTH output had variables divided by material type and no way to look at all similar values at the same time. In other words, the data had related values that spanned multiple variables. The images in Figure 1 are from a contrived dataset that shows a similar problem to what was being reported by the customer. In this example, a sphere projectile impacts a disk, and the damage variable is color mapped to the surface. Figure 1a shows the damage to the disk and Figure 1b shows the damage to the sphere. Finally, Figure 1c shows the combined results for which the customer was looking.

> "This tool has helped us better understand the material interactions involved in high-velocity armor design."
> – Brian Vonk

None of ParaView's built-in filters performed the necessary logic to combine the variables. Also, the ParaView calculator did not contain a maximum operator, which would have allowed the variables to be combined.

Fortunately, ParaView can be extended through the use of Python and C++. Since the desired results could be achieved through a relatively simple calculation (taking the maximum of two variables) and a short turnaround time was needed, Python was chosen for the task. If this were for a larger problem or a bigger audience, C++ might have been picked. C++ plugins in ParaView are faster and more features are available, but it must be compiled for each architecture on which it is run.

The DAAC produced a ParaView Python filter that can be applied to the dataset. The filter is loaded from an XML file that contains the Python code and a simple user interface. The user can change the input variables by typing them into the dialog boxes provided (not being able to have a drop-down list of variables is one of the downsides to using Python instead of C++). The resulting filter produced the results in Figure 2. The images show the sphere impacting the disk and the damage variable mapped over the entire dataset, not just one material. Future work may include wrapping the code in a Qt GUI and automatically combining all similar variables.

The DAAC is available for advice and support on pre- and postprocessing issues. We support many packages and data types. Please contact us at support@visualization.hpc.mil if you have an issue with which we can assist you.

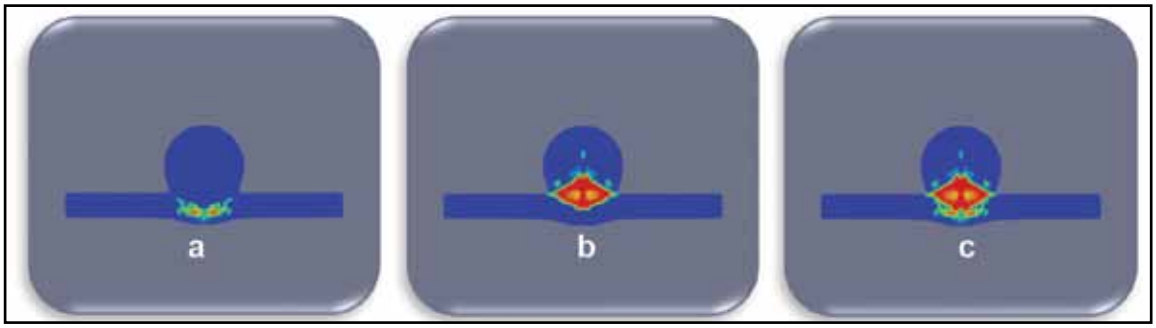

Figure 1. Damage to each material can be viewed only one at a time (a, b). Desired image is on the right (c)

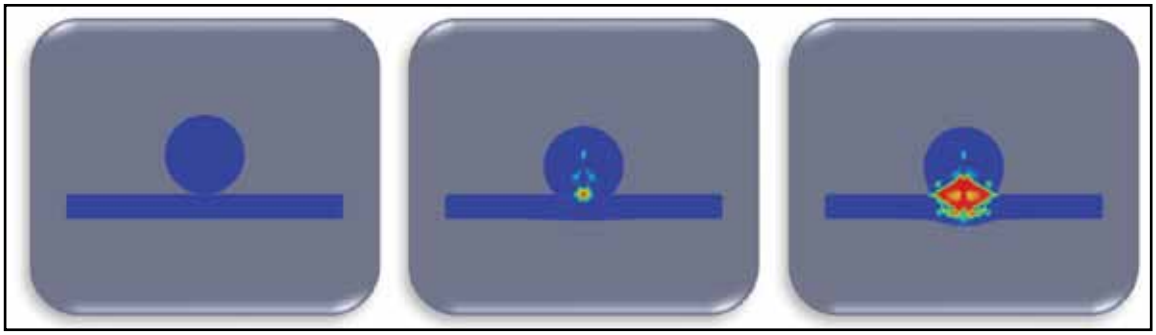

Figure 2. Frames from the completed animation showing a sphere impacting a disk. Results calculated using CTH

High-Resolution, Near-Real-Time Image Processing and Support

By Dr. Chris Sabol, Director, High Performance Computing Software Applications Institute for Space Situational Awareness, Air Force Maui Optical and Supercomputing Site, Maui, Hawaii

A team led by the Air Force Research Laboratory Directed Energy Directorate (AFRL)/RD has performed image enhancement software applications engineering development in support of Maui Space Surveillance System (MSSS) modernization at the Air Force Maui Optical and Supercomputing (AMOS) site (Figure 1). The AMOS site also includes and is provided important support by the Maui High Performance Computing Center DoD Supercomputing Resource Center (MHPCC DSRC). The work and efforts involved also include the AFRL modified use of an FY10 Dedicated High Performance Investment (DHPI) cluster award—and planned use for an FY12 DHPI award.

Key applications software and utilities are utilized at the AMOS site to process atmospherically blurred raw images of space objects obtained from ground-based, electro-optical telescope sensors, such as MSSS, and perform additional image processing to obtain high-resolution imagery information. The AMOS site is helping to increase the Nation's Space Situational Awareness (SSA) capabilities using advanced electro-optical-based collection and analysis technologies, and it is a national focal point for ground-based SSA, combining R&D with an operational mission.

AMOS is continuing to provide enhanced-image quality, utilizing the Physically Constrained Iterative Deconvolution (PCID) image processing and user-interface software packages that were developed by AFRL. PCID is an implementation of a multiple frame blind deconvolution technique. The PCID algorithm uses little or no *a priori* information about the space objects that are being imaged and incorporates features that increase the probability and speed of finding the global minimum, process non-idealities such as spatially cropped images due to jitter and/or object size relative to the detector size, and generate many of the PCID inputs automatically. The forward model used in the PCID algorithm is the standard linear imaging model. Multiple frames of data are collected by sensors using telescopes such as at MSSS, while an object passes overhead. Each collected frame is blurred by the earth's atmosphere. The name given to this blurring is point spread function (PSF). Frames are grouped together to form an ensemble where the object is constant throughout all frames. PCID jointly estimates the common object and the blurring from each frame by iteratively estimating the object and the blurring per frame. The object and PSF estimates more closely match the collected frames with each processing iteration cycle. Physical constraints, such as object positivity, use of support for the object and PSF, as well as noise weighting contribute to fast convergence to the object estimate.

Figure 1. Maui Space Surveillance System

PCID is best accessed by image analysts and other users through a Net-centric or Web user interface and job management system that has been co-developed with PCID called the Advanced Speckle Image Restoration Environment (ASPIRE). The output of interest is an enhanced image, or recovery, of the object. PCID developers also researched and created Cramér–Rao estimation theory and statistics lower bound (CRB) expressions for blind deconvolution and carried out a comparison study that showed that PCID closely approaches these theoretical limits when run in the nonblind mode. It also should be able to approach the theoretical limits in the blind mode when regularization and positivity is employed. A high-level graphic of PCID and ASPIRE is shown in Figure 2.

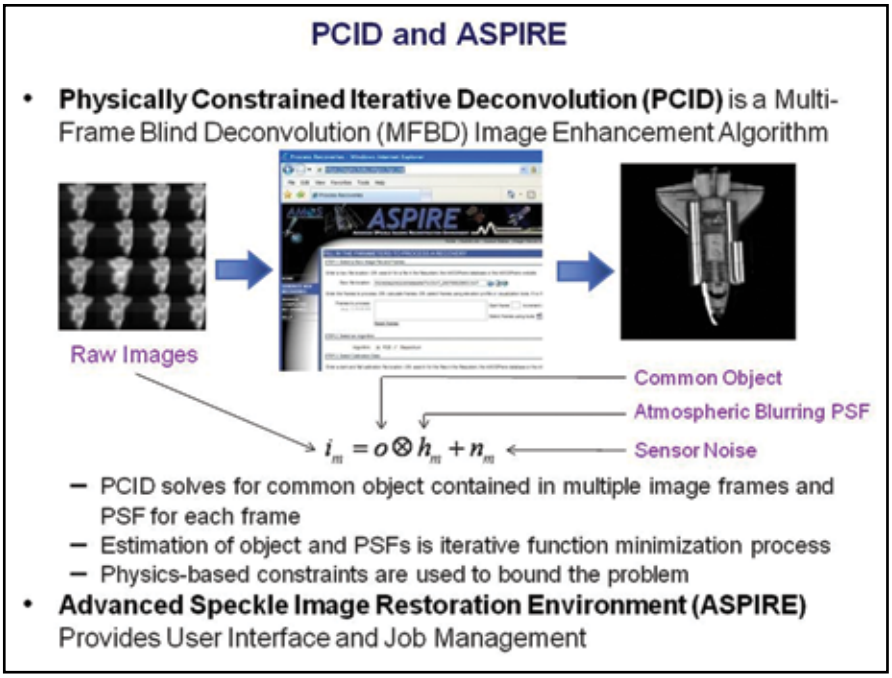

Figure 2. PCID and ASPIRE

Figure 3. FY10 dedicated high performance investment (DHPI) Kaku system

PCID and ASPIRE are continuing to utilize a DoD High Performance Computing Modernization Program (HPCMP) FY10 Dedicated High Performance Investment (DHPI) platform with characteristics outlined below.

- DoD HPCMP FY10 DHPI Award
  - *Kaku*
- Repurposed and upgraded former MHPCC DSRC *Jaws* nodes
  - 150 Power Edge 1955 compute nodes; Intel Xeon Woodcrest dual-core processors
  - 8 GB/core upgraded memory
  - 600 cores
  - Single Data Rate Infiniband, Gig-E interconnect
  - Linux Red Hat OS
  - 36 TB direct-attached storage

Our image enhancement team has continued to maintain and utilize PCID and ASPIRE to continue to perform R&D for improved performance. These packages are also used as described before to obtain high-quality images and to provide quick turnaround, high-throughput, postmission processing of an entire space object overhead pass in less than 30 minutes. The increased resolution and quality of imagery, which has been processed with the enhanced PCID algorithm and the FY10 DHPI HPC platform at the MHPCC DSRC, is shown in Figure 4.

The MSSS modernization and upgrades, which have occurred in calendar year 2011 and 2012, included new/upgraded infrastructure and sensors, including the 1.6 meter telescope new *FLASH* sensor and camera, which provides imaging and wide field of view (FOV) for multiple development and mission applications. This camera provides several new benefits that include using electron-multiplying charge coupled devices (EMCCD) for increased gain and lower readout noise. Modifications to PCID, ASPIRE, and another image enhancement algorithm, Parallel Bispectrum (PARSEC), were needed, as many new data headers in the camera sensor information files were not the same as previous MSSS sensors. The new MSSS *FLASH* sensor uses more of the instrument array as the FOV is increased so that sampling at larger fields of view is correct. Since *FLASH* operates at higher frame rates and larger image sizes than the sensor it replaced, HPC-backed image processing is even more valuable in support of MSSS data collections.

Many tasks were identified regarding the ASPIRE, PCID, and PARSEC software engineering that needed to be performed on the image enhancement calibration, ensemble generation, algorithms, and other utilities modifications. A review of calibration, ensemble generation (ensembles of camera picture frames), and modification of algorithms, and other software utilities work was needed and performed as outlined below.

ASPIRE, PCID, and PARSEC:

- Modify calibration, ensemble generation, algorithms, and other utilities
  - Review new *FLASH* sensor

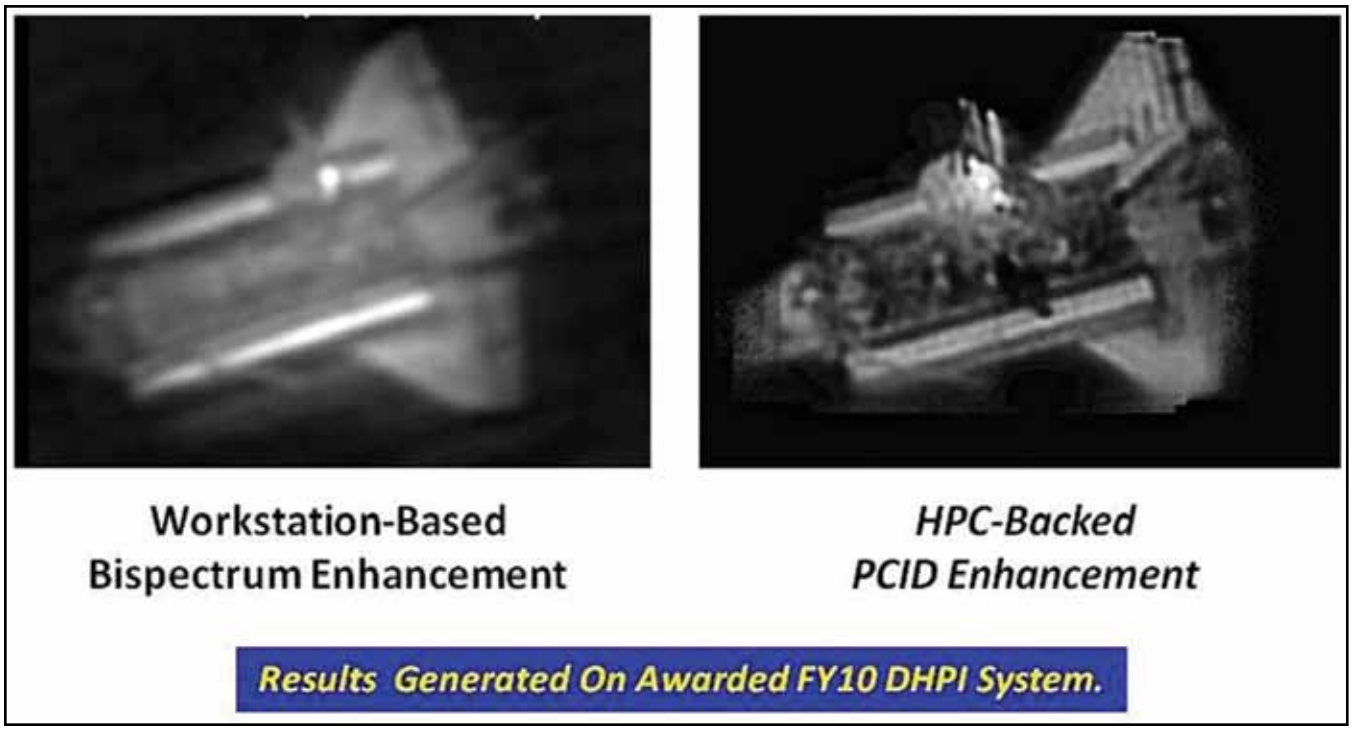

Figure 4. HPC-backed PCID provides increased image quality

calibration process to flexible image transmission (FITS) files

– Provide automatic calculations for object support, filter cutoff, PSF support, and others

– Modify ensemble generator for operation with *FLASH*, adjust run times, and batch submittals

• Ensure that Q data parameter scaling based on sampling rate of data and value used to scale all Fourier domain quantities is correctly set in parameter namelist

• Modify PCID and PARSEC submit pages on ASPIRE for efficient and valid processing

• Create new release versions PCID v11.0 and PARSEC v8.0

– New PARSEC adds "embed" to handle clipped frames and large object data for *FLASH* sensor

– Regression test and install on *Kaku*

– Provide validation support for *FLASH* sensor data processing while ensuring that GEMINI and AEOS data reduction is still valid

Significant improvements to image-processing enhancement have been realized over the past several years. The PCID software has realized more than a 10,000-times speedup compared with the original R&D code, and a speed-up of greater than 72 times has been achieved between calendar years 2008 and 2011.

PCID v11.0 and PARSEC v8.0 development was completed by April 2012. PCID v11.0 is comprised of 11,970 source lines of Fortran, C, and other files. PARSEC v8.0 is comprised of 3463 source lines of Fortran, C, and other files. The PCID v11.0 image-processing regression testing performance using standard datasets provides equal or better image recoveries than the previous PCID version. PCID v10 provided a 20 percent speedup in code execution, and PCID v11.0 provides more than a 10 percent additional speedup over the PCID v10 version. The team ported all software to *Kaku* and completed regression testing on this DHPI platform environment. The PCID automatic calculations for object support, filter cutoff, point spread function support, and other items were verified and adjustments made as necessary to these parameters to deliver the best overall recoveries using the automatic parameters.

MSSS modernization support tasks were performed by our team so that this latest version of PCID further improves the efficiency of the code as noted above while providing optimization for the new *FLASH* sensor at MSSS. ASPIRE was also updated to include a "no star" option if a reference star calibration is not needed or viable. This PARSEC update added correction for the object amplitudes based on values derived from the data that eliminates star measurement. ASPIRE, PCID, and PARSEC ensemble generator and parameters selection software modifications and new developments to process calibration data, calibrate *FLASH* data files, and interact with new data formats like HDF5. Code modifications were also performed to validate the Common Gateway Interface (CGI) script code component to account for the various new data types so that the software can flexibly read the header of FITS files to set the sensor in the validate CGI software.

Our team has successfully updated image enhancement software for MSSS modernization that is available for use. We continue to provide improved SSA technologies and capabilities that leverage HPC. We anticipate the arrival of the new FY12 DHPI HPC platform, *Lilikoi*, before the publication time of this article, when our team plans to prepare for and port image enhancement software and other utilities to the new platform for continued R&D. Additional work has also been performed to develop low-latency, near-real-time processing of data that are collected by the ground-based sensor during overhead passes of space objects. We would also like to continue to progress further towards providing even more capable SSA HPC service-oriented architecture (SOA) services and end-to-end NRT processing and demonstrations as may be directed. Thanks are expressed to the following AFRL team members: Kathy Borelli, Jason Addison, Adam Mallo, Brad Farnsworth, Ron Viloria, Scott Spetka, Bruce Duncan, and George Ramseyer.

Dedicated Support Partition Success

By Marie Greene, Air Force Research Laboratory Deputy Director, Maui High Performance Computing Center DoD Supercomputing Resource Center (MHPCC DSRC)

The MHPCC DSRC continues to host Dedicated Support Partitions (DSPs). Projects have run for more than two million CPU-hours and reserved more than 3000 cores. Successes to date include the following:

"The CREATE-AV DSP, ODEFN29462X01-Shadow Ops/Quality Assurance Software Development impact/ benefits provided immediate runtime access to HPC resources.

This accelerated beta release of CREATE-AV rotary wing software, Helios Version #3; release of CREATE-AV fixed-wing software, Kestrel updated Version #2.2.2; beta release of CREATE-AV fixed-wing software, Kestrel Version #3; and immediate debugging of software issues encountered by applications directly supporting ongoing acquisition programs, enabling rapid user support."

– Joe Laiosa

"We have used our DSP to develop the CREATE-AV Firebolt v2.0 high-fidelity physics turbomachinery CFD capability; perform validation of the CREATE-AV Firebolt v2.0 high-fidelity physics turbomachinery CFD capability for single-blade passage and full annulus simulations; and perform validation and alpha testing for the CREATE-AV Firebolt v1.0 0D engine capability coupled with Kestrel v3.0. The DSP has allowed the Firebolt team to make significant progress in the release of the v1.0 capability and in the development of the v2.0 capability. Both of these capabilities will allow acquisition engineers to assess propulsion/airframe integration well in advance of flight test. This represents a new capability for the DoD."

– Bobby Nichols

"We have been getting up to speed on using the Kestrel (CREATE-AV) CFD program and have generated some initial solutions that will eventually be used to facilitate fly-by-wire development on the F-15SA program."

– Dave Stookesbury

FIREBOLT – Module for propulsion systems in fixed- and rotary-wing air vehicles

SHADOW OPS – Computational tools to support acquisition programs that provide experience and establish connections and value

All project leaders with an active RDT&E computational project are eligible to submit a DSP proposal. All application software development efforts, large-scale weapons system test support, and other activities requiring substantial dedicated time on HPCMP resources that cannot be serviced through normal batch processing, interactive processing on the new utility servers, nor the Program's advance reservation service will be considered. Proposals may be submitted directly to the HPCMPO at require@hpcmo.hpc.mil.

From the Director's Desk – Frank Witzeman

The Technology Insertion 2011/2012 (TI-11/12) process for the DoD HPCMP was finalized recently with awards made to Cray, SGI, and IBM. For the AFRL DSRC, the portion awarded to Cray included an upgrade to our *Raptor* XE6 system acquired under TI-10. During June, every compute node (2732 total) socket was upgraded from AMD "Magny-Cours" 8-core 2.4 GHz processors to AMD "Interlagos" 16-core 2.5 GHz processors, with associated memory of 2 GB per core. *Raptor* effectively doubled in capacity from 16 cores per node to 32 cores per node, resulting in 87,424 cores total (up from 43,712 cores) and over 170 TB of memory (up from over 85 TB). Estimating its peak performance around 800 TFLOPS, *Raptor* is again one of the most powerful supercomputers in the world.

Cray Inc. was also awarded upgrades to the two XE6 systems (*Garnet* and *Chugach*) at the ERDC DSRC, with an additional option to combine *Raptor* and the ERDC DSRC systems into a single large system. The disk storage subsystems for all of the XE6s will also double in capacity for the integration, in order to sustain I/O performance and usable temporary workspace. The new, large, upgraded, and integrated Cray will consist of over 150,000 cores and will be capable of peak performance in the range of 1.5 PFLOPS. Presently, the objective is to have this leadership-class system in operation by early 2013.

The other TI-11/12 award to the AFRL DSRC included a new SGI Altix ICE X system consisting of 4608 Gemini twin-blade nodes installed in 32 M-Racks for a total of 73,728 cores and over 145 TB of memory (2 GB per core). Made up of Intel Sandy Bridge 8-core 2.5 GHz processors, the new SGI named *Spirit* is expected to have peak performance of over 1.5 PFLOPS. An interesting aspect of this system is the design of the cooling approach, where four M-Racks are combined with two cooling units to create a self-contained M-Cell that utilizes both forced cold air and chilled water flowing to/from each blade. *Spirit* will be installed in the new AFRL DSRC building and will be operational in early 2013.

In addition to the TI-11/12 awards, the AFRL DSRC will receive a new Appro cluster to experiment with ScaleMP's vSMP product to build a software-based, large, shared-memory environment. This experimental system called *Lancer* will consist of 160 nodes of dual Intel "Sandy Bridge" 8-core 2.5 GHz processors for a total of 2560 cores. Each core will have access to 8 GB of memory in the standard cluster mode (20 TB total). Initially, vSMP will enable the ability to generate an 8 TB single-system image for shared-memory applications and can leverage the 192 TB of Lustre-based usable storage workspace to generate virtual scratch space. In late 2012, *Lancer* will be made available to HPCMP users who demonstrate requirements for large, shared-memory applications, as well as those who are developing new methods that can exploit vSMP.

From the Director's Desk – Dr. Raju Namburu

Welcome to another issue of HPC Insights, and thank you for the opportunity to share some of the new projects and strategic high performance computing (HPC) initiatives being worked by the Army Research Laboratory DoD Supercomputing Resource Center (ARL DSRC) team. An important project that has kept us busy during the past few months is the preparation of our newly acquired HPC facility to host our Technology Insertion 2011/2012 (TI-11/12) systems arriving in September 2012. Our HPC facilities and engineering teams have done a remarkable job to refurbish the new facility infrastructure and lay the groundwork for current and future ARL DSRC HPC systems. I am pleased to report that we are on schedule for our two IBM iDataPLex systems to be deployed and ready for production by the end of this calendar year.

Our HPC initiatives include the recent strategic partnership we've formed with the Army Test and Evaluation community towards addressing their computational and I/O-intensive requirements and the demand for "big data" analytics for the Army's Network Integration Evaluation (NIE) Process. Our "Green" HPC initiative includes projects such as power-aware computing and efficient facility design that has already revealed several practical opportunities to improve the Center's HPC efficiency and operational overhead.

The ARL DSRC and DoD HPCMP Mobile Network Modeling Institute (MNMI) have developed a collaborative partnership with the Army Test and Evaluation Command (ATEC) to develop and implement scalable algorithms to process massive amounts of raw data (or big data) from live network experiments conducted at the Army's biannual Network Integration Evaluation (NIE) Process. The Army has a pressing need to improve this process, and the ARL DSRC is taking a proactive role with ATEC to address this problem.

With the proliferation of tactical sensors and data-gathering instrumentation in the battlefield, the volume of data has increased exponentially to the extent of almost "choking" available network bandwidth in tactical environments. The process to transfer, analyze, or "mine" this glut of data is both time-consuming in transfer and nearly impossible to analyze in a timely manner. The Army is using NIE, as a key part of the Agile Process, to assess potential network capabilities in a robust operational environment to determine whether they perform as needed, conform to the network architecture, and are interoperable with existing systems. The NIE ensures that the network satisfies the functional requirements of the force and relieves the end user of the technology integration burden. During the FY12.1 NIE process, ATEC collected data from different network devices executed under different scenarios. When receiving multiple terabytes a day, it takes longer to process the data than to collect it. In order to improve the performance and turnaround, the ARL/ATEC team has developed scalable algorithms to analyze the data on large-scale supercomputers. Our new approach using 100 cores seems promising and shows an order of magnitude improvement in analyzing data compared with the existing approach.

In our Green HPC Initiative, the ARL DSRC has been proactive in both local and Program-wide efforts in power-aware computing. With the average HPCMP HPC system utilization in the 80 percent range, opportunities exist to power off idle compute nodes, reducing energy costs for the idle nodes and associated cooling.

Scheduling of reservations, system maintenance events, and scheduling of large-node count jobs all create opportunities to power off idle nodes (idle nodes consume approximately 50 percent of the power of an active node).

Early returns from using the Energy Aware Scheduler (EAS) for HPC systems at both the ARL and MHPCC DSRCs show potential cost savings between 5 and 25 percent depending on the systems' utilization rate. The EAS work to date will be leveraged across other HPCMP systems, to include the incoming TI-11/12 platforms, and will undoubtedly further the Program's role as a leader in Green HPC.

An ongoing ARL DSRC Green HPC effort is the implementation of the "Keep-Alive" feature of the EAS. This allows specific classes on nodes (compute, graphic, and bigmem on the utility servers, for example) to be forced to remain powered on and idle. This provides an improved user experience, as nodes are available immediately for users because they do not have to wait for them to be powered on by the scheduler.

The Keep-Alive feature has been implemented on the ARL utility servers and SGI systems (*Harold* and *TOW*) and is showing great results so far. Additional nodes are getting powered off nights and weekends, and fewer nodes powered off during the normal workday.

Another successful power-aware, computing-related measure was implemented to improve performance on the ARL DSRC's SGI systems. Our systems engineering team encountered a discrepancy with the rate at which SGI compute nodes are powered on that caused a performance degradation in the Lustre file system when an excessive number of compute nodes are powered on all at once. By limiting the number of active "cpower" commands that the boot_nodes command see active at one time and metering out new cpower commands at a slower rate, we have resolved this issue.

We continue to explore new and innovative Green HPC techniques and partnerships within the national HPC community and to play a proactive role in integrating practical solutions for the resources and investments for our DoD community.

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

Existing Equipment Augments New Facility and Reduces Cost

By Michael Knowles, ARL DSRC

The ARL DSRC undertook an extensive building renovation to accommodate new HPCMP assets that arrived this fall. The new assets, referred to collectively as TI-11/12 (Technology Insertion 2011/2012) systems, include one 20K (*Pershing*) and one 18K (*Hercules*) core clusters from IBM. To accommodate these systems, extensive facilities infrastructure modification was required. These infrastructure modifications included emergency generators, UPS systems, electrical distribution systems, building cooling enhancements, network equipment, office design and implementation, and architectural changes (walls, floors, painting, windows). Building modifications commenced in March to accommodate the proposed systems. ARL personnel who performed these modifications were extremely diligent in the reuse and repurpose of existing equipment to minimize overall site modification and installation costs. These stewardship efforts are described herein.

To ensure the TI-11/12 systems are maintained at full operating capacity on a 24x7 schedule, emergency generators and UPS (uninterruptable power supply) systems were required. In March, ARL was awarded a new facility at Aberdeen Proving Ground (APG) that was immediately selected to house the new TI-11/12 systems. At that time, the new facility had one 2MW emergency backup generator that handled the overall building power as well as some smaller computer systems load fed by two existing 300 KVA UPS systems. The building load proved to be small enough to allow for one of the TI-11/12 systems (*Pershing*) to be supported by the generator. Significant modification of the existing generator was required to ensure it could support the new load for an extended period of time. Actually, the generator wasn't yet tied into the building at the time although wiring was in place. Additionally, a new fuel tank and enclosure were built to ensure the generator can run for more than 4 hours at a time. The *Pershing* power requirements exceeded the capacity of the existing UPS in the new facility, so ARL identified two UPS systems in the current facility that were underutilized. These UPSs were repurposed, relocated, and refurbished to support *Pershing's* primary computational cabinets. Both existing UPS systems from building the new facility were also relocated, refurbished, and repurposed to ensure *Pershing* disk racks and other building infrastructure can survive power outages. An older power distribution unit (PDU) was also reclaimed, refurbished, and repurposed for TI-11/12 support. Finally, rack-attached power in one of the existing new facility computer rooms was provided by a BUSBAR system that allows for individual power drops to racks. This unique distribution system was also removed, relocated, and repurposed for TI-11/12 disk racks.

In January 2012, two barely used 2MW generator systems became available at the Edgewood Area of Aberdeen Proving Ground. These generators were installed in 2001, but had only slight usage and were currently unused. The generators had 4180V alternators attached, however, and the new facility requirements were for 480V. These generators also had minimal fuel capacity (4 hours), so the enclosures and tanks had to be replaced. The two excess generators were relocated, refurbished, and repurposed to support *Hercules'* emergency power as well as future ARL power requirements. New fuel tanks and enclosures were fabricated for these generators.

The cooling requirements for each of the new TI-11/12 systems exceed 160 tons. The new facility had two redundant 210-ton chillers already installed. Calculations showed that the existing chiller loop for the new facility had the additional capacity to cool the *Pershing* heat load, but not that of *Hercules*. However, extensive modification to the existing chiller loop could be performed to support the existing building cooling load, as well as the new TI-11/12 requirements. Therefore, modification to the existing loop (including additional glycol), refurbishment of one of the existing chillers, additional pumps, and the addition of a new chiller (to handle the *Hercules* heat load) were required. The chillers will work in an N+1 mode of operation for redundancy. Finally, two computer room air handlers (CRAHs) were refurbished and repurposed to cool the area where the additional UPS systems were relocated.

While the work performed by the ARL team in the preparation of the new facility was extensive and completed in a short time frame, the extent of the stewardship of Government funds and effective reuse of equipment displayed a remarkable capacity to design and implement a cost-effective site modification while minimizing waste.

Shared License Buffer—Managing Shared Software Licenses Is an Important Part of Managing the System Workload

By Steve Thompson, ARL DSRC

The Shared License Buffer (SLB) has been controlling licenses for Abaqus, ACFD (Ansys CFD), and Cobalt since February 2012. Its job is to ensure that user jobs needing shared licenses are not started by PBS until they receive a license reservation and are thus guaranteed of getting the licenses they need. The SLB thus prevents the situation where a job starts, fails because it cannot get licenses, and has to be resubmitted by the user. To use the SLB, users simply enter a line into their PBS scripts specifying which licenses they need and how many. PBS handles contacting the SLB on behalf of the job and coordinates starting or requeuing the job with the SLB based on license availability for the duration of the job. The SLB also interfaces with ARS (Advance Reservation System) so that users can reserve licenses at the same time they reserve CPUs and be guaranteed that the licenses they need will be available during their ARS reservation.

The SLB currently also accepts requests and makes license reservations for Accelrys and Matlab, but enforcement is currently turned off for these two applications, meaning that there is no guarantee that the licenses reserved will be available. We hope to enable enforcement for these two applications soon. But even without enforcement, PBS still knows not to start jobs unless licenses are available, so users greatly improve their chances of avoiding failed jobs by requesting SLB licenses.

Gasp, Star-CCM+, and LS-Dyna were recently added to the SLB. Gasp users may reserve a "gasp" license (one per job), Star-CCM+ users may select a "ccmppower" license (one per job), and LS-Dyna users may select either an "lsdyna" license or an "mppdyna" license depending on the executable binary they will run (one license per CPU). For example:

```
#PBS -l gasp=1
#PBS -l ccmppower=1
#PBS -l mppdyna=32
```

In addition to accessing the SLB via PBS or ARS, users may also access the SLB directly to make what is called an "interactive reservation." An interactive reservation is done by running the SLB script SLB_make_resv.pl from the Linux command line. It allows a user to make a license reservation, though a number of policy restrictions apply. Interactive reservations are allowed only for Matlab, Abaqus, and ACFD. Interactive Abaqus reservations are limited to five licenses (enough for one CPU) and to 15 minutes so that users can perform pre- and postprocessing. Similarly, interactive ACFD reservations are limited to one acfd_solver license for 15 minutes. License reservations made through PBS have no such restrictions. For example, to reserve one MATLAB license for the default period of 4 hours starting now:

```
SLB_make_resv.pl -app matlab -n 1
```

Users may also make an interactive reservation for the future. See the SLB documentation posted on the CCAC website for details on additional capabilities and syntax.

ARS reservations also have a policy restriction: Star-CCM+ users may have only one ARS reservation at a time. If a user wishes to see a change in SLB policy, please contact the CCAC, the ESMT (Enterprise Software Management Team), or the UAG.

Several helpful capabilities are available to users via the "query" scripts provided by the SLB. Since the SLB tracks all license reservations, current and future, it has all the information necessary to provide answers about license availability in the future. The examples below show what type of queries can be run and the syntax for running them in SLB 2.0, which will be operational soon.

When will 768 cobalt licenses be available and for how long?

```
harold  >      SLB_query_when.pl -app cobalt -tok
"cobalt=768"
cobalt:  768 licenses available
   from  Thu Aug  2 08:21:36 2012 until forever
```

How many abaqus licenses will be available November 4 at 8:30am and for how long?

```
harold >    SLB_query_avail_t.pl -app abaqus -st
11/04_08:30 -tok abaqus
will query for Sunday  Nov  4  08:30:00 2012
abaqus: 244
   available from Sunday  Nov  4  08:30:00 2012
until forever
```

If I make a reservation from Aug 1 from 8:30am to Aug 2 at 10:45am, what is the maximum number of "anshpc" licenses I could reserve?

```
harold>  SLB_query_avail_block.pl -app acfd -st
08/01_08:30 -et 08/01:10:45 -tok "acfd_solver an-
shpc"
will query time window from Wed Aug 1 08:30:00 to
Thu Aug 2 10:45:00 2012
acfd_solver: 3
anshpc:  68
available from Aug 1 08:30:00 to Thu Aug 2
10:45:00 2012
```

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

The U.S. Army Engineer Research and Development Center (ERDC) is the Nation's premier laboratory for research and development in civil works and military infrastructure, including navigation, flood protection, force protection, and high performance computing data centers. In recognition of ERDC leadership, the Department of Defense has historically maintained a DoD Supercomputing Resource Center (DSRC) on the Information Technology Laboratory (ITL) campus at ERDC-Vicksburg.

Beginning in late 2012, the ERDC DSRC will host a leadership-class Cray XE6 supercomputer in support of the entire DoD research, development, test, and evaluation community. The system will rank among the fastest supercomputers in the world, providing 150,000 CPU cores for the solution of large-scale physics and engineering simulation problems. The system is approximately half the size of the Department of Energy's 300,000-core Cray XE6 located at Oak Ridge National Laboratory's National Center for Computational Sciences (NCCS), recognized as the Nation's fastest leadership-class supercomputer.

Unlike the NCCS system, which serves about 40 projects and 400 users in the open science community, the ERDC system will serve sensitive (but unclassified) DoD projects tied to Army, Navy, and Air Force research missions. Historically, ERDC has served about 900 users although this number is subject to debate based on the definition of an active user. All are required to obtain a National Agency Check (NACI).

The size of the new system presents significant management challenges. The DSRC currently hosts four smaller systems, the largest of which is 20,000 CPU cores, with a total less than 60,000 CPU cores. The new system will therefore require a different approach to systems administration, customer assistance, and technology integration. However, there is a historical precedent for meeting such challenges at ERDC. Over 10 years ago, ERDC integrated a large Cray T3E system using cabinets from several other installations. Just under 2000 CPUs, it was the largest parallel system among the DoD supercomputer sites and the focus of code migration efforts to transition from shared-memory vector systems to the distributed-memory parallel architecture.

An implicit requirement for supercomputer centers is to provide "capability" computing, or the opportunity to run jobs that use a significant fraction of the system. Scheduling policies, including job size and run time limits and priorities, are used to support capability computing. Historically, the DSRC has limited job size to about one-fourth of system size and provided no priority mechanism for larger jobs in the scheduling system. Maximum run time policies have gradually increased, such that jobs may run uninterrupted and un-checkpointed for up to 1 week. These scheduling policies cause larger jobs to experience longer queue waits than smaller ones.

Capability jobs require a priority scheduling policy in order to experience the same queue wait time as a smaller job. However, priority scheduling is efficient only with relatively rapid queue turnover, because the scheduler must accumulate idle resources for the next capability job. Queue turnover is a function of the maximum job run time policy. Thus, capability computing requires priority scheduling with relatively short maximum run times, compared with the current DSRC policies. In comparison, NCCS policies elevate queue priority for jobs that request 40 percent or more of the system resources, and limit the length of large jobs to 24 hours and small jobs (less than 256 CPU cores) to 2 hours.

From the Director's Desk – David Morton

As an old hardware geek, I am going to devote this article to our new MHPCC supercomputing system. Here are some of the details on our new supercomputer named *Riptide*. It was procured under the HPCMP Technology Insertion 2011/2012 effort. *Riptide* is an IBM 1350 IDataPlex system with 12,096 cores. The processor technology is Intel Sandy Bridge 2.6 GHz with dual 8 core processors. It has 2 processors and 16 cores per node with each core having 2 GB of memory per core. The system interconnect is based on InfiniBand FDR10 technology. Total peak performance is 246 TFLOPs. *Riptide* will be installed by the publication time of this article, and the goal is to make the system available for general use in February 2013.

While it may not affect HPC users directly, *Riptide* has several innovative energy efficiency technologies that are incorporated into the design and is seen as an energy efficiency demonstration vehicle for the HPCMP. Energy use has become a key technology driver within the HPC community, and if the current energy use growth curves continue, could become a financial constraint within the HPCMP, thus directly impacting users. It is becoming obvious we have to "bend" the energy growth (and related infrastructure costs) curves downward if the HPCMP is going to be able to deliver the same historical level of constantly increasing supercomputing performance.

One of those *Riptide* technologies is direct warm-water cooling. IBM has been able to design a system that uses an incredible, thermally efficient, direct water-cooled technology. Each node board has a number of water-cooled heat sinks that connect directly to the high-power components. This node-level water system uses a quick disconnect technology that enables easy removal. Because the water-cooled system is so thermally efficient, the inlet water temperature can be as high as 113 degrees F. This means that chilled-facility water is not required; the water used in this system can just be pumped outside where a simple heat exchanger can cool the hot water coming from the system to an acceptable level for reuse in the system. No refrigeration will be required for over 90 percent of the energy being dissipated from the compute nodes on *Riptide*. This technology is interesting, as there will be a large reduction in refrigeration use (typically, 20 to 40 percent of total infrastructure energy use), as well as resulting reduction in facility infrastructure upfront costs.

Overall, when compared with the current *Mana* system, we believe that *Riptide* coupled with some other facility improvements will deliver 2.4X capability at approximately 50 percent of the energy usage. The 2.4X is primarily driven by processor improvements, but the overall metric is appealing. Historically, each new system delivers more performance, but also uses more energy.

As you can tell, I am excited about the new *Riptide* system. I believe that it will serve our users well while demonstrating technologies that could have long-term impacts on the HPCMP.

Interns from Universities and Military Academies Receive Research Experience

From May 21 – August 17, 2012, 14 undergraduate and graduate students participated in research internships at the MHPCC DSRC. Programs supported by the MHPCC DSRC included the DoD HPCMP Military Academies, AFRL/DE Scholars Program, Akamai-Center for Adaptive Optics (CfAO), and the Maui Economic Development Board. Universities represented were the U.S. Air Force Academy, U.S. Naval Academy, University of Buffalo, University of Hawaii, University of Southern California, Oregon State, and California Polytechnic State University.

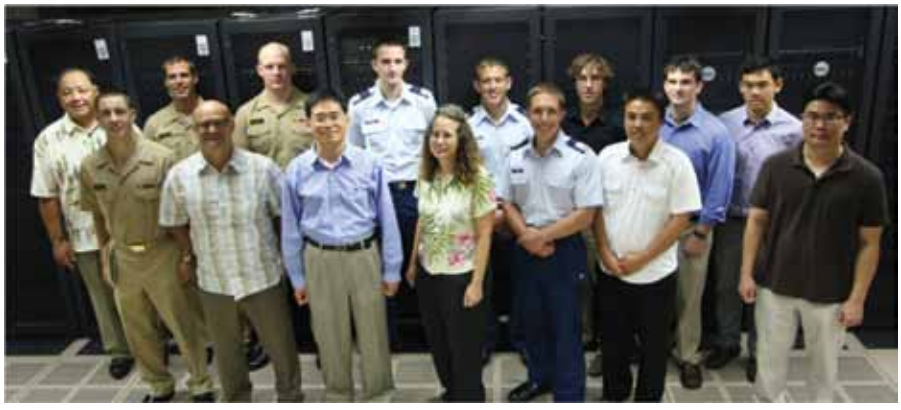

Front Row (left to right): Ian Sonnenberg (USNA), Dave Morton (MHPCC/AFRL), Hyung Suk Kang (USNA Faculty), Laura Ulibarri (AFRL/MSSS), Michael Jones (USAFA), Nathan Hara (California Polytechnic State University), and Aaron Ahue (University of Hawaii)

Back Row (left to right): Gene Bal (MHPCC/UH), Benjamin Rowe (USNA), Chance Rauscher (USNA), Aaron Gibson (USAFA), Brandon Mueller (USAFA), Michael Andrle (University of Buffalo), Matthew Whittaker (University of Buffalo), and Andrew Yau (University of Southern California)

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

Each collegian presented their research findings to the AFRL/MHPCC staff. Many of the interns also served as beta testers for the MHPCC DSRC Portal accessing the Kestrel software from the Portal.

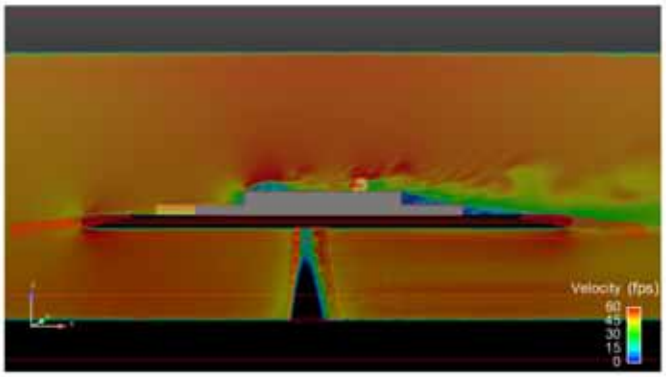

Simple Frigate Shape 2—Simulations Performed at AFRL MHPCC

Modeling the Falcon Telescope Network

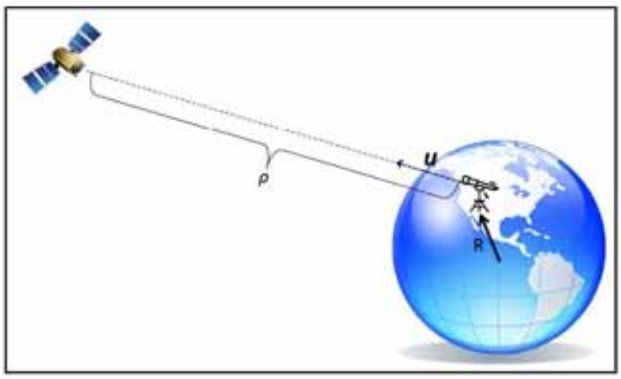

Track Initiation for Optical Space Surveillance

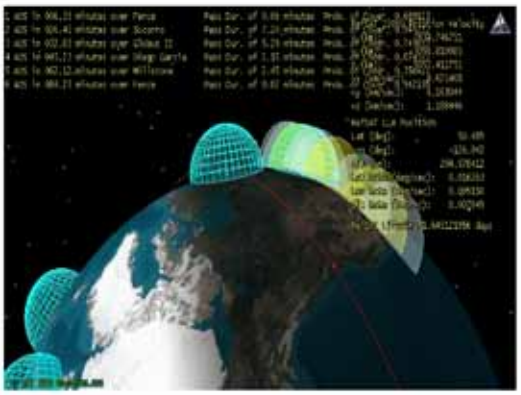

Reentry Analysis Using ReDVU

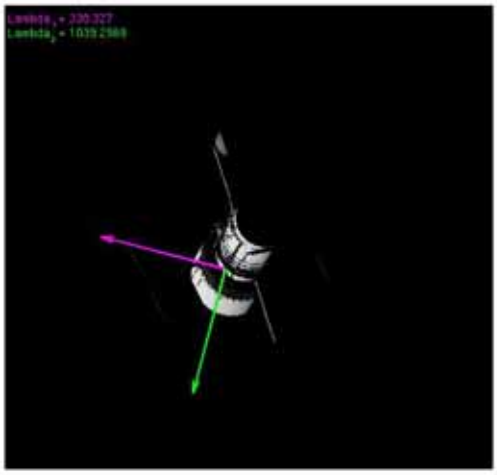

Advanced Satellite Characterization

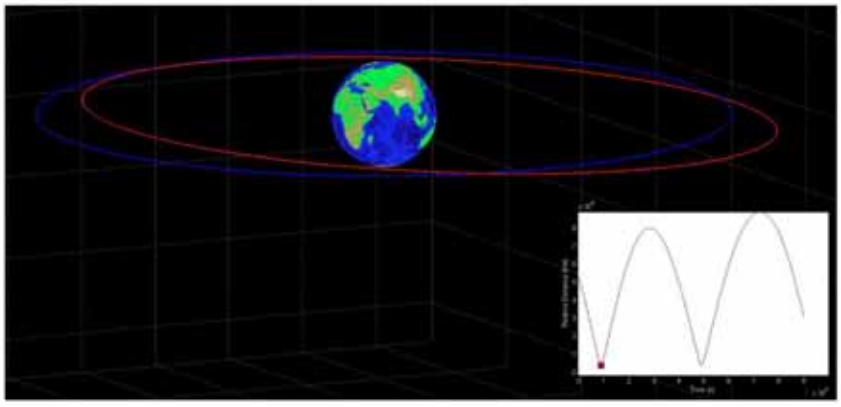

Conjunction Analysis in Space

From the Director's Desk – Tom Dunn

In this issue of HPC Insights, you'll find two user success stories that underscore the importance of HPC and details of Technology Insertion 2011/2012 (TI-11/12) systems that are being installed and tested even as I write this article.

In their article "High-Fidelity Simulations of Bubble, Droplet and Spray Formation in Breaking Waves," Zhaoyuan Wang, Jianming Yang, and Frederick Stern describe how wave breaking is of great importance to ship hydrodynamics, which can produce strong turbulence with large amounts of air entrainment, bubbles, droplets, jets, and spray. Dr. Travis A. Smith's article describes how the Naval Research Laboratory (NRL) and the Naval Oceanographic Office (NAVOCEANO) implemented the Coupled Ocean/Atmosphere Mesoscale Prediction System (COAMPS) version 5.0 at the Navy DoD Supercomputing Resource Center (DSRC) in support of the Talisman Sabre naval exercise in July 2011.

Also of special note, the Navy DSRC will more than triple its computing power this year when the three new TI-11/12 supercomputers are available for allocated use. Two of the iDataPlex systems will be identical, each consisting of 18,816 Sandy Bridge Intel processor cores, 37 TB of memory, and 2.8 PB of disk storage available for computational modeling and research. A third iDataPlex system will have 4032 of the same processor cores, 8 TB of memory, and 576 TB of disk storage. The peak computational capabilities of the two larger systems will be 351 TFLOPS each, and the third system will be capable of 75 TFLOPS. The Navy DSRC website has been updated to provide this information. Please visit www.navo.hpc.mil/hardware/ and www.navo.hpc.mil/docs/ to learn more.

In determining a theme for system names, we connected our location—NASA's Stennis Space Center in South Mississippi—with the Navy roots of our parent command, Naval Meteorology and Oceanography Command (NMOC). The theme chosen for the new systems will honor astronauts who have served in the Navy: namely, Biloxi, Mississippi, native Fred Haise, a retired U.S. Air Force officer who also served as a Navy and Marine Corps aviator and the Apollo 13 pilot; Commander Susan Still Kilrain, a naval aviator who piloted two shuttle missions and more than 30 different aircraft; and Captain Eugene Cernan, a naval aviator and the last person to set foot on the moon.

These supercomputers will enable the DoD science and research community to test and model defense systems that cannot be modeled in the real world due to time and financial, physical, or safety constraints; and in some cases, they can accomplish this work in a matter of hours as opposed to the days, weeks, or even months that traditional research methods can require.

Got VIZ Questions? Ask Our Staff – support@daac.hpc.mil

The Data Analysis and Assessment Center (DAAC) resources are free of charge to you as a user of the High Performance Computing Modernization Program (HPCMP). These resources are not just the hardware and the supported software dedicated to data analysis; they also include the most valuable resource of all—the data analysis expertise of the DAAC staff.

The uses of data analysis or data visualization vary widely among the large HPCMP user community. Consequently, DAAC offers a wide range of data analysis services. These services are characterized by three broad categories: community, collaborative, and custom.

The community service is primarily for researchers who are either new to visualization and wish to get started with minimal effort or want to learn more about how to effectively visualize their data. The embodiment of this service is our website, http://www.daac.hpc.mil.

The collaborative service is for people who want one of our team members to help in visualizing their data. At this level, a staff member will work closely with you in deciding on what you want to communicate with your visualizations, recommending tools and methods to produce the visualizations, and generating a work flow for you so that you can generate these on your own.

Finally, the custom service is for users that require high-quality images or animations of their data for which they also need to show the data in context or natural environment. These kinds of visualization are effective for high-level briefings to your upper management and to your sponsors.

Image courtesy Dr. Balu Sekar, WPAFB

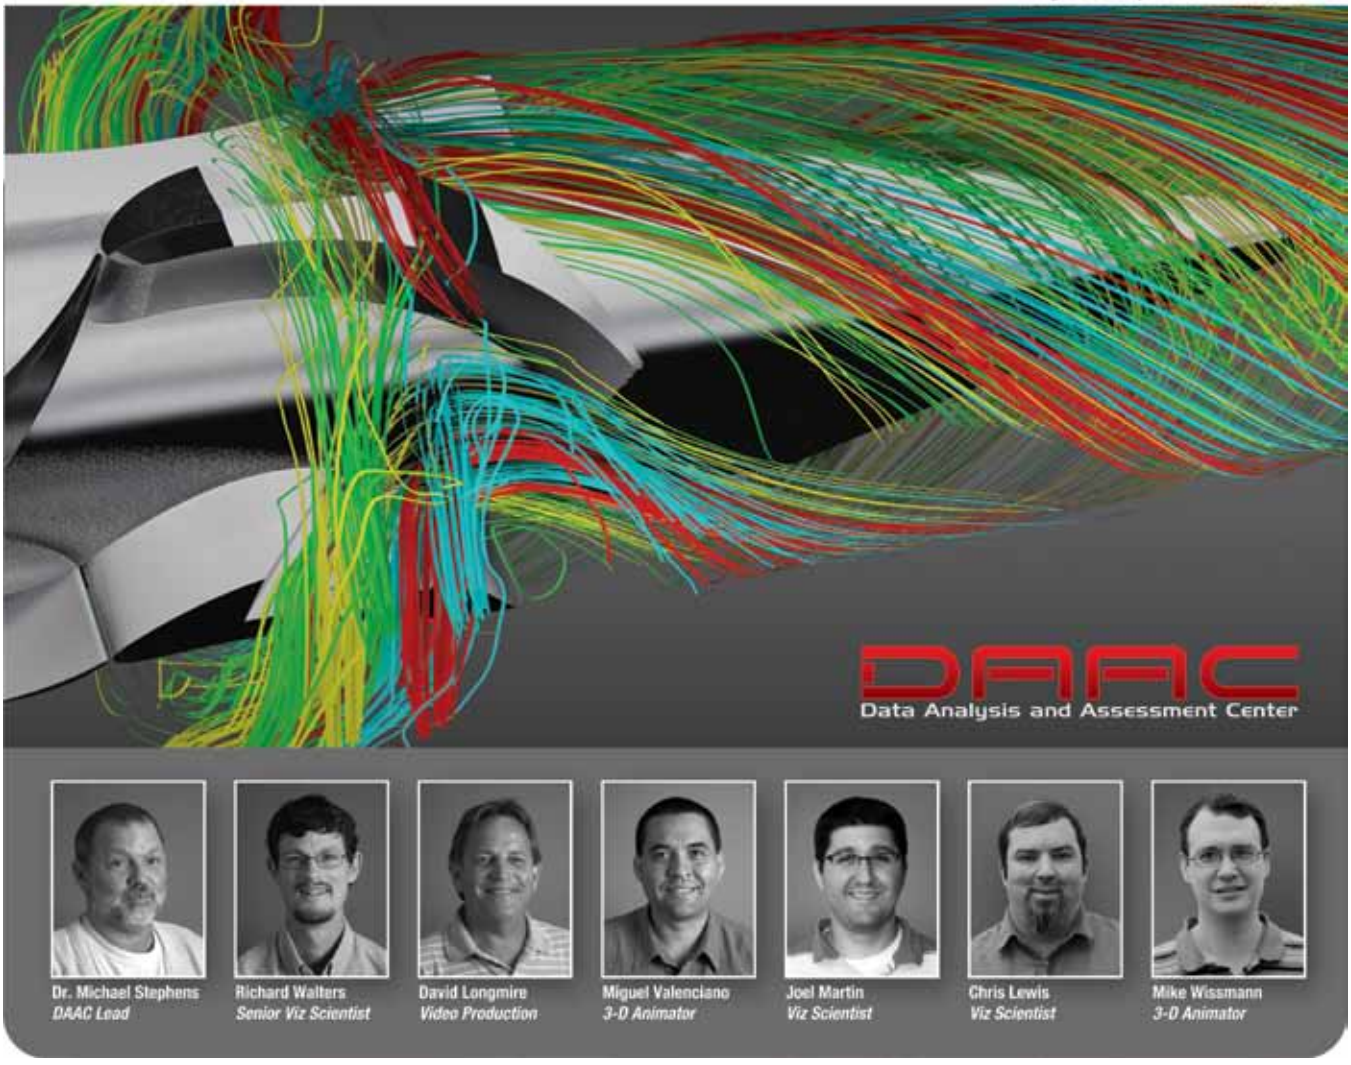

HPCMP "Green Team" Wins 2012 Chief of Engineers Award of Excellence*

The Facilities Community of Practice (COP) team of the DoD High Performance Computing Modernization Program (HPCMP) was recently announced as winner of the U.S. Army Corps of Engineers 2012 Chief of Engineers Award of Excellence entitled Green Dream Team Award. The award recognizes excellence in sustainability, design, and construction achievement. The HPCMP Facilities COP received the award for creating an interagency Facilities COP team, the "Green Team." The team is comprised of representatives from each of the five DoD Supercomputing Resource Centers (DSRCs) and meets monthly to share best practices in supercomputing facilities operation and to plan energy-aware initiatives. It has proposed numerous changes to the supercomputing center infrastructures, consistent with best practices. As the practices have been implemented, each Center has seen a marked increase in efficiency and cost savings. The Chief of Engineers has requested that the team submit their nomination to the GreenGov Presidential Awards Program.

One of the best practices recommended by the COP was the use of magnetic-levitation chiller compressors with economizers. This upgrade demonstrates a designed savings of approximately 30 percent on mechanical power consumption. Additionally, the COP implemented a requirement that supercomputer vendors use high-voltage power supplies in systems they sell to the Centers. This requirement has resulted in savings by reducing the amount of copper used in the systems and the labor needed to connect them to power distribution systems. To maximize efficiency, a 13,800-volt line interactive uninterruptable power supply designed to operate at 98.5 percent efficiency was installed at the U.S. Army Engineer Research and Development Center (ERDC) DSRC. It was also designed for use in an outdoor location, which further reduced the cooling load within the facility and provided a double benefit, i.e., increasing efficiency and decreasing mechanical load.

COP-recommended savings efforts will be seen in the near future, as three of the five Centers are scheduled to install water side economizers. This change will reduce the energy used to mechanically produce chilled water while using the ambient air conditions to deliver "free" cooling inside the facility. This will reduce the run time of mechanical chillers by 3000 hours at the ERDC site alone, resulting in a savings of $225,000 per year.

In 2010 the MHPCC DSRC applied a set of improvements including increasing chilled water temperature, applying air flow containment measures, and turning off unneeded systems, reducing the site's power usage effectiveness (PUE) from 1.65 to 1.39. These changes represent a potential savings of $376,000 per year. Additionally, this site installed a dual-track, triple-junction concentrated photovoltaic power array as a demonstration project. A follow-on to this project is a 1.2MW solar photovoltaic array. The power gained from the sun will be used to supplement conventional power sources while helping the COP gain a greater understanding of how renewable energy can help meet future supercomputing energy needs. The estimated savings at the MHPCC DSRC site total $600,000 per year.

COP research efforts are also focused on potential energy savings areas such as the use of DC power instead of AC utility power for supercomputers, coupled with improved energy recovery methods. Adopting DC power will reduce power transformation losses, while recovered energy can be fed to other energy-consuming systems. The COP is also looking at innovative ways to monitor Center progress in energy efficiency. The Program recently supported DoD investment in an energy

General Bostick recently presented the COP Team of the HPCMP with the U.S. Army Corps of Engineers' 2012 Chief of Engineers Award of Excellence in the category of Sustainability. Accepting the award is Greg Rottman.

*Information presented here comes from the Awards Nomination Package prepared by Greg Rottman.

monitoring system to gather data related to the operation of systems being installed at each Center. This follows the philosophy that "you can't manage what you can't measure." Future coordinated efforts include automated controls that will dynamically adjust water flow and chilled water temperature based on ambient conditions at the Center sites. This will minimize mechanical chiller operation in lieu of free cooling from the atmosphere.

The HPCMP Facilities COP is a highly effective strategy for promoting the pursuit of energy efficiencies in large supercomputer centers. The three Services (Army, Navy, and Air Force) leverage shared lessons learned as they face common problems in support of a common DoD mission. This approach channels otherwise competitive tendencies between Services into coordinated plans and recommendations. Where Centers were once insular, the COP has emerged as a cooperative mechanism to combat the common threat of energy costs. The challenge is to continue the drive to technological leadership, recognizing that the energy budget will be increasing while the fiscal budget will not. The approach can be replicated elsewhere in the Federal Government where similar common problems and missions unite otherwise separate and distinct agency programs. The COP has worked to share technical information, vendor expertise, and lessons learned.

The HPCMP Green Team members include Greg K. Rottman, HPCMP Executive Agent; Brian Schafer and Mark Poe, AFRL DSRC, Wright-Patterson Air Force Base, Ohio; Robert Sheroke and Mark Perdue, Army Research Laboratory, Aberdeen Proving Ground, Maryland; Paula L. Lindsey and Charles Langford, ERDC DSRC, Vicksburg, Mississippi; Carl Shelton, MHPCC DSRC, Maui, Hawaii; and Rob Thornhill and William Cook, Navy DSRC, Stennis Space Center, Mississippi.

"World-Class Network and Supercomputing Capabilities"

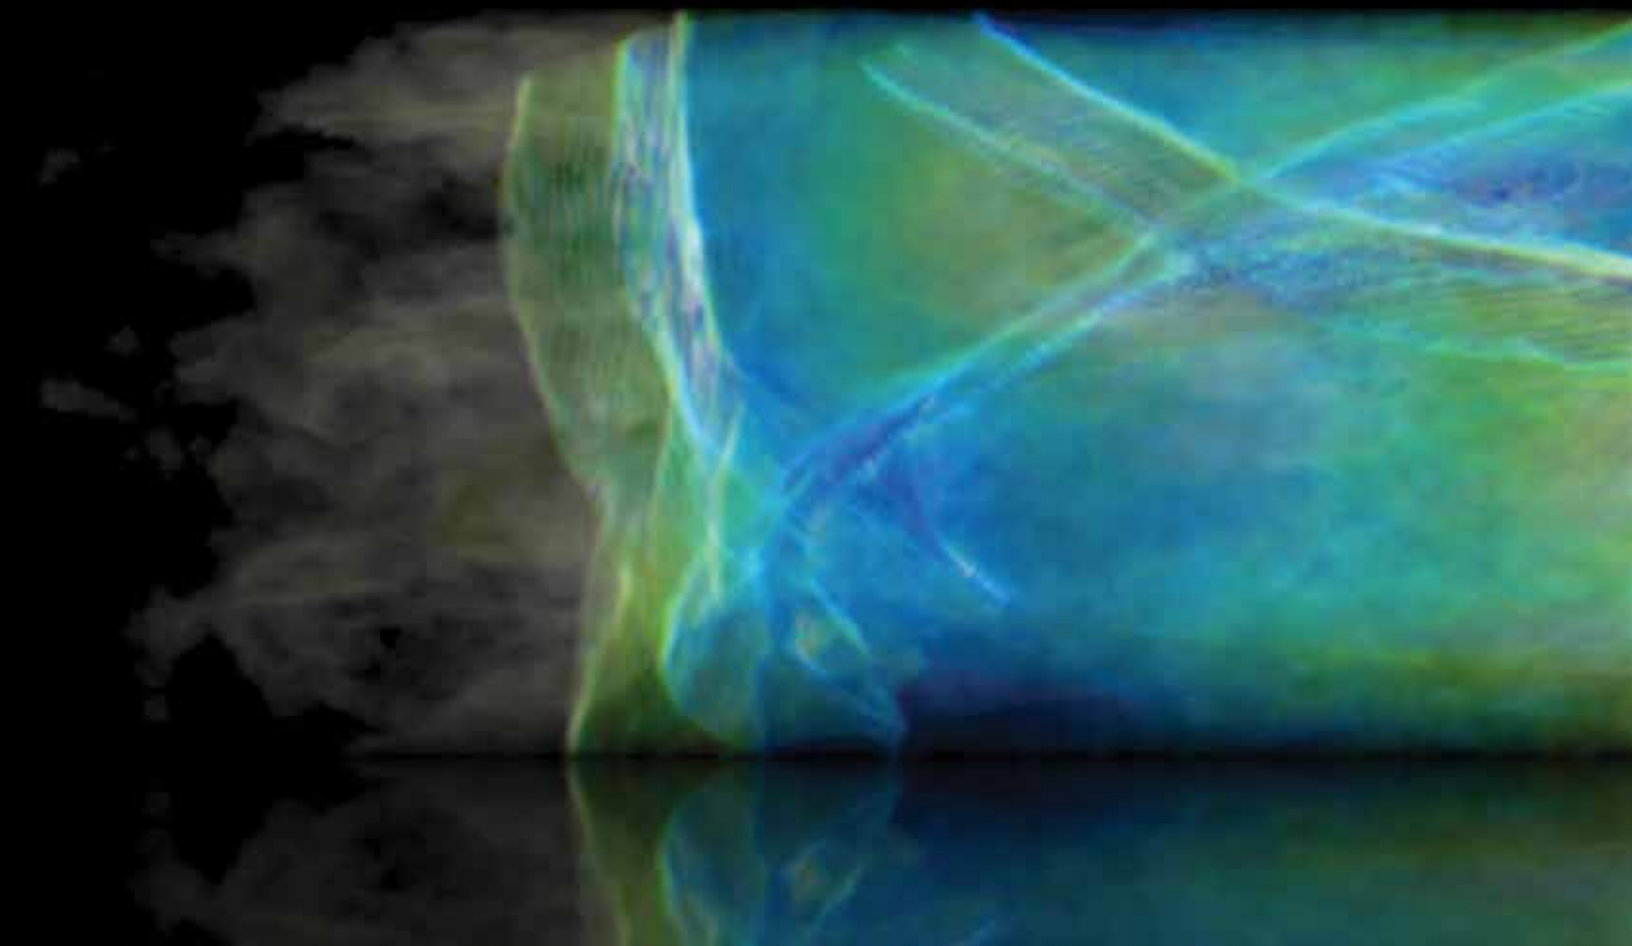

Scientists and engineers throughout the U.S. leverage the capabilities of the High Performance Computing Modernization Program (HPCMP) to solve the most time-consuming computational problems. They know that the HPCMP's supercomputing centers continually architect, deploy, and sustain their equipment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

CCAC Accounts Center
1-877-222-2039

http://www.hpc.mil

DOD HPC
DEPARTMENT OF DEFENSE
HIGH PERFORMANCE COMPUTING
MODERNIZATION PROGRAM