




MISRC

JOURNAL

SPRING 1999



Inside This Issue:

Explosive Detonation • Airborne Laser Project • Nanowires • High Energy Fuels • Topside Communications • Portable Batch System • VAMPIR • Parallel Programming Tools • Adaptive HPC Software

WHO WE ARE

The U.S. Army Corps of Engineers, Waterways Experiment Station (CEWES) Major Shared Resource Center (MSRC) is part of the High Performance Computing Modernization Program (HPCMP) and is located in the Information Technology Laboratory at CEWES in Vicksburg, MS. As a world-class facility, the CEWES MSRC employs a technical staff to provide full-spectrum computational support for DoD researchers, from Help Desk assistance to one-on-one collaboration. More than 4,000 computational scientists and engineers are involved in the HPCMP with immediate access to DoD HPC capabilities, regardless of their locations across the nation, via the Defense Research and Engineering Network.

Other services include a diverse and well-equipped Scientific Visualization Center for visualization expertise and capability. In addition, the Programming Environment and Training component provides for transfer of cutting-edge HPC technology and training and development activities for acquiring HPC skills.

EDITORIAL STAFF

- ✓ Interdisciplinary Coordinator
Mary L. Hampton
- ✓ Outreach Coordinator
Michelle Morgan Brown
- ✓ Chief Editors
Michelle Morgan Brown
Mary Gabb

Approved for public release; distribution is unlimited.

Design and layout provided by the Visual Production Center, Information Technology Laboratory, U.S. Army Engineer Waterways Experiment Station.


The contents of this journal are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval for the use of such commercial products.

CONTENTS

Explosive Detonation in Concrete Structures	1
DoD Challenge Project Studies Airborne Lasers	4
Amorphization and Fracture of Silicon Diselenide Nanowires	6
Design of New High Energy Fuels	8
Topside Communications Get Boost from CEWES	10
Portable Batch System Aids Common User Environment	13
VAMPIR Comes to CEWES MSRC	14
Repository of Scalable Parallel Programming Tools for CEWES MSRC Platforms	16
CEWES Developing Adaptive HPC Software Solution	19

THE FRONT COVER:

Integrated Topside Design, a Department of Defense High Performance Computing Challenge Project. These images show the antenna patterns for a mid-starboard high frequency whip broadcasting at 10.0 MHz. The Integrated Topside Design project is working on new ways to meet requirements for more communications capability with greater imagery and data transfer capacity, while also working to meet aggressive signature reduction goals for U.S. Naval surface combatants. (See article on page 10.)



Explosive Detonation in Concrete Structures

*Kent T. Danielson, Ph.D.,
Mark Adley, Ph.D.,
Stephen Akers*

Using high performance computing (HPC) and scientific visualization, scientists at CEWES evaluate the damage resulting from blast propagation in reinforced concrete structures. Professor Kent Danielson of Northwestern University and the Army HPC Research Center and Dr. Mark Adley and Mr. Stephen Akers of CEWES have developed computational procedures that significantly utilize the modern parallel computing platforms at the CEWES MSRC to simulate such events with advanced constitutive models. These simulations also produce large amounts of data that are best interpreted by visualization. Although the main focus of this project is to assess damage for a specific application, other objectives are to a) assess the overall accuracy of the approach for such simulations, b) evaluate the recently developed microplane constitutive model for concrete, and c) determine the feasibility of routinely performing such analyses with the aid of parallel computing.

The explosive detonation in a reinforced concrete slab is depicted by the finite element model in Figure 1, which consists of 995,192 hexahedral elements and 1,030,089 nodes. The event was experimentally staged at CEWES. The C-4 explosive was placed in a cylindrical cavity at the center of the slab. Quarter symmetry was assumed for

the calculation. The fully coupled explosive-structural analysis uses a microplane constitutive model for the concrete, an elasto-plastic model for the reinforcing steel, and a JWL equation-of-state model for the C-4 explosive. Ignition of the explosive is treated by a programmed-burn algorithm. These procedures were implemented into a parallel finite element code, ParaAble, developed by the authors. ParaAble uses the METIS partitioning software to distribute partitions of elements onto separate processors for computation (e.g., the partitioning example in Figure 2). An overlapping algorithm is used to communicate partition interface data concurrent with element computations for partition interiors.

The development of the microplane concrete constitutive model is a joint Northwestern/CEWES effort

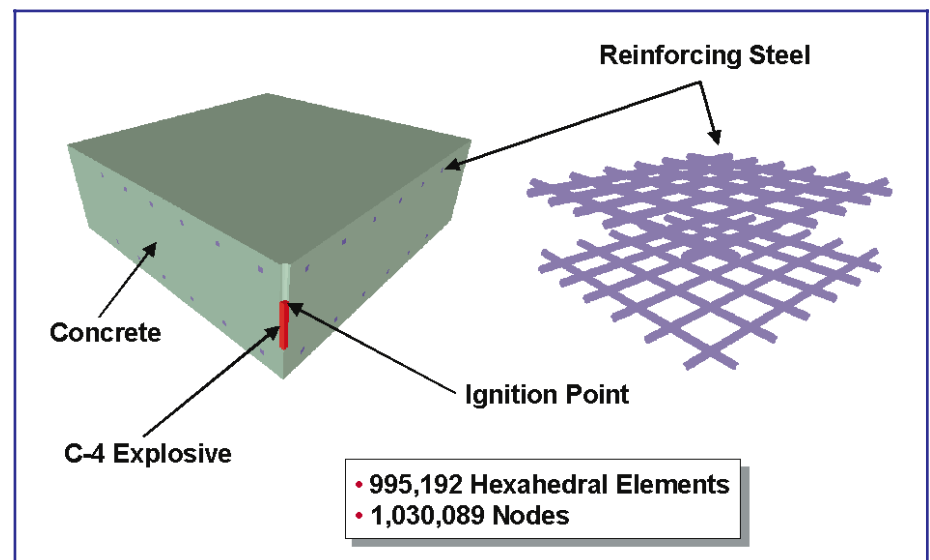


FIGURE 1. Model of the explosive detonation in a reinforced concrete slab.




FIGURE 2. METIS partitions for 128 processors.


by Professor Zdenek Bazant (and students) and the authors. The fundamental nature of the microplane model yields distinct advantages, as data are more accurately fit with simpler experiments, and greater confidence is provided for general applications over common three-dimensional constitutive theories. Previous experiences have demonstrated its accuracy for other applications. A major disadvantage is that the microplane model is computationally intensive (by more than an order of magnitude greater than typical elasto-plastic models). Leveraging parallel computing, however, its use with the large-scale finite element model depicted in Figure 1 was made reasonable. The scalability was excellent, as the analysis required about 8, 4, and 2 CPU-hours on 128, 256, and 512 processors, respectively, of the Cray T3E-1200 at the CEWES MSRC. As a result of the overlapping algorithm, communication time for partition interface data was insignificant, thus achieving the near perfect levels of parallel efficiency. Although the analysis

was performed with less analysis time on 512 processors, the turn-around times for the 128- and 256-processor runs were much faster. The 512-processor run required the entire machine and had to be scheduled appropriately.

Scientific visualization plays an important role in such simulations. Examination of model validity (e.g., material definitions, boundary conditions, partitioning) is crucial. Evaluation of predicted quantities is also improved by visualization (Figure 3). Display of deformed shapes showing scalar quantities (e.g., strains, pressure, damage) as the explosive event evolves provides the analyst with a better understanding of the structural response in such cases. With the assistance of the scientific visualization staff at the CEWES MSRC, visualization software was linked with the ParaAble output database. Despite the large model and frequent mirroring of the image for symmetry, typical display procedures (e.g., zoom, rotate, pan) were rendered in a matter of

seconds with simple click-and-drag mouse operations.

In summary, the parallel computational and visualization capabilities of the CEWES MSRC were invaluable to this project. Analyses that would require over 1,000 serial computing hours were performed in only a few hours on the HPC plat-

forms. Examination of large data sets predicted by these analyses was also rapidly made by advanced visualization software. These items greatly improved the productivity of the participants in conducting research activities important to our national defense. 

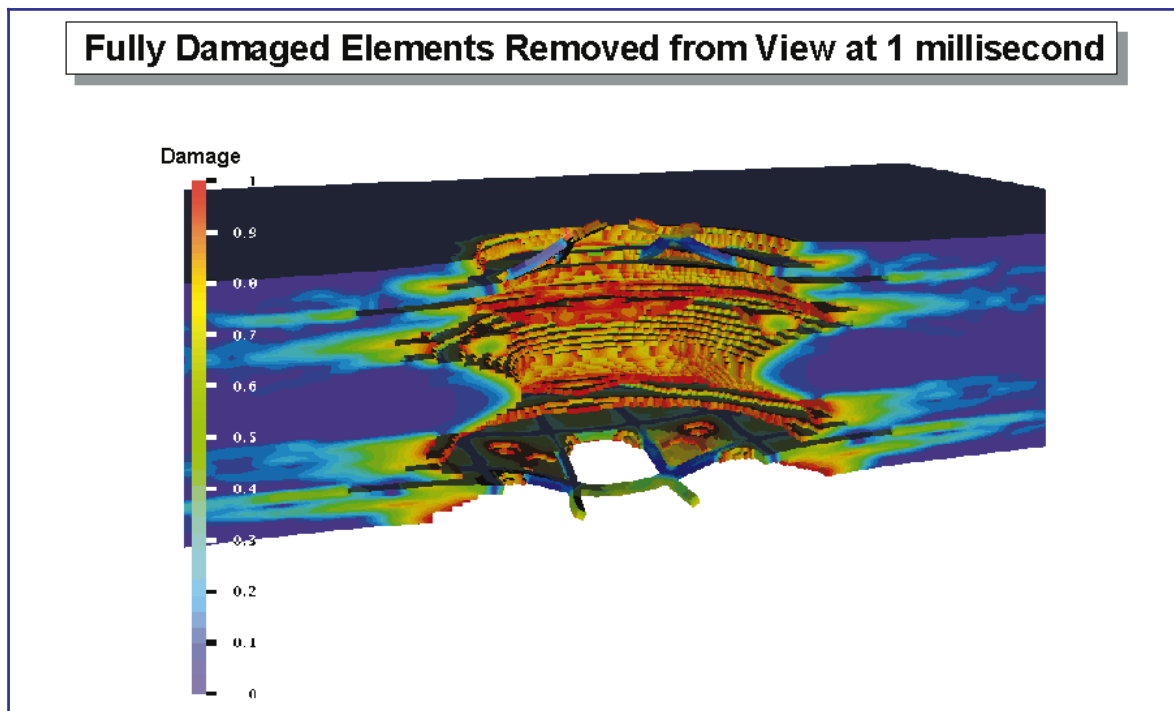


FIGURE 3. The color range from 0 (blue, undamaged) to 1 (red, completely damaged) indicates damage.

Dr. Danielson is a research associate professor of mechanical engineering at Northwestern University in Evanston, IL. Dr. Adley and Mr. Akers are research civil engineers in the Structures Laboratory at the U.S. Army Engineer Waterways Experiment Station in Vicksburg, MS.

The authors would like to thank Richard Strelitz, Ph.D., Richard Walters, and the rest of the CEWES MSRC scientific visualization staff for aiding in the creation of the images

DoD Challenge Project Studies Airborne Lasers

*Richard A. Strelitz, Ph.D.,
Joseph Werne, Ph.D.*

Lasers are invaluable in targeting, ranging, positioning, communications, and potentially as weapons. Effective use depends on the precision and consistency with which a coherent beam can be focused on a target. Scattering and refraction due to atmospheric effects can dramatically degrade an initially coherent beam, complicating tracking and focusing. To explore these effects, the Airborne Laser (ABL) Challenge project, sponsored by the High Performance Computing Modernization Program, models stratospheric turbulence and its impact on active optical tracking and the consequences for adaptive optics systems. The objective is to establish design parameters for a boost-phase, theatre-missile defense system utilizing a directed-energy laser weapon carried on a 747 aircraft. The ABL project has been allotted 100,000 node-hours on the CEWES MSRC Cray T3E in this year alone, as well as twice that amount at the NAVO MSRC and an equivalent amount on the IBM P2SC at the Maui High Performance Computing Center.

The ABL project considers two complementary facets of ABL design: 1) a directed energy (DE) component that concerns the modeling of optical propagation phenomena associated with turbulence in the upper tropo-

sphere and lower stratosphere, and 2) a geophysical effort that seeks to develop numerical simulations of that turbulence. The second component is sponsored by the Air Force Space Vehicles Directorate (VS). The DE thrust is based on the application of Maxwell's equations in a random medium. The results of this part of the project will provide quantitative estimates of the fundamental limits of accuracy of the ABL for a given characterization of turbulence. The second component of the project utilizes the classical Navier-Stokes equations to simulate conditions in the stratosphere to provide insights as to the nature and distribution of stratospheric turbulence.

The great challenge of the VS effort is to achieve sufficient separation of large- and small-length scales in the simulated turbulence so that a meaningful connection to stratospheric balloon and radar measurements can be made. Because reliable sub-grid models of stratified turbulence are not currently available, direct simulation of the fluid equations is required, and so-called "large-eddy" simulation (in which sub-grid turbulence processes are modeled) is not appropriate. Recent results^{1, 2} of the direct simulation offer explanation and allow interpretation of stratospheric turbulence measurements, promising both valuable input for the DE work and a

1 Werne and Fritts. (1999). "Stratified shear turbulence: Evolution and statistics," *Geophysical Research Letters*, 26, 439.


2 Werne and Fritts. (1999). "Turbulence and mixing in a stratified shear layer: 3D K-H simulations at Re=24,000." XXIV General Assembly of the European Geophysical Society, The Hague, 19-23 April 1999.

testbed for the development and improvement of theoretical descriptions of anisotropy in stratified and shear environments. In addition to providing fundamental design and evaluation tools for ABL, the DE and VS simulations have intrinsic scientific value and application to the Department of Defense (e.g., fundamental results on the strength and distribution of stratospheric turbulence and on the phenomenology of optical propagation through strong turbulence; evaluation of advanced tracking and adaptive optics concepts; and imaging, surveillance, and communication on nearly horizontal terrestrial paths).

The computational challenge of the turbulent simulation studies is to link together the small-scale and the large-scale phenomena using physics, and not untested “scaling theories.” This means obtaining accurate solutions with reliable time-stepping algorithms and large numbers of spectral modes to represent all scales of stratified shear turbulence. The laser-beam path is affected by thermal-gradient fluctuations, which determine the refractive index; too coarse spatial resolution invalidates the turbulence simulation by disallowing important anisotropic shear and stratification effects that impact beam focus. In addition, the VS problem is highly intermittent, with solutions exhibiting skewed and strongly non-Gaussian statistical distributions for the flow fields, again, requiring the inclusion of many length scales to properly describe the turbulence.

Visualization of the solutions alone can be challenging. For example, 32 MB of texture memory was required to create the image in Figure 1 (This image was created using VIZ, a three-dimensional volume-rendering software package from Reveal Soft-

ware LLC, Boulder, CO). This was the only way that the full volume could be used and still allow the user to change vantagepoint and representation in essentially real time. Complexity is an important issue in the data and its visualization as well. The topic of interest, turbulence, is evanescent and irregular, a sort of structured chaos. Depicting the distribution of velocity, vorticity, and the thermodynamic quantity potential temperature is difficult enough on a two-dimensional grid, within a small volume, or a steady-state system. When none of these conditions apply, the visualization rapidly grows in scope and complexity and can be done only on the largest machines. The data sets used in the ABL project are very large; the solution grid alone is 720x240x1440, with 32 bits of data at each node (or 1 GByte). In comparison, it should be noted that volumetric models like those encountered in medical or seismic tomographic imaging (e.g., computerized axial tomography or CAT scan) are often on the order of 256^3 at 8-bit precision, or nearly a factor of 64 less. Furthermore, the processing algorithms of tomography used to extract a three-dimensional structure from a series of cross sections are linear and easily parallelized while the nonlinear Navier-Stokes equations are not.

Scientific visualization is also crucial in depicting the three-dimensional nature of the data. The shear-layer dynamics begin as a two-dimensional phenomenon, but rapidly succumb to three-dimensional instability and subsequent turbulence and mixing. Without the ability to view the data in three dimensions by using volume rendering and selectable transparency, it would be difficult, if not impossible, to see how and why the turbulence patterns develop. 




FIGURE 1. Volume visualization of stratospheric turbulence. Yellow depicts the viscous dissipation, i.e., intense regions of small-scale variations in the velocity field. Blue shows the thermal dissipation where temperature fluctuations are most intense. Turbulent mixing is responsible for the absence of intense thermal fluctuations in the interior, while entrainment dynamics at the layer's edges maintain intense thermal fluctuations there (See back cover for larger image.)

Dr. Strelitz is a member of the scientific visualization team at the CEWES MSRC. Dr. Werne is a research scientist with Colorado Research Associates in Boulder, CO.

Amorphization and Fracture of Silicon Diselenide Nanowires

*Phillip Walsh,
Rajiv Kalia, Ph.D.,
Akihiro Nakano, Ph.D.,
Priya Vashishta, Ph.D.*

The study of nanowires is becoming more important due to their relevance to the field of nanotechnology. Magnetic arrays of nanowires could soon be used in data storage devices. Nanocomposites consisting of mixtures of nanowires (fibers) and nanoclusters are promising candidates for new materials combining the high chemical and temperature resistance of ceramics with the high ductility of metals. Nanowires will also be used for connecting components of micromachines as the scale of this technology grows smaller. For any of these applications, it is important to study the structural properties of nanowires. Theoretical and experimental efforts in this direction have therefore increased over the last few years.

The crystalline structure of silicon diselenide makes it an ideal prototype for nanowires and a convenient starting point for computer simulations. Its structure consists of exclusively edge-sharing tetrahedral units which form a one-dimensional chain-like structure along its c axis. The bonds between the chains are weak so that a nanowire may be formed by simply separating a number of chains from the bulk.

Computer simulations using the molecular dynamics (MD) method were carried out on 120 processors of the IBM SP at the CEWES

MSRC. The MD method consists of solving Newton's equations of motion iteratively for each particle in an N particle system under the influence of a potential determined by the positions of the remaining $N-1$ particles. This method is ideal for studying structural properties of materials from an atomistic perspective. The large amount of computing resources required to simulate nanometer scale materials ($N=1$ million to $N=100$ million atoms) makes DoD high performance computing (HPC) sites ideal for this type of simulation.

The simulated nanowire consisted of 1,204 chains with more than 4 million atoms, was initially 350 nanometers long, had a circular cross section, and had a diameter of 21 nanometers. The nanowire was thermalized and heated to 100 °K over 45,000 MD steps; strain was then applied along the length of the wire. A structural transformation occurred near 7 percent strain. The a-b unit cell structure, which had originally been rectangular with lattice constants $a = 9.669$ angstroms and $b = 5.998$ angstroms, transformed to a square structure with $a = b = 7.2$ angstroms. This resulted in an overall transformation of the cross section from a circular shape to an elliptical shape.

The failure process for the nanowire is interesting. At about 4.2 picoseconds (ps) after reaching a critical strain of 15 percent, some of the tetrahedral bonds in the outermost chains break, seeding a process of solid-state amorphization

The authors are members of the Concurrent Computing Laboratory for Materials Simulation at Louisiana State University in Baton Rouge, LA. Mr. Walsh is a research assistant with the Department of Physics and Astronomy. Drs. Vashishta and Kalia are professors in both the Department of Physics and Astronomy and the Department of Computer Science. Dr. Nakano is a professor in the Department of Computer Science.

which spreads from this initial point throughout the cross section of the wire and along its length. During amorphization, edge-sharing, intra-chain bonds break and chains cross-link to form corner-sharing tetrahedral units in a manner characteristic of amorphous silicon diselenide. This characterization of the amorphization process has been confirmed through bond angle and structure factor calculations performed on the amorphized region. The result is an amorphous region of the wire separated by two relatively weak crystalline-amorphous interfaces, which is where the final fracture occurs.

process at 4.2 ps after critical strain. The length of the wire in the figure is compressed by a factor of ten. The temperature of most of the wire remained at 100 °K (blue) while the region of the breaking bonds climbed to over 2000 °K (yellows and reds). Figure 2 shows the spread of the thermal spike at 5.7 ps, 7.2 ps, and 8.4 ps after critical strain. In Figure 3 (16.8 ps), the amorphous region extends throughout the cross section of the nanowire and continues to spread along its length. A similar analysis of the stress tensor components has also been done. The experience gained in doing these simulations




FIGURE 1. Local temperature distribution at 4.2 ps after critical strain. Temperatures range from 100 °K (blue) to over 2000 °K (red).





FIGURE 2. Local temperature distributions for (from left) 5.7 ps, 7.2 ps, and 8.4 ps after critical strain.

One way to follow the spread of local amorphization is to map the local temperature (i.e., local kinetic energy) as it changes with time. Breaking bonds results in a sudden release of large amounts of kinetic energy. Figure 1 shows the thermal spike that began the amorphization

will enable similar simulations of silicon nitride and silicon carbide nanowires to be performed at HPC sites in the immediate future. Simulations of structural properties of nanocomposite materials consisting of mixtures of nanowires and nanoclusters are already underway. 




FIGURE 3. Local temperature distribution at 16.8 ps after critical strain.

Design of New High Energy Fuels


Mark S. Gordon, Ph.D.,
Graham D. Fletcher, Ph.D.

For several years, the U.S. Air Force has had an active research effort that focuses on the design and delivery of new high energy fuels (the HEDM project for high energy density materials). Such potentially useful new fuels must have a specific impulse (I_{sp}) that is significantly better than those of existing fuels, such as liquid oxygen/liquid hydrogen or monopropellants such as hydrazine (N_2H_4). A large specific impulse, measured in units of time, is provided by fuel specimens that have a small mass and a large energy content measured by the heat of the reaction that would occur in the en-

gine chamber. Not reflected in the I_{sp} , but also important, is that any energy barrier separating the proposed new fuel from any products must be at least 30 kilocalories per molecule (kcal/mol).


The Air Force Office of Scientific Research HEDM program includes efforts in both theory and experiment. Indeed, one of the strengths of the program is the highly interactive nature of the participating scientists. The Gordon theoretical chemistry research group at Iowa State University has participated in the HEDM program for several years. Two years ago, Dr. Robert Schmitt, SRI International, proposed the molecule shown in Figure 1 as a potential HEDM candidate. At that

FIGURE 1. The molecule shown has been identified by Robert Schmitt (SRI) as a possible new high-energy species. As part of the Air Force High Energy Density Materials (HEDM) program, the molecular structure and energetics of this compound have been investigated using second order perturbation theory (MP2) and a large atomic basis set. These calculations would be impossible without the recent development of highly scalable code for MP2 gradients and without access to the CEWES T3E. Preliminary indications and ongoing calculations suggest that this molecule may be a viable fuel candidate.



time, the molecule had not been synthesized. Having 10 nitrogen atoms connected to each other is unusual, and such a structure is expected to be highly energetic. Therefore, the two parties entered into a collaborative effort in an attempt to characterize this interesting compound.

The first step was to predict its molecular structure and its heat of formation. This information is sufficient to provide an estimate of the I_{sp} . Computational techniques were used to determine the geometry and vibrational frequencies using Hartree-Fock theory and a reasonable basis set. These preliminary calculations predicted a heat of formation for the compound of 457 kcal/mol. This lead to an I_{sp} for the substance as a monopropellant of 329 sec, which is much higher than the I_{sp} of 240 for hydrazine. Because the expected products of the proposed compound are N_2 and CO_2 , it is also environmentally much friendlier than hydrazine.¹ The proposed fuel appears to be quite promising.

The next step is to improve the accuracy of the predictions, by repeating the calculations at a higher level of theory, namely second order perturbation theory (MP2). This entails a demanding series of calculations because the molecule is so large. Fortunately, the team has recently developed and implemented into the electronic structure code GAMESS² (General Atomic and Molecular Electronic Structure System) a highly scalable code for MP2 energies and gradients.³ This code, which is available on the CEWES T3E, allows researchers to execute these complex calculations. Indeed, they are currently in progress. Typical runs using 128 processors last 129,600 sec, and to date the project has used over 30,000 node-hours on the CEWES T3E. Once the calculations are completed, and assuming the predictions are not altered greatly from those obtained at the Hartree-Fock level of theory, the project will proceed to investigate the energy barriers separating this species from various breakdown products. 

Drs. Gordon and Fletcher are professors of chemistry at Iowa State University in Ames, IA.

1 The I_{sp} for the fuel LOX/RP1 is only 300 sec.

2 M.W. Schmidt, K.K. Baldrige, J.A. Boatz, S.T. Elbert, M.S. Gordon, J.H. Jensen, S. Koseki, N. Matsunaga, K.A. Nguyen, S. Su, T.L. Windus, M. Dupuis, and J.A. Montgomery, Jr., "The General Atomic and Molecular Electronic Structure System", *J. Comp. Chem.*, 14, 1347 (1993).

3 G.D. Fletcher, M.W Schmidt, and M.S. Gordon, "Developments in Parallel Electronic Structure Theory", *Adv. Chem. Physics*, in press.

Topside Communications Get Boost from CEWES

Charles W. Manry, Jr., Ph.D.

The topside of a modern U.S. Navy surface combatant is a sophisticated assortment of weapons, electromagnetic (EM) radiators, and other hardware. Large numbers of antennas, transmitters, and receivers are required to meet radar, electronic warfare, information warfare, and communication requirements. An increasing inventory of EM systems is constantly being added to meet requirements for more communications capability with greater imagery and data transfer capacity (Figure 1). These requirements, as well as aggressive signature reduction goals, create new demands for Integrated Topside Design (ITD) for Naval surface combatants. Answering these demands is critical, because the combat effectiveness of Navy ships is limited by the ability to provide advanced Integrated Topside Designs.

Designing and optimizing a complex electromagnetic environment is slow, expensive, and error-prone. Eighty percent of the “affordability” decisions are made before a detailed design is available for a new platform. It is only through the application of concurrent engineering and its associated simulation-based design environment that 21st Century Integrated Topside Designs can be implemented. Thus the need for advanced simulation, visualization, and optimization tools that exploit high performance computing (HPC) are critical to the required design tools (Figure 2). The goal of the Electromagnetic Interactions GenERalized (EIGER) development was to create a computational EM framework for the analysis of 21st Century Integrated Topside Designs. The EIGER framework that has been developed incorporates a wide variety of numerical and analytical




FIGURE 1. USS Radford with Advanced Enclosed Mast/Sensor (AEM/S) mast. (Courtesy U.S. Navy and SSCSD).

techniques into an efficient, scalable EM analysis code of unprecedented versatility. Because of its careful initial design, EIGER has achieved, in a very short time, a relatively mature status as a general EM modeling tool.

In order to effectively utilize HPC resources in the application of EIGER applied to ITD problems, a DoD MSRC Challenge award of 202,000 CPU-hours on the CEWES IBM SP system was awarded. Both Space and Naval Warfare Systems Center, San Diego (SSCSD) and CEWES MSRC personnel are working together to execute simulations that show the capabilities of HPC and EIGER used to solve ITD problems (Figure 3). EIGER development is supported by multiple

sponsors, including the Department of Energy, and is a DoD Common High Performance Computing Software Support Initiative (CHSSI) project.

The SSCSD supports the command, control, communications, computers, intelligence, surveillance, and reconnaissance missions of the Navy (<http://www.spawar.navy.mil/sandiego/welcome.page>). In support of these missions, SSCSD has maintained a unique capability to support the Navy Integrated Topside Design efforts (<http://bobcat.spawar.navy.mil/d85/index.htm>). Advanced EM modeling has been and will continue to be critical to the success of these technology initiatives. 



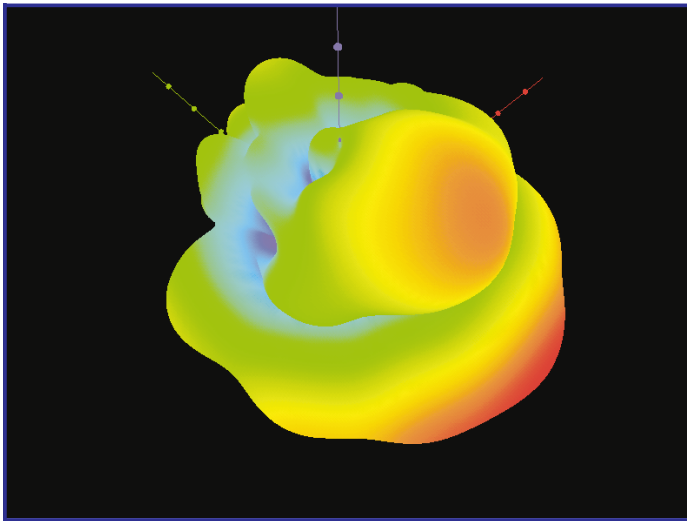


FIGURE 2. Artist conception of amphibious assault ship (LPD 17 class) with AEM/S masts. (Courtesy U.S. Navy and SSCSD).


Dr. Manry is an engineer in the Electromagnetics and Advance Technology Division at the Space and Naval Warfare Systems Center in San Diego, CA.



Current flow on ship displayed using “jet” color map of currents magnitudes. Blue represents the smallest current values and red represents the highest. The source of energy is the mid-starboard high frequency whip broadcasting at 10.0 MHz



Antenna broadcast pattern for the antenna shown above. The distance of the surface from the center is the strength of this pattern. The color range from purple (lowest) to red (highest) shows broadcast power.



The color map for this image represents the phase shift of the broadcast signal.

FIGURE 3. This sequence took slightly over 1 wallclock hour using 128 processors and approximately 12 GBytes of memory on the CEWES IBM SP for this one frequency. There were 21 sequences in the set; therefore, it took over 21 wallclock hours to perform.

Portable Batch System Aids Common User Environment

Portable Batch System

James Patton Jones

With the installation of the Portable Batch System (PBS) on the Cray T3E, the CEWES MSRC moves a step closer toward its goal of implementing a common user environment across all MSRC systems. PBS, developed at NASA Ames Research Center, is a batch queuing system which systematically sorts and executes user jobs on available computing resources. Specifically, it was designed to replace the Network Queuing System (NQS), also developed by NASA, and to provide a seamless environment for job processing among disparate computing systems. It was this ability for seamless integration that caught the attention of the High Performance Computing Modernization Program in 1997. In July of that year, CEWES deployed PBS on its IBM SP. Subsequently, PBS was installed on the SGI PowerChallenge and Origin2000 systems.

With the same queuing system on the different supercomputers, users are able to focus on science, rather than learning a different way of running their jobs on each system. In addition, PBS allows a user to submit a batch job to any system on which he has a valid account, further simplifying its use. Furthermore, from a single window of the PBS graphical user interface, xpbs, users can view the status of all their batch jobs wherever they may be running within the entire site or even multiple sites, such as in the MSRC MetaCenter, a joint ASC/CEWES MSRC project

to support job queuing across sites (see CEWES MSRC Technical Journal, Fall 1998).

The CEWES Cray T3E is the latest system to operate under PBS. Earlier this year, the CEWES MSRC engaged in efforts to better integrate PBS with the T3E and to implement specific local job scheduling policies. PBS replaces Cray's Network Queuing Environment (NQE) extensions to NQS. Ironically, last summer Cray announced that it would be discontinuing support for NQE, which prompted many sites using Cray systems to consider PBS. The new CEWES MSRC enhancements make it even more desirable to switch to PBS on the T3E.




FIGURE 1. The PBS graphical user interface, xpbs, offers a single interface to different supercomputers which allows the user to select which queues and jobs to display, while providing point-and-click batch job submission.

Portable Batch System

Mr. Jones is a systems analyst for MRJ Technology Solutions and is currently serving as a NASA liaison at the CEWES MSRC.

To aid in the migration from NQS to PBS on systems such as the Cray T3E, PBS provides the 'nqs2pbs' conversion utility which translates an NQE/NQS batch job into a PBS job. The resulting batch job will actually be valid for both NQE/NQS and PBS, allowing the user to submit the same job to either batch system, providing fur-

ther flexibility in transitioning to PBS. PBS is now distributed and supported by MRJ Technology Solutions, the NASA research and development contractor which developed PBS. Information on using and acquiring the PBS software package is available on the World Wide Web (<http://science.nasa.gov/Software/PBS/>). 

VAMPIR

VAMPIR Comes to CEWES

Clay P. Breshears, Ph.D.

VAMPIRtrace and VAMPIR are two programming tools that work together to measure the performance of parallel programs and to display this data in various graphical formats. To use the tools, the VAMPIRtrace message passing interface (MPI) profiling library is linked with a user's code. When executed, the instrumented MPI calls generate trace data about the time a call was initiated and how long the function call lasted. The library hooks into the MPI profiling interface and achieves low tracing overhead by keeping the output data within the local memory of the process. Upon completion of the code execution, a post-processing step gathers all trace data from each processor and writes it to disk in a single file. Even when run on multiple processors, the effects of distributed clock drifts are corrected automatically in order to keep a high correlation between each separate processor's trace data.

Performance data is automatically gathered for each individual MPI call used in a program. In a code, especially long-running codes with many calls to MPI routines, a user may only be interested in a particular portion of the program. Tracing may be enabled and disabled at the discretion of the user by inserting calls to VAMPIRtrace application programming interface (API) functions. Thus, performance analysis can be concentrated on specific areas of a code. This will also control the amount of trace data generated (which can easily become large for programs that run on many processors or for a large amount of time). Also, since any parts of the code that are not MPI routines are considered the same, the API provides calls for defining, starting, and stopping user-defined activities. In this way, performance data can be gathered on specific user-written routines besides the default MPI activity.

The VAMPIR visualization tool analyzes the trace data generated by

VAMPIRtrace. There are three main types of visualization displays (Figure 1):


- The timeline display shows process states over time and communication between processes by drawing lines to connect the sending and receiving process of a message. Message patterns and relative amounts of time spent waiting for messages, completion of global communication routines, or other facets of execution can be easily seen with this display.
- The statistics display shows the cumulative statistics for the complete trace file in pie chart form for each process. The percentage of execution time taken up by all communication or one particular MPI routine can be shown with this display.
- The process state display shows every process as a box and displays the process state at a selected point in time. This display allows the user to identify how many processes are executing MPI calls or user code or are standing idle.

Details and features of each of these displays can be configured by the user with pull-down menu options. Such customizations can be saved to a configuration file that is read each time VAMPIR is started.

VAMPIR includes flexible filter operations to reduce the amount of information to be displayed, as well as rapid zooming and forward/backward motion in time to allow the user to focus quickly on arbitrary time intervals. Thus, the user can easily identify performance bottlenecks at the appropriate level of detail. Information displayed includes message-passing, collective communication, and application subroutine execution.

For help in using VAMPIRtrace and VAMPIR on any of the the CEWES MSRC HPC platforms, contact the CEWES MSRC Customer Assistance Center by telephone at 800-500-4722 or by e-mail at info-hpc@wes.hpc.mil.

Dr. Breshears is a research scientist at the Center for Research on Parallel Computation and the PET on site lead for scalable parallel programming tools at the CEWES MSRC.



VAMPIR and VAMPIRtrace provide three main displays: timeline, statistics, and process state.

Repository of Scalable Parallel Programming Tools For CEWES MSRC Platforms

*Shirley V. Browne, Ph.D.,
Clay P. Breshears, Ph.D.*

High performance computing (HPC) users interested in parallel programming tools can take advantage of the scalable parallel programming tools (SPP Tools) software repository developed as part of a Programming Environment and Training (PET) effort at the CEWES MSRC. The SPP Tools repository lists and summarizes nearly forty tools including benchmark programs, distributed processing tools, math libraries, and parallel processing tools such as debuggers, performance analyzers, and parallel I/O systems. The repository is located at http://www.nhse.org/rib/repositories/cewes_spp_tools/catalog. Some of the entries are vendor or commercial tools while others have been developed by government-funded research projects. For selected tools, there are site-specific usage information and tutorials as well as discussion forums.

Examples of Tools

An example of a tool included in the SPP Tools repository is the TotalView debugger, which is available at the CEWES MSRC in two versions: TotalView for the IBM SP and SGI/Cray Origin2000 from Etrus, Inc., and Cray TotalView for the Cray T3E. Instructions for using the interactive TotalView debugger with CEWES MSRC

queuing systems and compilers are provided along with links to a web-based tutorial and hands-on exercises.

Another tool listed in the repository is the VAMPIR performance analysis tool from Pallas in Germany, which is installed on all CEWES MSRC HPC platforms. Instructions for linking an application with the VAMPIRtrace profiling library, running the instrumental program, and viewing the resulting trace data using the VAMPIR visualization tool are provided in the form of a step-by-step tutorial. Figures 1 and 2 show screen shots from using VAMPIR to analyze and improve performance of the Icepic code from the Radio Frequency Weapons Prototyping DoD Challenge project.

User discussion forums have been set up for both TotalView and VAMPIR to provide a mechanism for users to post comments, questions, and experiences. The forums are intended to allow users to learn from each other and to facilitate communication between users and the tool developers.

Information about platform-specific debuggers and performance analysis tools is included, along with hints and tricks about how to use these tools in specific situations, such as with MPI programs. A good strategy for a new user would be to use a cross platform tool such as TotalView or VAMPIR first, and then

Parallel Programming Tools

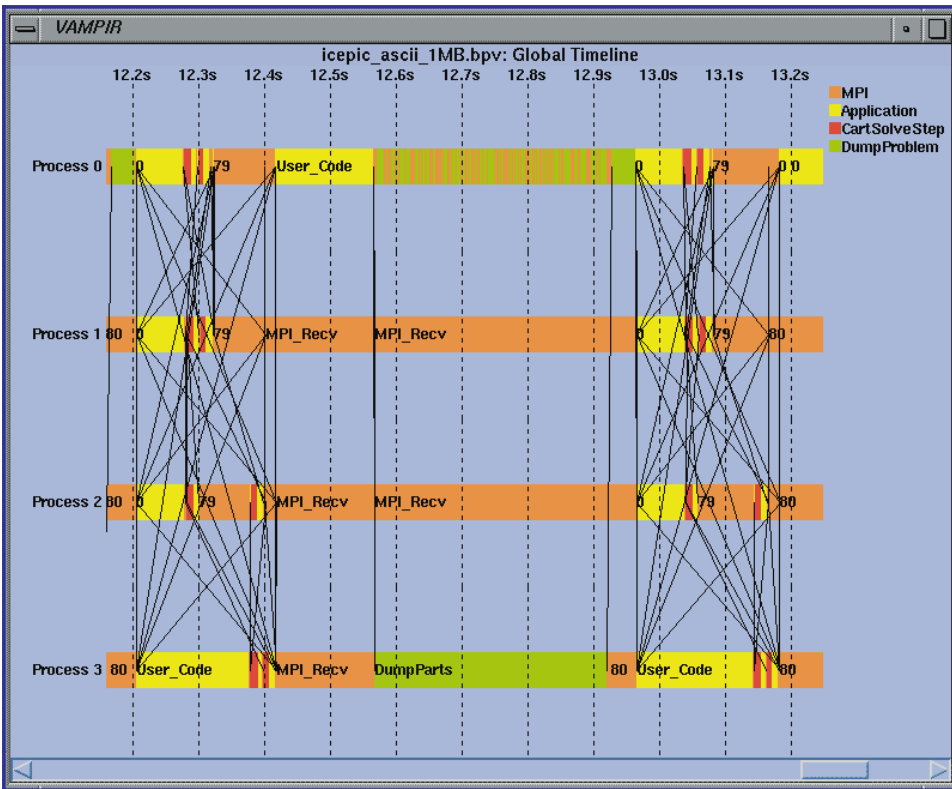


FIGURE 1. VAMPIR visualization of Icepic execution before algorithmic changes (Image courtesy of Jerry Sasser, Ph.D., Air Force Research Laboratory, Kirtland AFB, NM).

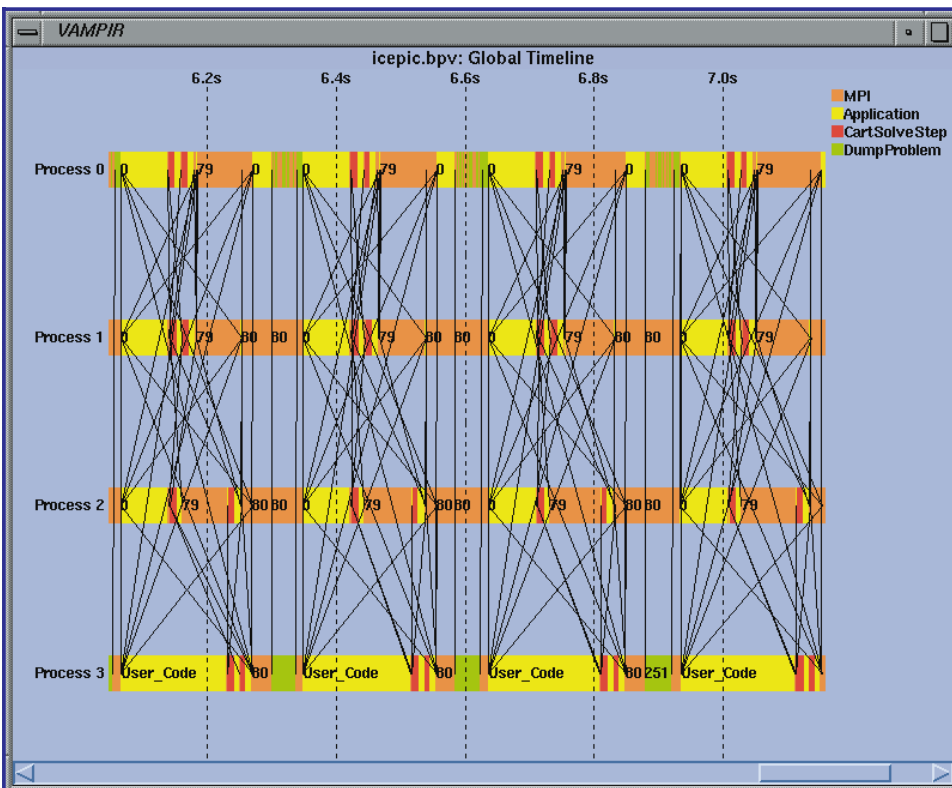


FIGURE 2. VAMPIR visualization of Icepic execution after algorithmic changes (Image courtesy of Jerry Sasser, Ph.D., Air Force Research Laboratory, Kirtland AFB, NM).

Parallel Programming Tools

Dr. Browne is Associate Director of the Innovative Computer Laboratory at the University of Tennessee in Knoxville. Dr. Breshears is a research scientist at the Center for Research on Parallel Computation and the PET on site lead for scalable parallel programming tools at the CEWES MSRC.

switch to a platform-specific tool if platform-specific features are needed.

The math libraries section of the catalog includes the LAPACK and ScaLAPACK linear algebra libraries for shared memory and distributed memory machines, respectively, along with information about the vendor versions of these libraries. Users are urged to use the tuned vendor version of a routine where available. Additional routines not yet included in the vendor versions are available from the research versions. Other available math libraries include SuperLU_MT, a multi-threaded, sparse matrix solver package.

Future Tools

In addition to the tools already installed, the catalog gives a preview of coming attractions. For example, the Virtue virtual reality environment for collaborative analysis of large-scale performance data is currently being


installed at the CEWES MSRC and will be available soon (Figure 3). Plans are to provide a converter from the VAMPIRtrace format to the format understood by Virtue, in order to allow large-scale trace data not easily viewable by VAMPIR to be visualized using more scalable three-dimensional (3-D) representations. For example, a 3-D time tunnel

display can be used for viewing state changes and communication behavior of parallel programs.

Deployment Information

Linked to the SPP Tools repository is a software deployment grid (http://www.nhse.org/rib/repositories/cewes_spp_tools/catalog/grid.html), which lists software packages and their deployment status with respect to CEWES MSRC machines. The grid allows users and system administrators of MSRC systems to keep track of the various versions and locations of software installed on various machines. By simply clicking on the appropriate grid entries, users can immediately find details about deployment, as well as instructions on how to use the software and local contact information for support questions.

The SPP Tools repository and deployment grid have been implemented using the Repository in a Box (RIB) software developed by the federally-funded National High-performance Software Exchange (NHSE) project (<http://www.nhse.org>). RIB is being used to set up an interoperable network of software repositories at the DoD MSRCs, as well as at other government and academic high performance computing sites.

Users are encouraged to contact the CEWES MSRC with suggestions about additional tools or information that they would like to see made available. Users can contact the CEWES MSRC Customer Assistance Center by telephone at 800-500-4722 or by email: info-hpc@wes.hpc.mil. 




FIGURE 3. Virtue time tunnel and call graph displays of a parallel message-passing application. (Image courtesy of Daniel Reed, Virtue project lead at University of Illinois-Urbana/Champaign.)

CEWES Developing Adaptive HPC Software Solution

Alan K. Stagg, Ph.D.

One of the most challenging problems in the field of computational mechanics is obtaining sufficiently accurate solutions at reasonable cost. The central issue is the requirement to numerically capture detailed physics that may be localized within the computational domain. Examples of such local phenomena that require enhanced grid resolution include shock waves in compressible flow and dynamic concentration fronts in groundwater flow. A popular strategy is to locally refine the grid based on the features of interest and to remove grid resolution (coarsen) where it is no longer required. The advantage of this approach is that small-scale features can be captured without using a costly, fine grid over the entire problem domain.

Adaptive grid schemes of this type have been popular for a number of years; however, they are difficult to implement on scalable, parallel architectures, and no general purpose libraries are available. The issues related to implementing such schemes on parallel systems are just now being addressed in the research community, and much work is needed to identify the best methods. Researchers at CEWES have recently developed an innovative approach that simplifies the inherent difficulties associated with parallelizing these adaptive grid schemes.

The CEWES MSRC, located in the Information Technology Laboratory (ITL), is collaborating with the


Coastal and Hydraulics Laboratory (CHL) at the U.S. Army Engineer Waterways Experiment Station to develop and demonstrate this advanced technology in the Common High Performance Computing Software Support Initiative (CHSSI) code ADH, a state-of-the-art, multidisciplinary flow code. Project leaders are Alan Stagg (ITL) and Jackie Hallberg (CHL). The ADH code is serving as a preliminary testbed for development and evaluation of the parallel adaptation schemes that have been designed to support large-scale calculations of critical interest to the DoD. The new grid software is enabling refinement and coarsening for both tetrahedral and triangular grids on parallel architectures that support MPI message passing. Software development has been completed, and testing with two- and three-dimensional grids is underway. Preliminary results indicate high efficiency of the adaptive grid algorithm. A locally refined grid generated with the software is shown in Figure 1. Here an error indicator has been used to mark elements in the lettered regions, and the grid has been refined repeatedly to increase resolution in these areas. Colors represent the levels of refinement during which triangular elements are added. Dark blue elements are original elements, while red and white elements have been added in the last steps of refinement.

The potential benefits of this advanced technology are multifold. Reductions in computing time by one or two orders of magnitude are possible compared to using a

Adaptive HPC Software

Dr. Stagg is a computer engineer at the CEWES MSRC.

uniform, fine grid to capture flow features. With the parallel refinement strategy, solution time for complex problems could be reduced from months to days, or even hours. In addition, this technology will provide DoD users with the capability to solve problems with very high local resolution, mak-

ing possible calculations that would otherwise be intractable. These efforts continue to demonstrate the CEWES MSRC leadership and commitment to providing leading-edge software solutions for attacking the nation's most challenging computational problems. 




FIGURE 1. Locally refined grid based on an “ITL” error indicator.

Recently Published Reports:

- Bova, S. W., and White, J. B., III. "Where's the Overlap? An analysis of Popular MPI Implementation."
- Bova, S. W., Breshears, C. P., Cuicchi, C. A., Demirbelik, Z., and Gabb, H. A. "Dual-level Parallel Analysis of Harbor Wave Response Using MPI and OpenMP."
- Browne, S., Ho, G., and Mucci, P. "PAPI: Portable Interface to Hardware Performance Counters."
- Fahey, M., and Nagle, D. "Cray Fortran pointers vs. Fortran 90 pointers and porting from the Cray C90 to the SGI Origin2000."
- Heiland, R., and Baker, M. P. "Co-processing: Experience with CUMULVS and pV3."
- Mastin, W. "1998 CEWES MSRC PET Training Activities."
- Nagle, D. "Migration issues: Moving from serial to parallel processing."
- Nagle, D. "The effect of Fortran 95 PURE and ELEMENTAL procedures on parallel execution."
- Nagle, D. "Using Fortran procedure interfaces."
- Shih, A. M., and Baker, M. P. "Management Strategies for Scientific Data: Assessing Utility of HDF for CEWES MSRC users."
- Wheeler, M. F., Dawson, C., Li, J., Parr, V. "UT-PROJ: The University of Texas Projection Code for Computing Locally Conservative Velocity Fields."
- White, J. B. "Reading sequential unformatted Cray C90 files on an SGI Origin."
- White, J. B., and Bova, S. W. "Where's the overlap? Overlapping communication and computation in several popular MPI implementations."
- Willhoite, B., and Nagle, D. "Using Fortran 90 features for cache optimization."

CEWES MSRC 1999 Training Schedule*

July

Open MP and Pthreads
Enlight for CFD and CSM Applications

August

Workshop on Parallel Algorithms
Distance Training Workshop

September

How to Use Parallel Linear Algebra Library Routines
Advanced Performance Optimization Tools and Techniques

October

Grid Generation and Adaptive Grids
IBM POWER3 SP Parallelization Workshop

December

Using the SGI Origin2000 for Code Development and Analysis

* Additional courses may be offered. Please check the CEWES MSRC web page at <http://www.wes.hpc.mil>

For more information about the reports or planned training courses listed above, please contact the CEWES MSRC Customer Assistance Center by telephone at 1 800 500 4722 or by e mail at info_hpc@wes.hpc.mil. If you would like to attend one of our training courses or if your organization has specific training needs, please let us know.



For further information, contact:


CEWES MSRC
Customer Assistance Center

Web site

<http://www.wes.hpc.mil>

E-mail: info-hpc@wes.hpc.mil

Telephone: 800-500-4722
800-500-HPCC



STRATOSPHERIC TURBULENCE. These images show results from a DoD High Performance Computing Challenge Project in support of the AirBorne Laser (ABL) program. Scattering and refraction due to atmospheric turbulence can dramatically degrade an initially coherent laser beam, complicating ABL system performance. The images indicate regions of intense density fluctuations (blue) along with regions of intense turbulent mixing (yellow) in a stable stratified region of the atmosphere experiencing wind shear. Side and perspective views are shown. The three-dimensional snapshots shown are at times of roughly 7, 11, and 18 minutes. The DoD simulations represent the largest and highest-resolution computations of stratified shear turbulence yet conducted. Results are used to improve understanding and interpretation of turbulence measurements in the stratosphere so that ABL design can be facilitated. (See article on page 4.)

U.S. Army Engineer Waterways Experiment Station
ATTN: CEWES-IH/CEWES MSRC
3909 Halls Ferry Road
Vicksburg, MS 39180-6199