ERDC MSRC RESOURCE

Major Shared Resource Center

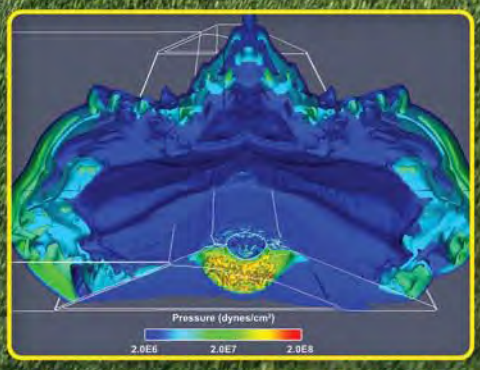

Inside This Issue:

Computational Studies of Cratering on Earthen Dams Using HPC

HIGH PERFORMANCE COMPUTING FOR THE WARFIGHTER.

FALL 2008

from the director . . .

Dave Stinson
Acting Director, ERDC MSRC

In looking through the articles for this issue of the *Resource*, I am reminded of how critical high performance computing (HPC) is to the development of technology and to our Nation, which is so dependent on this technology. The breakthroughs in science that are possible with the computational power we now have available are mind-boggling. With the acceptance of Jade, our Cray XT4, in July, the Engineer Research and Development Center Major Shared Resource Center (ERDC MSRC) has around 120 peak teraFLOPS of computing on the floor (see article on Jade). This kind of computational power is for the first time allowing us to design new materials by piecing together molecular components.

The material science research going on at ERDC involving carbon nanotubes has been identified as being the "Major Fundamental Research Accomplishment" for ERDC. The advances from this research in the areas of advanced composites, woven fabrics, and textiles with military and industrial applications can lead to the next fundamental shift in how we all work and live. I hope you enjoy reading about this in the article on new materials.

Realizing the importance of HPC to the Department of Defense (DoD), the Army, and the Corps of Engineers, the ERDC MSRC has committed to putting an infrastructure in place to support HPC for the long term. Investments in space and scalable cooling and power have positioned the MSRC to support virtually any HPC requirement in support of the High Performance Computing Modernization Program (HPCMP) for the foreseeable future. Our focus at the ERDC MSRC continues to be on providing capability for the largest science and engineering problems in the DoD. The article on ERDC infrastructure highlights some of the work being done in this area.

As I write this article, the advertisement for permanently filling the ERDC MSRC Director position is finally on the street. So, by the time you read this article, the ERDC MSRC should have a new director. I appreciate the opportunity that I have had as Acting MSRC Director for the past year and half to work with the HPCMP and Shared Resource Center staffs and to serve the Information Technology Laboratory, our MSRC staff, and users of our systems. I really appreciate the value of a great staff and the fantastic support they have given me, and I am looking forward to what the future will bring.

Dave Stinson
Acting Director, ERDC MSRC

About the Cover: Artistic rendering of a computed crater on an earthen dam. Back cover: Isosurface of pressure during the crater-forming explosion. For related story, see page 2.

Contents

Computational and Experimental Studies of Cratering on Earthen Dams

By Byron J. Armstrong, Dr. Stephen A. Akers, Dr. Gordon W. McMahon, and Denis D. Rickman, ERDC Geotechnical and Structures Laboratory

Acts of terrorism in the United States and across the world have shown the necessity of protecting high-value assets in the civilian as well as the military community. Protecting these assets not only includes the direct protection of structures and people, but also sometimes requires protection of resources whose failure could drastically affect both. One such scenario would be the attack on an earthen dam on a reservoir near a city.

The Geotechnical and Structures Laboratory (GSL) at the Engineer Research and Development Center (ERDC) is conducting a two-phase research effort to (1) examine the vulnerability of earthen dams to terrorist attacks and (2) to propose methods of mitigating the effects of such attacks. This article concentrates on the first part of this effort concerning the response of an earthen embankment to high-explosive attacks.

The GSL is using a combination of scaled experimental tests and computer simulations to examine the effects of attacks on earthen embankments. One phase of the experimental program consists of a series of six scaled tests conducted at Fort Polk, Louisiana. Two delivery methods are considered for these tests. The first consists of a crest attack with a vehicle-borne improvised explosive device (VBIED). The second scenario represents a waterborne VBIED. Two experiments were conducted for the first scenario and four experiments for the second. For each test, a reservoir was constructed and filled with water and a scaled bomb placed either on the embankment crest or in the water near the embankment. A limited amount of instrumentation (stress gages and accelerometers) was placed in the earthen embankment. After each bomb was detonated, the resulting crater was surveyed both along the

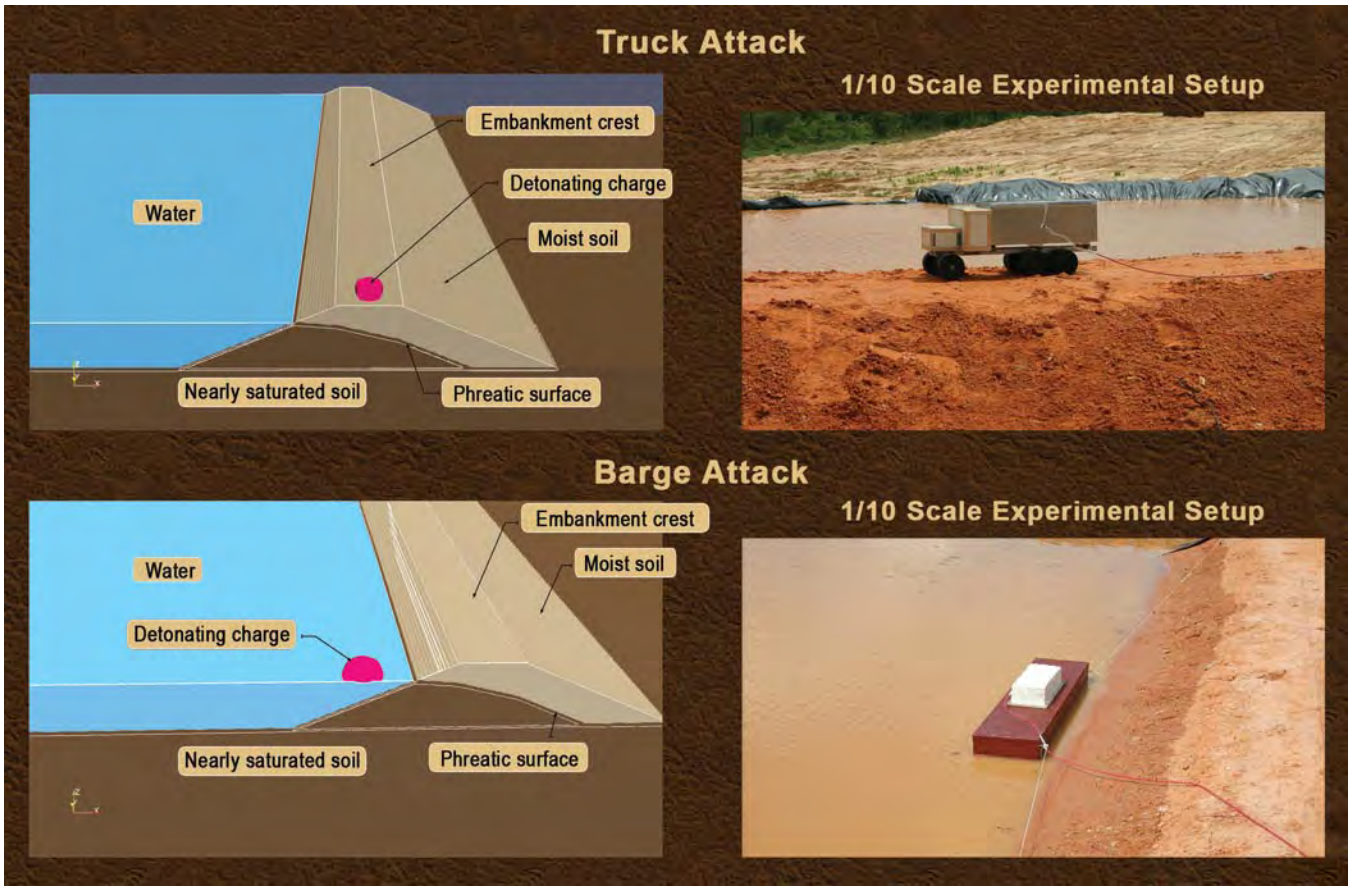

Figure 1. Experimental and calculational configurations

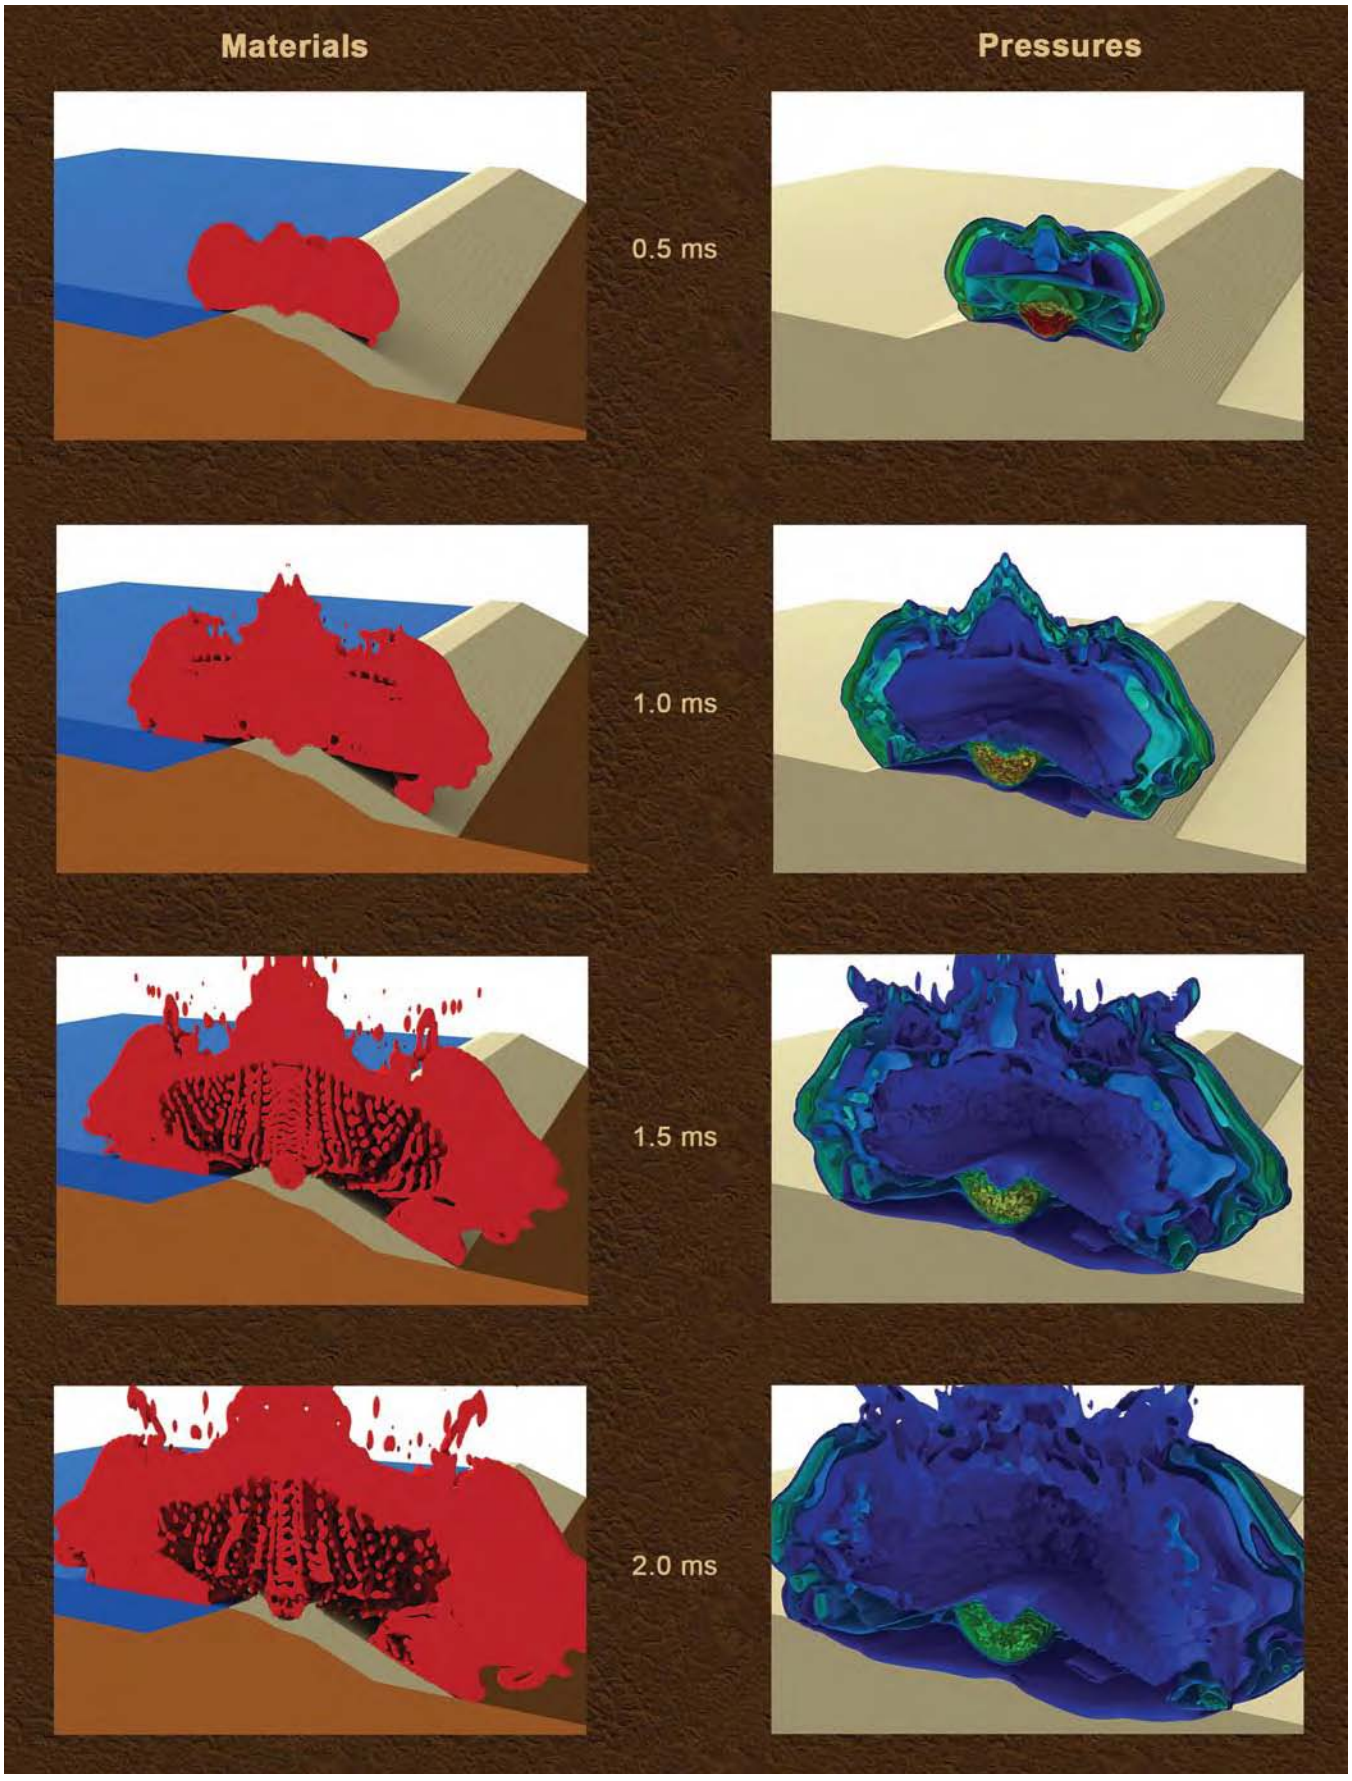

Figure 2. Early material displacements and loadings from truck attack

center of the crest of the dam as well as along a cross section going through the original charge location. For some of the waterborne VBIED detonations, additional profiles were surveyed closer to the location of the explosive. Soil properties for both the moist soil (above the phreatic surface) and nearly saturated soil (below the phreatic surface) were determined by laboratory testing.

Two of the numerical simulations replicating scaled experiments are presented in this article (Figure 1). The first simulation represents a VBIED attack against the crest of an earthen embankment, and the second represents an attack of a waterborne VBIED at a location near the embankment. These calculations were performed to examine the effectiveness of first-principle codes in predicting crater formation and to validate available first-principle codes and constitutive models against the experimental results obtained from the Fort Polk tests. The calculations concentrated on modeling the response of an earthen embankment to the high-explosive attack. This involved explicitly modeling the bomb detonation, the shock wave and high-pressure expansion of the gaseous detonation products, the interaction of the shock wave and detonation products with the soil and water near the bomb, and the resulting displacement of the soil forming the crater.

The numerical simulation of cratering is a challenging problem that pushes the capabilities of many of the current computational codes, material models, and experience of the analysts. Accurate numerical modeling of crater formation must address both the high-pressure, short-duration loading of the initial shock wave and expanding detonation gas as well as the later time, slower moving soil as the final crater forms. During the early time loading, modeling the detonation physics and initial crushing of the soil are dominant. Later in time, the lower stress and lower velocity regions, where material strengths dominate, appear to control the actual final formation of the crater. This requires the calculation to capture the microsecond loading of early times as well as the millisecond loadings and motions during crater formation. Both the memory

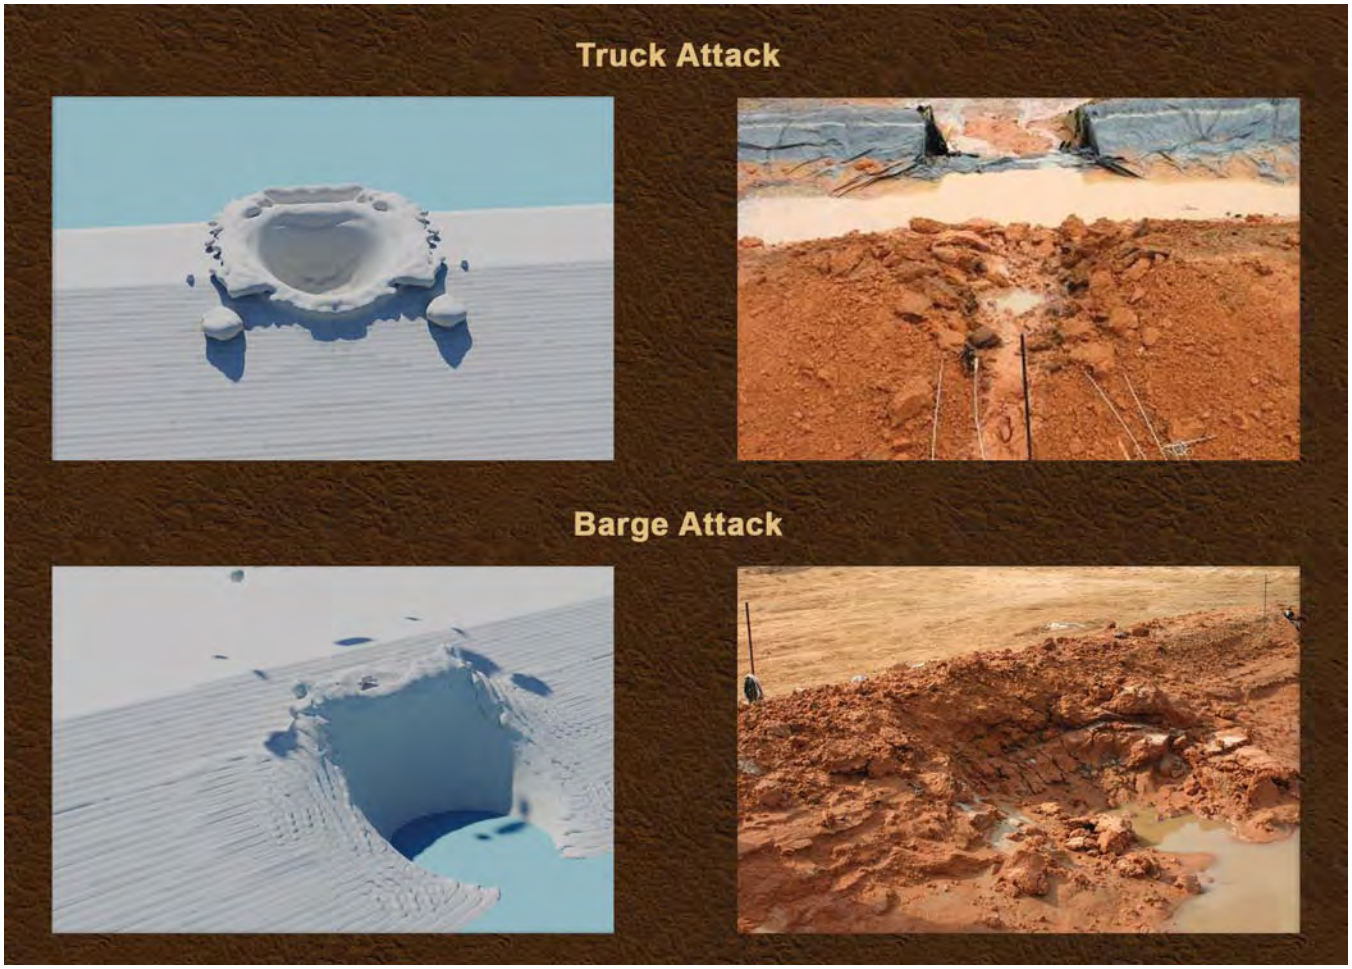

Figure 3. Comparison of calculated and experimental craters

requirements and the long run times dictate that these calculations be run on state-of-the-art massively parallel computers such as the Cray XT3 and Cray XT4 computers at the ERDC High Performance Computing (HPC) Center. Calculations of this type typically require 256 or more processors, 2 to 4 days of wall-clock time, and over 500 gigabytes of disk storage.

The numerical simulations reported here were performed using CTH. CTH is an Eulerian shock physics code that numerically solves the partial differential equations describing the conservation of mass, momentum, and energy. It does this in a structured Eulerian mesh fixed in space and uses equations of state to close the coupled system of equations. CTH was designed to treat a wide range of shock-wave propagation and material-motion phenomena in one, two, or three dimensions. Eulerian codes are able to handle severe material distortion, but accomplish this by sacrificing their ability to keep track of material interfaces.

The simulations modeled the wet soil of the embankment, the nearly saturated soil below the phreatic surface, the explosion of the bomb, air, and water in the reservoir. The material model used for the earthen materials in the calculation was based on Mie-Grüneisen Equation of State using the P-Alpha model to simulate the volumetric crush-up response of the soil and a Geo-Yield strength model to simulate the shear response of the material. The explosive was modeled using a standard JWL equation of state for C-4 with default parameters and a detonation velocity of 8.193e+5 cm/sec. The water was modeled with the default CTH Mie-Grüneisen equation of state for water, and the air was modeled using the standard Sesame tabular equation of state. Gravity was used in the simulation, and the CTH gravitational stability model was used for setting the initial pressure state of materials to pressures expected at corresponding depths.

Figure 2 shows early time CTH materials and pressures at times of 0.5, 1.0, 1.5, and 2.0 milliseconds for the VBIED attack on the crest of the embankment. The early time material plots essentially show the development of the fireball resulting from the explosion. In this figure, the explosive charge and detonation products are shown in red, the two soil types in shades of tan and brown, and the reservoir water is shown in blue. The corresponding pressure plots show the propagation of the blast wave through the air, the soil comprising the dam, and the water in the reservoir. Early crater formation can be seen by observing the interface of the soil and air/detonation products.

Figure 3 shows the resulting craters for both the numerical simulations and the experimental tests. The crater size for the numerical simulation for the crest attack with a VBIED is somewhat smaller than that measured in the experiment. The agreement of crater sizes is somewhat better for the waterborne VBIED. Work is currently progressing to refine material models in CTH as well as the utilization of more advanced codes to better predict cratering.

Dr. Michael Stephens, Kevin George, and Miguel Valenciano of the ERDC Data Analysis and Assessment Center were very helpful in visualizing the results of the simulations. The research reported herein was conducted as part of the Critical Infrastructure Security Program, U.S. Army Corps of Engineers, Headquarters, Office of Homeland Security. Permission from the Director, Geotechnical and Structures Laboratory, to publish this article is gratefully acknowledged. The research was conducted using an allocation of HPCMP Challenge processing hours at the ERDC.

Using HPC to Accelerate the Insertion of New Materials into DoD Systems

By Dr. Charles Cornwell, HPC PET Team at the Army Research Laboratory and HPTi, and Team Member of the ERDC Carbon Nanotube Technology for Military Engineering Team

In the future, significant developments in nanotechnology will occur at a scale of less than 100 nanometers where materials and devices are built from molecular components. High performance computing (HPC) will play a key role in advancing our understanding of the behavior of new materials at the nanometer scale. The U.S. Army Engineer Research and Development Center (ERDC) is using high performance computers to design materials at the nanometer scale in order to optimize their properties and accelerate the production and insertion of these materials within DoD systems. Miniaturization is now entering the nanoscale regime where we approach the size range of individual atoms and molecules. Characterizing the physical properties of nanoscale materials can be extremely difficult, and high performance computers have been used extensively in this area of scientific research. Over the last decade, computer simulations have emerged as powerful tools to complement theoretical and experimental studies of fundamental materials properties. Manufacturing materials with atomic-scale precision and control involves some formidable theoretical and computational challenges.

One of the more promising nanoscale materials is the carbon nanotube (CNT). CNTs are elongated versions of Bucky Balls that were discovered in 1985 by Dr. Robert Curl, Sir Harold Krato, and Dr. Richard Smalley, and for which they received the 1996 Nobel Prize in Chemistry. CNTs are very small in diameter, and range from about 0.4 nm or 16 billionths of an inch, to over 70 nm. To put this into perspective, a sheet of notebook paper is about 75,000 CNT diameters thick, depending on which variety of CNT you choose.

The mechanical properties of CNTs, including high strength, high stiffness, toughness, and low density, make them an ideal strengthening material in advanced composites, woven fabrics, and textiles with military and industrial applications. CNTs can now be spun

ERDC Carbon Nanotube Technology for Military Engineering Team, Molecular Dynamics Simulation Group: (Pictured left to right) Dr. Charles R. Welch, Program Director, and Richard Haskins, ERDC Information Technology Laboratory (ITL); Dr. Charles Cornwell, HPC PET Team at Army Research Laboratory and HPTi; Dustin Majure, ERDC ITL; and Dr. Robert Kirgan, ERDC Environmental Laboratory (EL). Group members not pictured are Dr. Nicholas (Jabari) Lee, ERDC ITL, and Dr. Anthony Bednar, ERDC EL

into continuous fibers that, in theory, can be produced to any length. The strength of these fibers exceeds the strength of spider silk, making them the strongest fiber known to man. Spider silk is a natural material that researchers have long tried to mimic because of its renowned toughness. The CNT fibers are reported to be 20 times as tough as steel wire, 17 times as tough as Kevlar, and 4 times as tough as spider silk (*Science News*). New composite materials will provide one to two orders of magnitude improvement in material properties such as tensile, compressive, and shear strengths, elastic strain limit, shear and Young's moduli, and strength-to-weight ratio over existing construction materials. The CNT fibers can improve the performance of materials used in applications including fiber-reinforced, high-strength concrete, structural steel, light-weight portable structures, construction materials, armor protection, and high-strength-to-weight materials for aircraft components.

The impetus for nanotechnology comes from a new generation of analytical tools such as the atomic force microscope and the scanning tunneling microscope that allows researchers to observe and manipulate individual atoms and molecules. However, the extremely small dimensions of CNT fibers impose a tremendous challenge in the experimental study of their structure and mechanical properties. The internal structure such as the diameters, length, orientations, and chiralities of the CNTs determines the properties of the fibers, but there are practical limits in the range of material parameters that can be explored experimentally and in the capability of instrumentation to provide a comprehensive picture of relevant phenomena that describe the properties of CNT fibers. This is especially true in the case of critical transient behavior.

In recent years, significant progress in modeling the performance and processing of materials has been made. These advances have primarily been in the area of application of continuum methods where empirical constitutive relations are used to describe the behavior of materials. These methods can be used to predict the performance of a material following the deformation process and the temperature history of the material. However, these models have had little impact on the design and optimization of material properties because they cannot quantitatively determine, predict, or manipulate the internal structure of the material. CNT fibers are manufactured via advanced processing techniques that greatly influence the atomic structure of the constituent CNTs and thus its material properties. Manipulating the atomic arrangement of the fibers is the primary mechanism that materials scientists and

Nobel Laureate Robert Curl (right) visits the ERDC High Performance Computing Center in November 2007. His host, Dr. Bob Welch (left), is Program Manager of the Carbon Nanotube Technology for Military Engineering Program. Professor Curl's co-discovery of Carbon 60 is considered to have paved the way for research and discovery of carbon nanotubes. Curl, along with Sir Harold Krato and Professor Richard Smalley, received the 1996 Noble Prize in Chemistry for their discovery

engineers have to optimize the properties and performance of the CNT fibers.

Realistic numerical simulations provide both quantitative and qualitative insights into the properties of materials at the atomic and molecular levels. The models and simulations are sometimes thought of as taking up a middle ground between theory and experiment, with some commonalities with both. The calculations are based on our theoretical understanding of the interactions between atoms and molecules. However, the results are often analyzed in much the same way as experiments. Theory relies on analytical and numerical approximations that reduce the answer to complex problems to a model tractable in a closed-form solution. This process usually restricts analytical solutions to highly idealized problems that are difficult to interpret in terms of experimental results.

Numerical methods and computers solve problems of a more complex nature by replacing the need for approximation with elaborate calculation schemes. It was not until the introduction of powerful computers that modeling and simulation found widespread application. Theoreticians are no longer restricted to working with idealized models that are amenable to closed-form solutions, but can deal directly with detailed and highly realistic approaches to the original systems. The level of complexity of problems that can be solved using HPC has extended the range of problems that can be

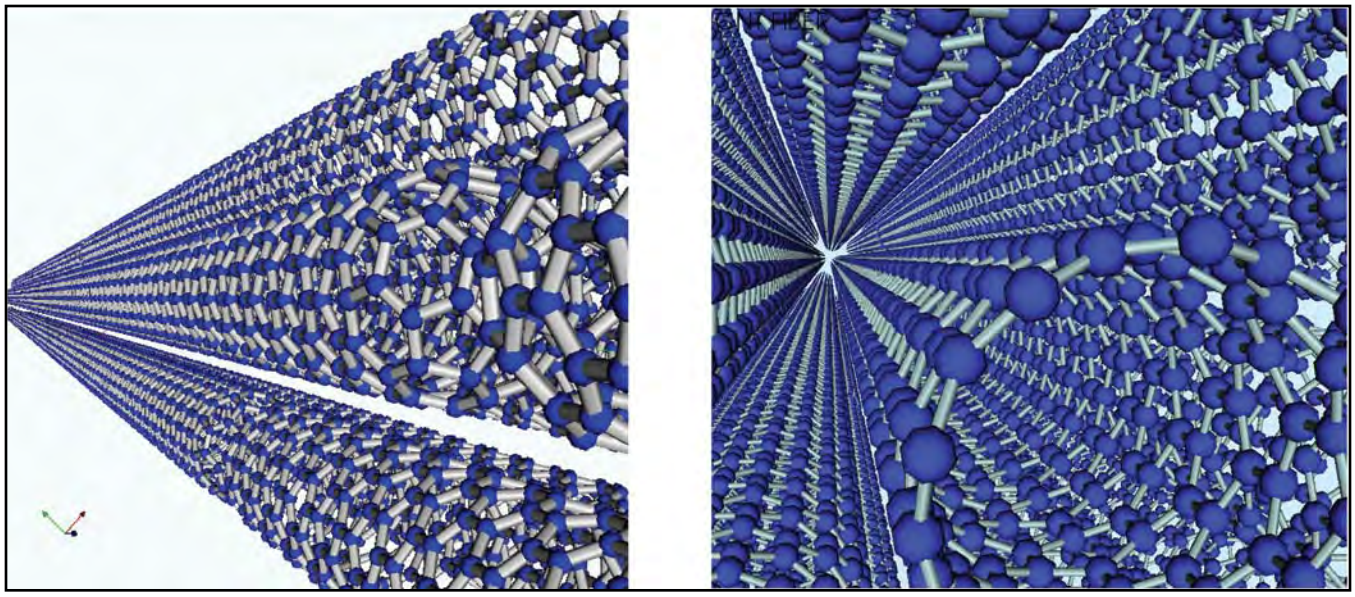

Figure 1. Largest simulation performed on CNTs fibers 8000Å long consisting of CNTs with average length of 2000Å. Fiber required millions of atoms to construct

simulated to encompass real systems. High performance computers allow researchers to perform calculations on digital models that are faithful representations of the actual material studied in the laboratory.

Researchers at ERDC are developing models and techniques to study and improve the properties of CNT-based construction materials. The initial goal of the project is to use modeling and simulation to produce a design for a one-million-pounds-per-square-inch CNT fiber. We cannot expect to model a macroscopic system with on the order of 1023 atoms using atomistic simulations. Fortunately, this is not required. Recent advances in experimental techniques provide methods for the direct measurement of the mechanical properties of individual CNT fibers. With the advent of nanotechnology and the HPC resources at ERDC, researchers can now model a representative ensemble of CNTs consisting of millions of atoms (Figure 1) and perform calculations on fibers at the nanometer scale using atomistic calculations that are representative of the fibers under study in the laboratory. The simulations are not meant to replace experimentation, but can assist in the interpretation of experimental results and guide one in the direction of the most promising avenues of exploration.

This increase in size and complexity of these simulations requires algorithms that run efficiently on high performance parallel computer platforms. The calculations presented here were performed using the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS). LAMMPS is a classical molecular dynamics code designed to run efficiently in an HPC environment. It was originally developed in the 1990s for a cooperative research and development agreement (CRADA) led by Steve Plimpton at Sandia National Labs. The calculations were performed on the Cray XT3 (Sapphire). The largest fiber simulations consisted of approximately one million atoms and used up to 2000 processors. The calculations required to produce the results in this article used an estimated half-million CPU hours and were completed over the course of a couple of months. The same calculations performed on a single-processor machine of comparable speed would take tens of years to complete.

The program began with the study of individual CNTs to determine their mechanical properties when placed under strain. The calculations show that the tensile strength of the CNTs is considerably larger than the tensile strength of CNT fibers measured experimentally. Experimental observations indicate that the strands of a CNT fiber occur in hexagonal closest packed (HCP) bundles of parallel CNTs of similar radius. It is believed that fibers composed of parallel aligned CNTs fail, which is due to the CNTs sliding past each other because of the weak van der Waals force that binds the CNTs together. It has been suggested that the strength of the fibers can be increased by increasing the length of the constituent CNTs and that the maximum theoretical strength of CNT fibers would be obtained when the shear force is equal to the intrinsic breaking strength of the constituent CNTs. Therefore, we began our study of the failure mechanics of CNT fibers with a bundle of CNTs to study the behavior of a single CNT as it is extracted from a bundle. We construct a bundle

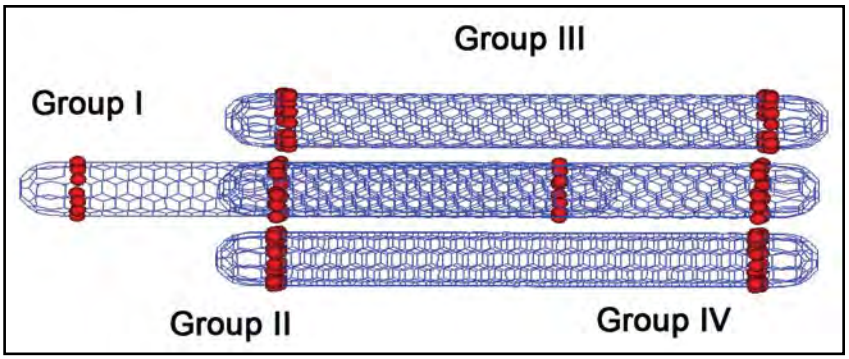

Figure 2. Experimental observations indicate strands of CNT fiber occur in hexagonal closest packed bundles of parallel tubes of similar radius. Longitudinal axes of CNTs lie parallel to the z-axis. Four groups of atoms used to control and monitor response of system during the extraction process

of parallel CNTs, arranged in an HCP configuration with a one-core CNT surrounded by six on its perimeter (Figure 2). This geometry does not represent the local environment of CNTs in a fiber, but the extraction of a single CNT from a bundle should exhibit some of the basic characteristic behaviors of the CNTs in a fiber when they slip past each other.

The simulations explore the effect of increasing CNT length on the evolution of yield stress in an attempt to determine the critical contact length needed to achieve maximum fiber strength. Figure 3 provides typical results for the cases studied. In Figure 3, the total tensile force across the complete cross section of the central carbon nanotube at the Group I atoms is shown as a function of contact length. The tensile force is generated in the central CNT by the extraction process as a result of the shear forces between the central CNT and the six perimeter CNTs in the bundle. The tensile force is plotted as a function of the contact length of the CNTs. The first part of the force-contact length curve is linear and is as we had expected. However, as the contact length increases, the tensile force becomes asymptotic to some value that is a function of the chiralities, orientations, and diameters of the CNTs involved, and does not continue to increase with length. Thus, the criteria for critical contact length, that is, the length by which the shear forces between the carbon nanotube bundle and the central nanotube exceed the tensile strength of the central CNT, are never met for CNT bundles. This result is counter-intuitive and is the consequence of non-uniform straining of the CNTs along their length, allowing the central tube to be extracted via a series of strain-slip events. These results apply to slow or nearly static extraction of the central tube and so correspond to creep in a macro material.

The largest value for the shear force and local strain was observed in the bundle comprised of 1000Å long CNTs. The contact length for the system was 973.61Å, and the shear force was 4.722 (eV/Å). The maximum local strain at the group I ring was 0.88 percent. Previous simulations on individual CNTs indicate that the CNTs break at approximately 17 percent strain. The strains produced during the extraction process are well below the strain required to break the CNTs. In all of the simulations, the core CNT slipped before the shear force reached the intrinsic breaking strength of the core CNT. These simulations indicate that the load transfer between parallel aligned CNTs is not sufficient for the core CNT to reach the critical length before slipping occurs. However, there is a considerable difference in the local environment and boundary conditions between the CNT in the above simulations and that of CNT fibers.

To better understand the failure mechanics of fibers, we construct a fiber from multiple strands comprised of an ensemble of parallel CNTs (Figure 4). The strands of the fiber consist of multiple CNTs that are arranged in an HCP configuration with the longitudinal axis of the CNTs running parallel to the z-axis. The data obtained from the simulations are used to calculate the stress-strain relationship for a series of fiber configurations.

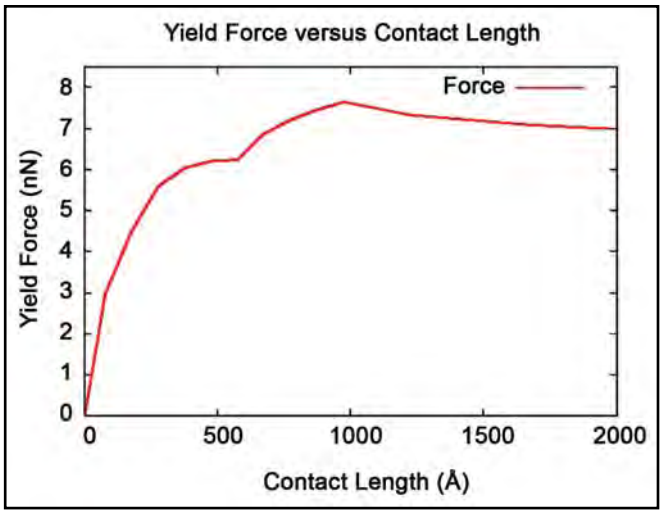

Figure 3. Yield force in central CNT (Group I atoms) versus contact length for 200-2000Å CNT bundles. The cusps in the graph at 600Å and 1000Å indicate regions of rapid change in response to slip events

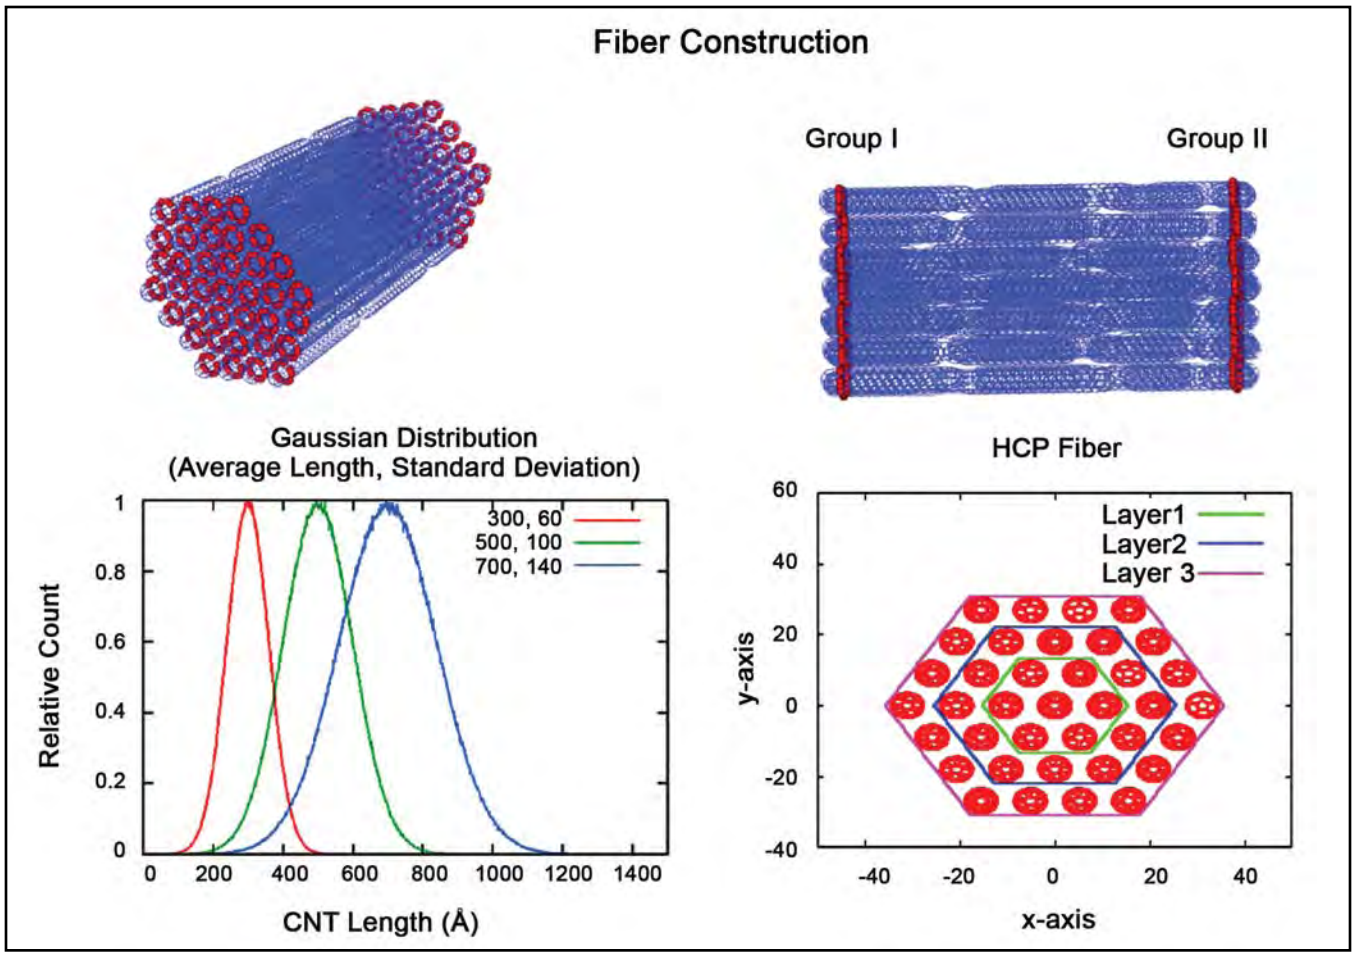

Figure 4. CNT bundle: The two groups of atoms (UL,UR) are used to constrain the fiber during the quasi-static stretching. Example output for the distribution function (LL) used to generate the CNTs gives an indication of the range of acceptable CNT lengths for the fibers. The fiber cross section (LR) shows the hexagonal arrangement of the CNTs, and the cross-sectional area used in the calculation of the stress

The stress is given by $\sigma = F/A$, where A is the cross section of the fiber and F is the force required to maintain the constraint on the system. The area of a regular hexagon for the number of layers in the fiber (Figure 4) is used to calculate the cross section of the fiber. The center of mass distance between the atoms in groups I and II is used to calculate the length of the fiber. We ran simulations on a series of CNT fibers to investigate the effect of CNT length distribution on the stress-strain relationship in parallel-aligned CNT fibers under strain. The data obtained in these simulations are used to establish a relationship between the average CNT length and the elastic modulus and tensile strength of the fibers.

Figure 5 is a plot of the stress versus strain for fibers 1600, 4000, and 6400Å long, comprised of CNTs with average lengths 400, 1000, and 1600Å, respectively. The yield strength of the fibers represents an upper limit to the load that can be applied and is defined as

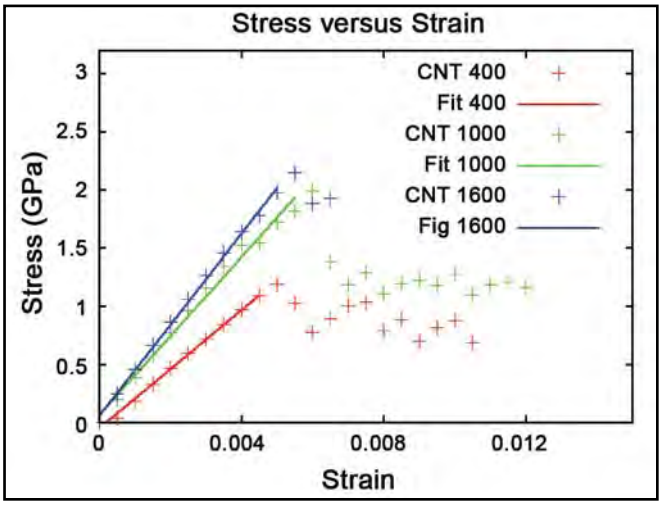

Figure 5. Stress-strain for several fiber configurations. The fibers are constructed using CNTs with an average length of 400, 1000, and 1600Å. The fiber length is four times the average CNT length, and the standard deviation in length is 20% of the average CNT length

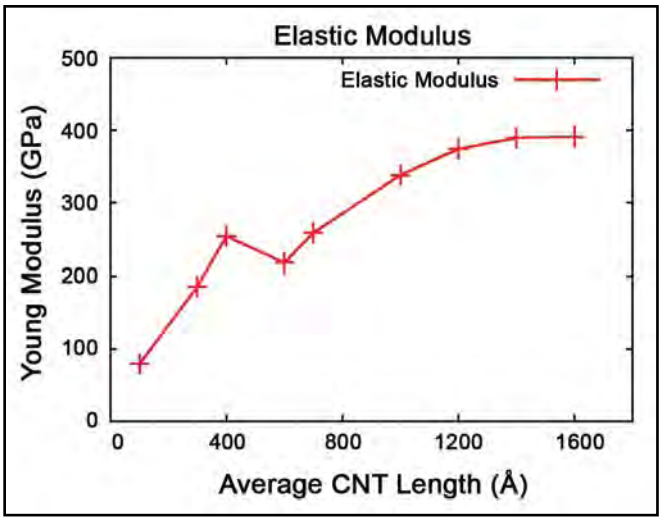

Figure 6. Young's modulus for all fibers. Fiber lengths ranged from 400 to 6400Å. The average CNT lengths ranged from 100 to 1600Å

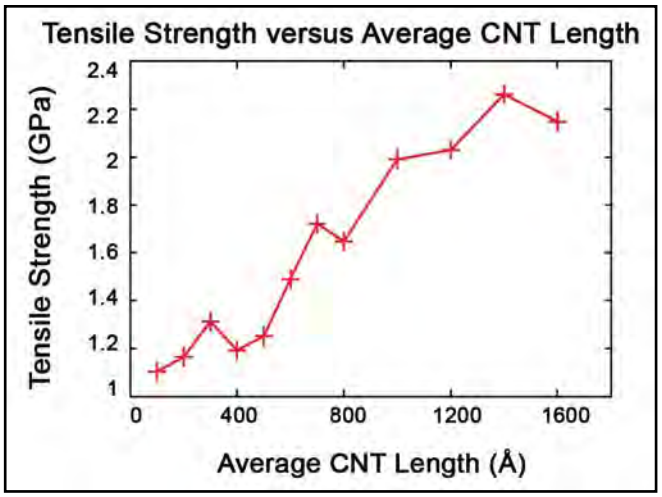

Figure 7. Tensile strength for all fibers. Fiber lengths ranged from 400 to 6400Å. The average CNT lengths ranged from 100 to 1600Å

the tensile strength of the fiber. A least squares fit to the data points in the linear region of the plots is used to calculate the elastic modulus of the fibers. Data from the individual plots provide the information necessary to determine the elastic modulus and tensile strength as a function of the average CNT length (Figures 6, 7).

The maximum elastic modulus obtained from the simulations is 430 GPa for fibers with an average CNT length of 1600Å. Figure 6 indicates that the value of the elastic modulus does not continue to increase with an increase in the average length of the constituent CNTs. Therefore, 430 GPa is close to the maximum obtainable elastic modulus for this fiber configuration. Maximum yield strength of 2.3 GPa is obtained for fibers with an average CNT length of 1400Å (Figure 7). The yield strength in Figure 7 continues to increase with an increase in the average CNT length, but the rate of increase begins to taper off. Therefore, the tensile strength will not continue to increase with an increase in the average length of the constituent CNTs. However, it is difficult to estimate the maximum value for the tensile strength from the data in Figure 7.

These analyses show that for the slow straining of the CNT-CNT structures, the interaction forces between CNTs become asymptotic with respect to contact length for perfectly aligned CNTs subject to van der Waals forces only, and the tensile strength of the CNTs is never reached. This asymptotic value is dependent on the orientation, chirality, and diameter of the CNTs involved. The asymptotic behavior is the result of the non-uniform axial straining of the CNTs by the extraction process, allowing discrete strain-slip events to occur. For short contact lengths, the interaction force is linear, and for medium contact lengths, the interaction force is nonlinear and increasing with contact length. The simulations indicate that the elastic properties of the CNTs limit the shear force required to initiate sliding between CNTs and provide a mechanism for releasing strain energy. The results suggest that initially there is an advantage to increasing the length of the constituent CNTs in a fiber, but ultimately the load transfer mechanism is the factor that will determine the maximum strength of parallel-aligned CNT fibers. Consequently, achieving the maximum strength in a parallel-aligned CNT filament of arbitrary length will require enhanced load transfer between the CNTs. Such enhancements might include fiber twisting, CNT entanglement, or cross-bonding between CNTs. Current research is focused on fiber designs that facilitate load transfer between the CNTs in the fibers by increasing the CNT-CNT interactions (Figure 8).

Numerical simulations are becoming more realistic and providing both quantitative and qualitative insights into the properties of materials at the atomic level. The use of HPC is becoming an integral part of materials science research and is essential for developing realistic material models that have the predictive capability necessary for the design and production of novel materials with atomic-scale precision and control. The ERDC research program spearheading this quest for CNT-based super material is known as the Carbon Nanotube Technology for Military Engineering Program. The research is strongly supported by the HPC PET and is a collaboration among the Information Technology Laboratory, the Construction Engineering Research Laboratory (CERL), and the Environmental

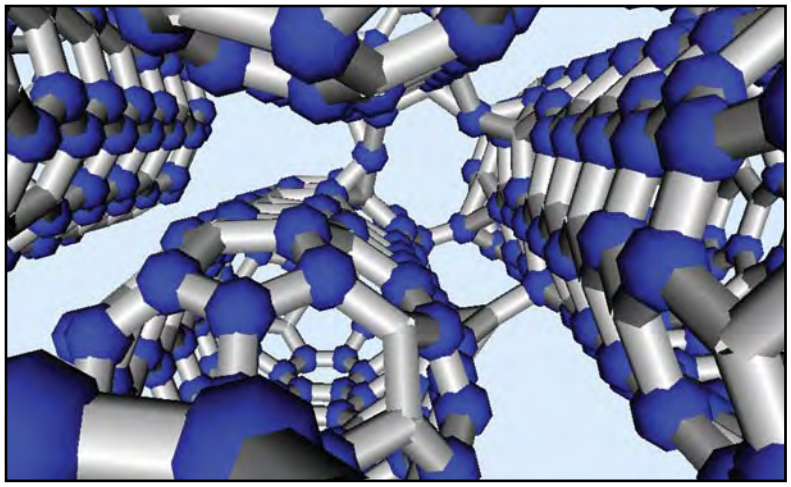

Figure 8. CNT fiber with cross-linking between the CNTs. The interstitial carbon atoms form chemical bonds with atoms in adjacent CNTs in the fiber. Experimentally, this can be accomplished by irradiating the fibers with high energy particles to knock carbon atoms out of the CNT lattice. These carbon atoms then form bonds with atoms in adjacent CNTs

Laboratory. It is funded under the Environmental Quality and Installations Program and was recently cited by the ERDC Director as being the "Major Fundamental Research Accomplishment" for ERDC. The Program Director is Dr. Charles R. Welch, the Technical Director is Marty Savoie, and the Senior Manager is the CERL Director, Dr. Ilker Adiguzel. The initiative is trying to link theories for the properties and performance of CNT materials, at the atomic and molecular level, with the processing steps required to produce high performance CNT-based materials. The use of HPC improves the quality of research at the DoD research centers and is fundamentally changing the way new materials are developed.

The potential benefit of producing molecular-designed materials is unprecedented. We are only limited by our imagination, ingenuity, and the ability to compute the enormous volume of calculations required to solve these problems.

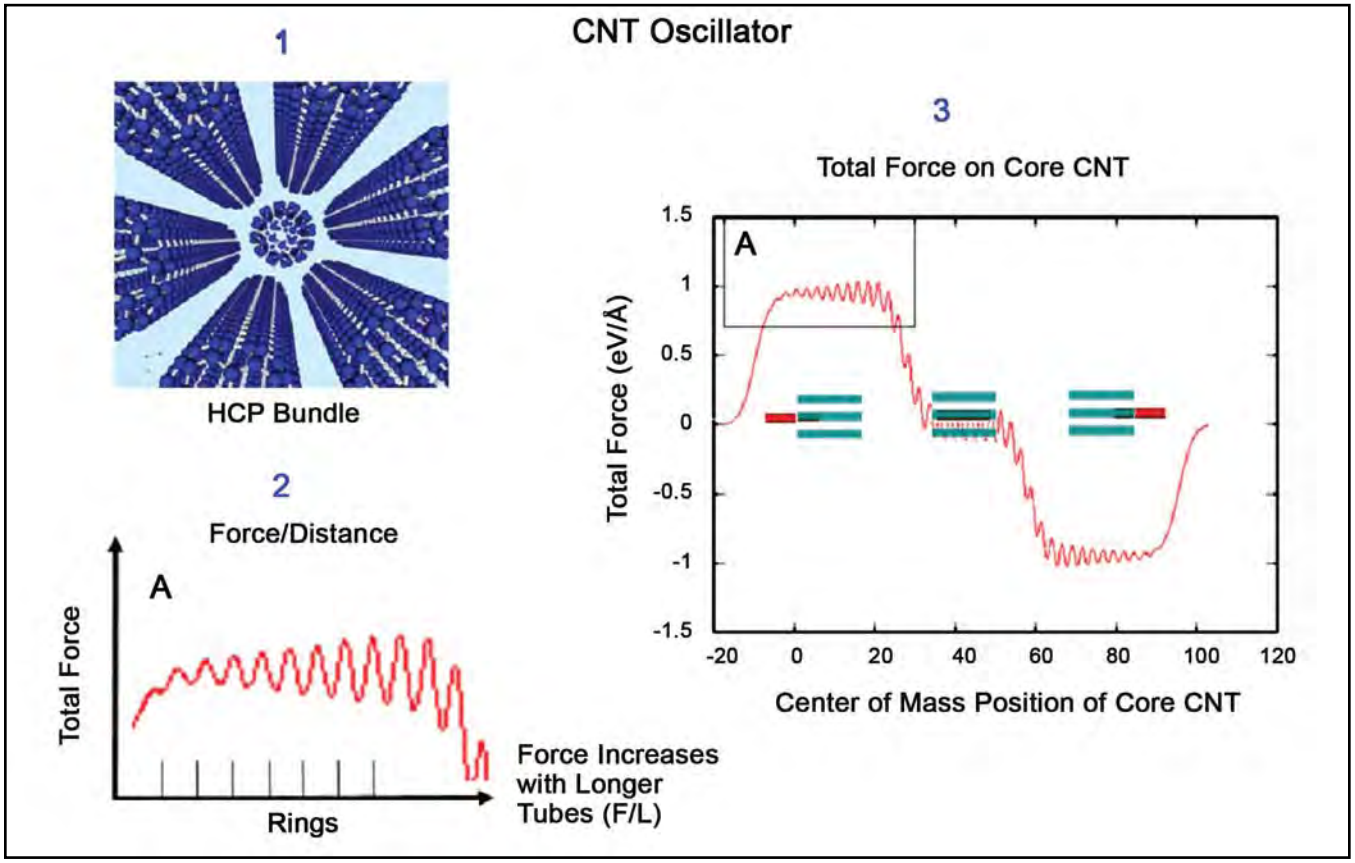

Figure 9. (1) HCP bundle of (5,5) CNTs, (2) insert showing the increase in shear force with increased contact length between core and perimeter CNTs, (3) shear force on core CNT as it is translated through center of bundle

Wave Information Studies (WIS) Pacific Region Hindcast

By Barbara Tracy, ERDC Coastal and Hydraulics Laboratory, and Deanna Spindler, Science Applications International Corporation (SAIC) at the National Oceanic and Atmospheric Administration (NOAA) National Centers for Environmental Prediction (NCEP)

Pacific coastal wave information is important for defense and civil coastal engineering and wave climate applications. The Wave Information Studies (WIS) uses numerical wave hindcasts to produce wave information for a dense network of points near the coast in addition to deepwater locations.

The WIS, Coastal Field Data Collection Program, Coastal and Hydraulics Laboratory at ERDC has recently completed a 1981-2004 Pacific basin hindcast at a 0.5-degree grid resolution. This wave information is applicable for studies in the Hawaiian Islands and U.S. territories, but finer resolution is needed for the U.S. mainland Pacific coastline. WIS is currently producing a 2000-2005 Pacific regional hindcast consisting of a three-grid system to define the Pacific coast wave climate. This application uses the 0.5-degree Pacific basin wind fields and a new set of 0.25-degree wind fields especially developed for the U.S. Pacific mainland coastal area by the meteorologists at Oceanweather, Inc. (OWI). This study is utilizing new wave modeling technology in the form of a nested-grid, full-spectral third-generation wind-wave model WAVEWATCH III (version 3.12) developed at NOAA/NCEP (Tolman 2007). This WAVEWATCH III technology

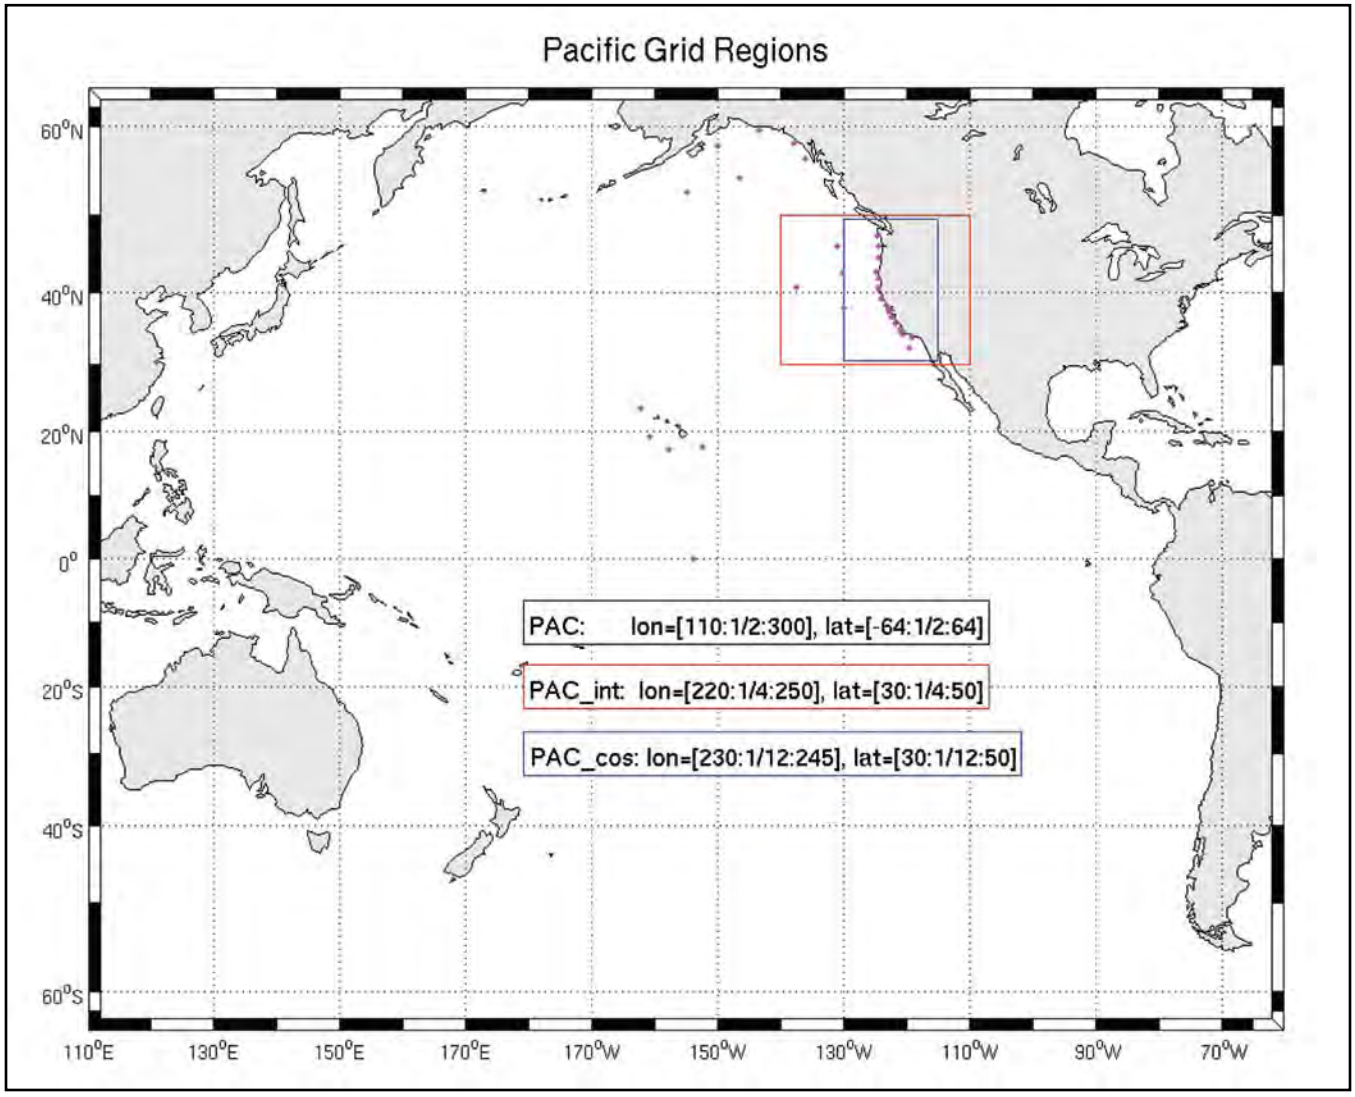

Figure 1. Set of three nested grids for Pacific regional hindcast (Spindler and Tolman 2007)

allows energy to go back and forth through the boundaries between the nested grids. One computer run will include three grids to produce wave information for a 1-year Pacific Regional wave hindcast using the MPI parallel computing facilities at the ERDC MSRC. This Pacific wave information will be invaluable for U.S. west coast engineering and wave climate applications.

Figure 1 shows the set of three nested grids for the Pacific regional application. The full Pacific basin grid with 0.5-degree resolution is labeled "PAC" in Figure 1. Bathymetry is required for each of the grid points in all these grids. Bathymetry for the PAC grid was derived from the operational NOAA bathymetry grid based on gridded relief data from ETOPO2 (Tolman 2003). An island obstruction grid derived from the operational NOAA information was also used. This obstruction grid simulates the presence of an island at a grid point by blocking some of the energy in the proper directions and is an important component of a full Pacific basin simulation. The `PAC_int` and `PAC_cos` grids also use bathymetry information from ETOPO2. The `PAC_cos` grid was chosen to cover the shelf break. Land-sea masks and obstruction grids to simulate islands were also prepared for the two finer grids. Figure 2 shows the `PAC_cos` grid bathymetry and available measurement stations near the Pacific coast.

WAVEWATCH III was run on the CRAY XT3 (Sapphire) at the ERDC MSRC in Vicksburg, Mississippi. Original hindcast tests were done with monthly input wind files, but the usual WIS hindcast procedure works with 1 year of input and output information per run. This resulted in much consultation with the high performance computing (HPC) experts at the ERDC MSRC. Compilation consisted of a standard compile for MPI parallel processing, and the hindcast was run in a batch environment. Several environment variables were set on the CRAY XT3 in order for the runs to work.

These included the following:

`MPICH_PTL_OTHER_EVENTS=16384`
and
`MPICH_PTL_MATCH_OFF=1`

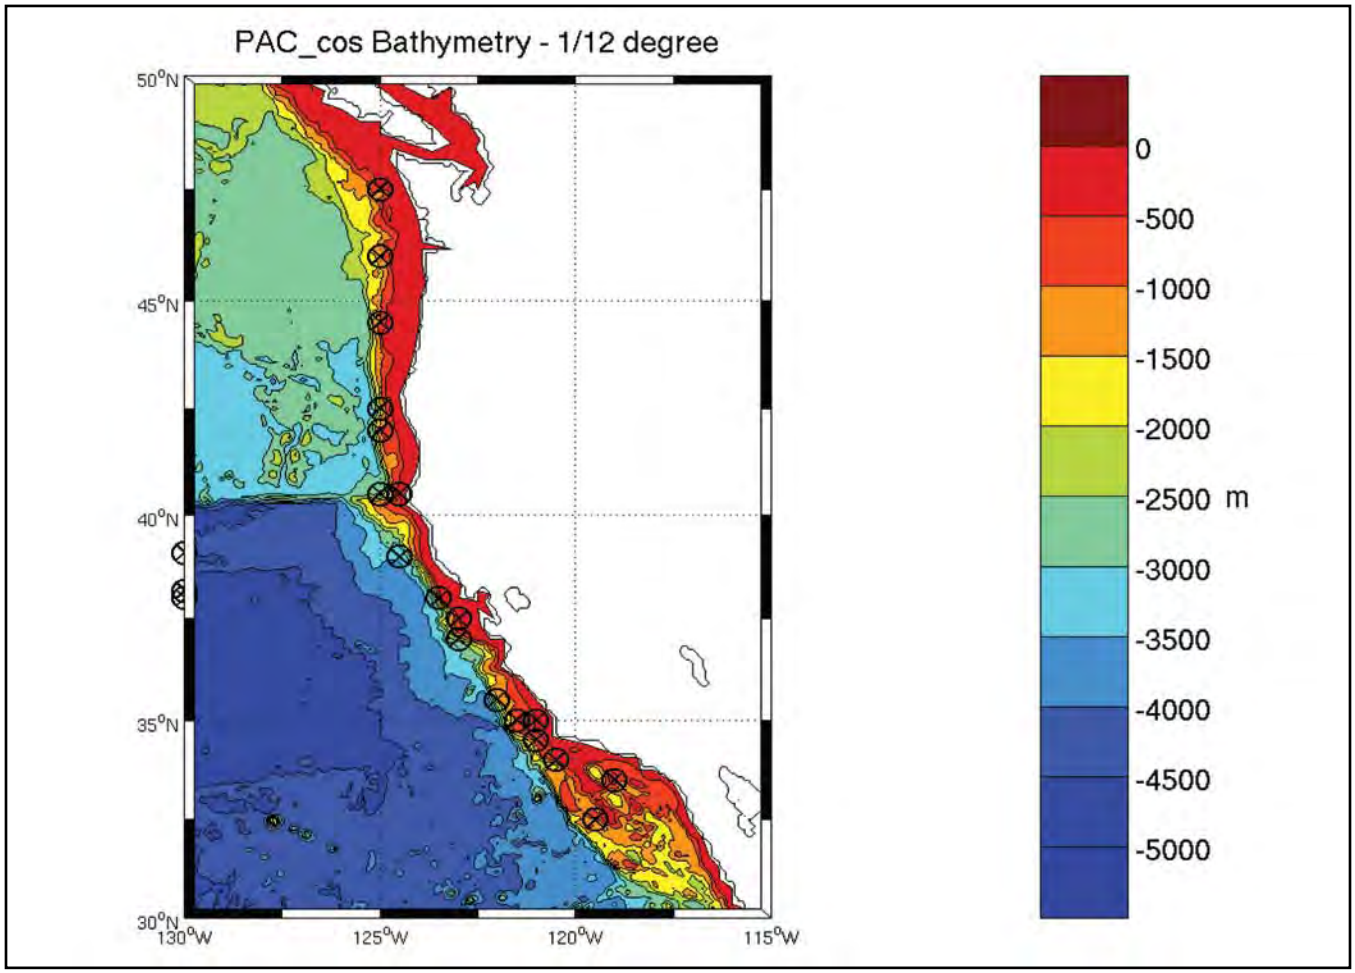

Figure 2. `PAC_cos` bathymetry contour plot. Depth values in meters. Circled locations indicate measurement sites (Spindler and Tolman 2007)

A 1-year hindcast run for the year 2000 with 64 processors took 43 hours. Profiling tools currently indicate no memory leaks, but tests will continue to speed up the total run time.

WIS hindcast information is compared with all available measurement devices to validate numerical results. Bulk wave parameters such as significant wave height, peak period, mean period, and wave direction plus wind speed and wind direction for coincident hindcast and measurement locations are compared in monthly line and scatter plots. Statistical measures such as bias, root mean square (RMS) error, scatter index (SI), skill score (SS), and correlation (Corr.) are computed for all these quantities. WIS also uses spectral validation tools in the form of the Wave Model Evaluation Diagnostics System (WaveMEDS) to validate the numerical spectral results with the measured spectral results. WIS used the WaveMEDS (Hanson et al. 2008) system to evaluate Pacific basin performance of several numerical wave hindcast models resulting in the choice of WAVEWATCH III for the Pacific basin hindcast.

Figure 3 shows Pacific regional line plot comparison results for the National Data Buoy Center (NDBC) location 46042 (36.75N, 122.42W near Monterey, California) for December 2000. Figure 3 shows excellent agreement between the hindcast and measured wave and wind information.

Table 1 provides December 2000 comparison statistics from the previous Pacific basin hindcast at NDBC 46042, and Table 2 includes Pacific regional hindcast December 2000 comparison statistics for the Pacific regional hindcast 46042 plot in Figure 3. Variables included in the tables are WS (wind speed) in m/sec, HS (energy-based significant wave height) in m, TP (peak wave period) in sec, and TM (mean wave period) in sec. Bias is computed so a positive number means that the hindcast is higher than the measurements. SI is a measure of the data's scatter, and a smaller number indicates better agreement. A perfect SS is 1.00, so numbers close to 1.00 indicate good agreement. Wind speed shows definite improvement in the regional hindcast since the regional windfields include blended

Table 1. NDBC 46042 December 2000 PAC Basin Hindcast

	Bias	RMS	SI	SS	Corr.
WS	0.37	2.39	46	0.90	0.54
HS	0.61	0.52	21	0.93	0.91
TP	0.22	2.05	15	0.99	0.54
TM	-0.02	1.00	8	1.00	0.80

Table 2. NDBC 46042 December 2000 PAC Regional Hindcast

	Bias	RMS	SI	SS	Corr.
WS	-0.25	0.85	16	0.99	0.94
HS	0.41	0.44	18	0.97	0.93
TP	0.23	2.10	15	0.99	0.53
TM	0.23	0.93	8	1.00	0.82

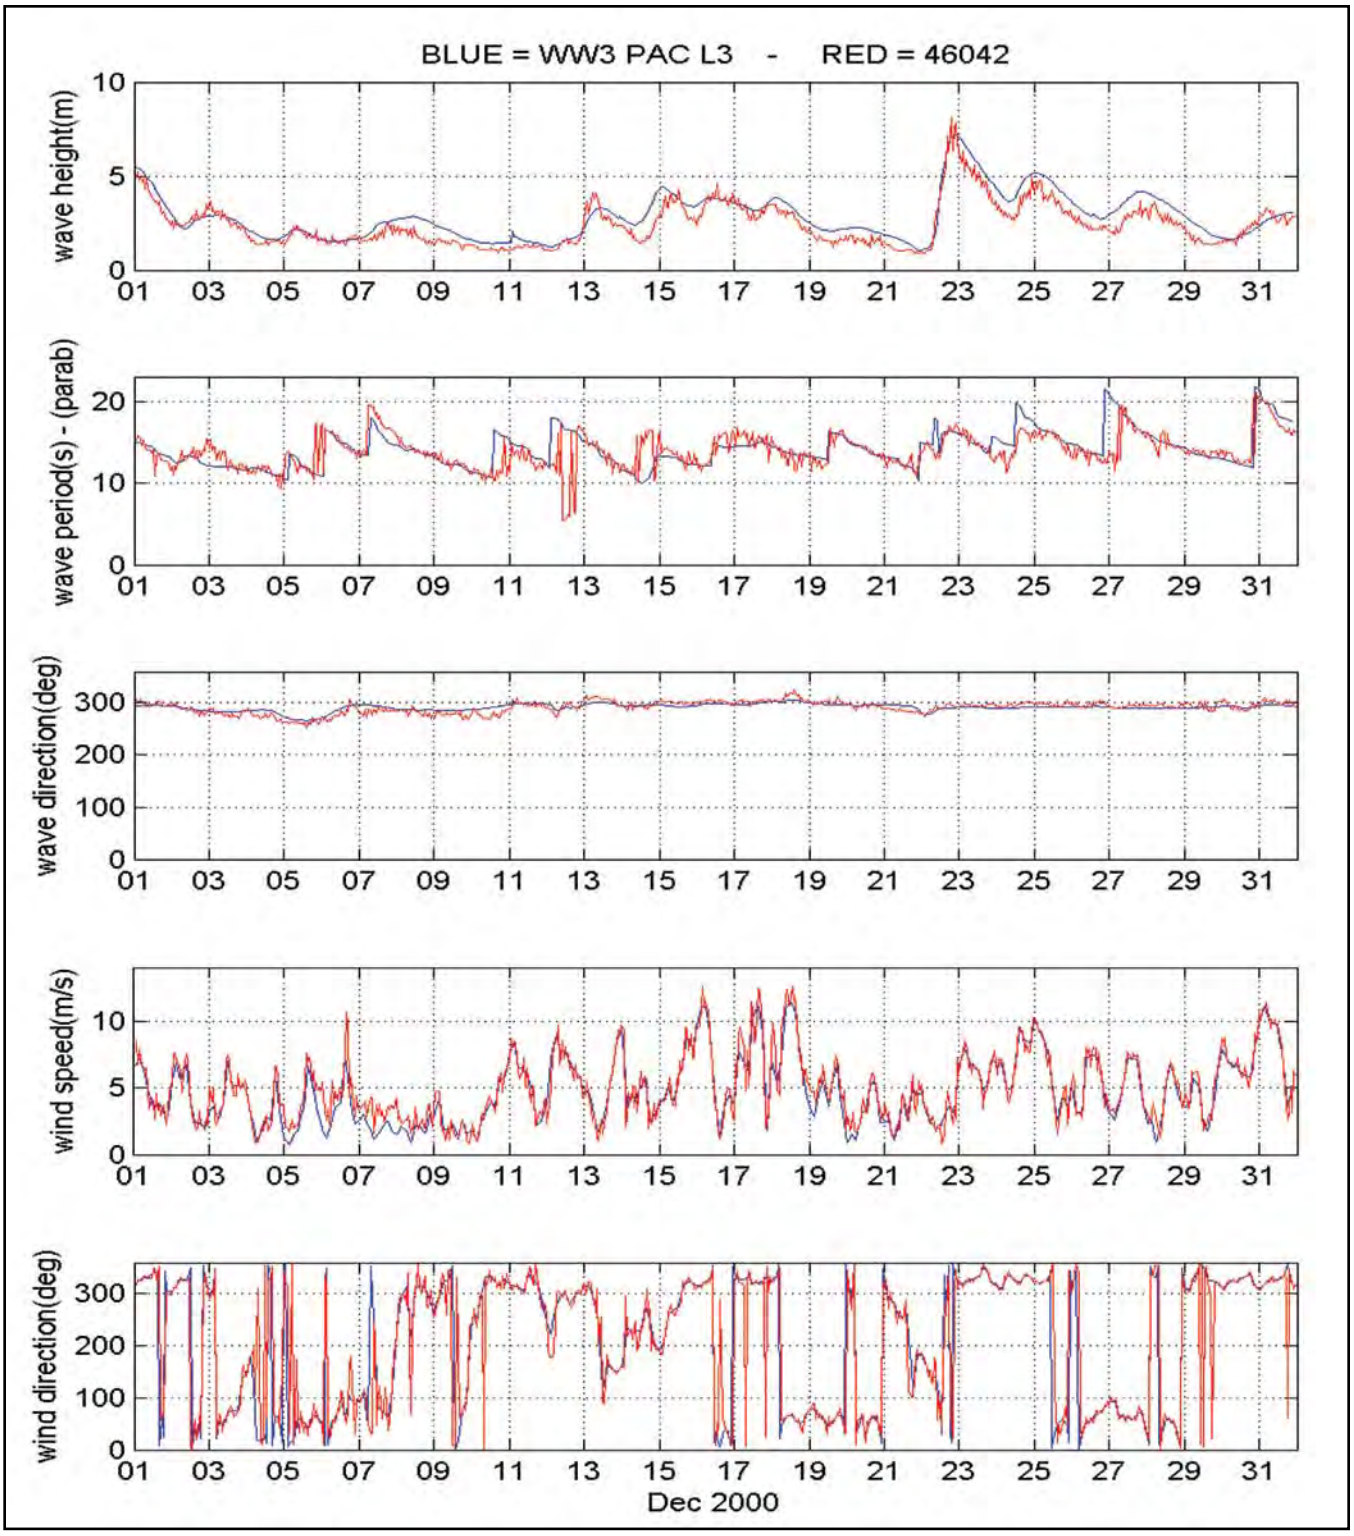

Figure 3. Comparison plot of hindcast versus measurement results for NDBC 46042 (36.75N, 122.42W near Monterey, CA). Plots show significant wave height, peak wave period, overall mean wave direction, wind speed, and wind direction for December 2000

information close to the coast. The regional hindcast significant wave height shows improvement in comparison with the basin hindcast. The wave period statistics show similar results for both hindcasts. This is logical because the spectral peak and mean wave period are driven by the energy that arrives in the regional area from the basin hindcast, and this is the same for both simulations.

A Pacific storm event for October 2000 was chosen as an example to show how the Pacific regional hindcast handles storm waves. Figure 4 from the *Mariners Weather Log* (Bancroft 2001) shows the storm on October 27, 2000, as seen by GEOS10 infrared satellite close to the time of maximum intensity. The *Mariners Weather Log* noted a ship report of seas of 10.1m and an east wind of 55knots at 54N, 136W on 1800 October 27. The storm then drifted east to the Oregon coast on October 29. Figure 5 shows Pacific regional hindcast results compared with measured results at Cape Elizabeth (NDBC 46041 at 47.34N, 124.75W) on the Washington coast. Figure 5 shows a maximum wave height of near 9m on October 29 and good agreement between the hindcast and measurements. Figure 6 shows the contoured wave height results for October 29 for the 0.25-deg grid (GRID 2). Note the area of 12m waves offshore. Figure 7 shows the contoured wave height results for the 1/12-deg grid (GRID 3). Gray's Harbor and Cape Elizabeth (46041) on the Washington coast are shown in Figure 7. Figure 8 shows a zoomed GRID 3 wave height contour plot to show more detail at 46041 and Gray's Harbor. Measurements at Gray's Harbor are not available for this storm, so hindcast results at this location will be especially helpful.

Figure 9 shows a plot of the partitioned vector wave components for the October 27th storm at Pacific regional output station 11 (46.83N, 124.25W) located close to Gray's Harbor on the Washington coast. A significant wave height plot is shown below the vector plot for reference. Maximum significant wave height from this storm was almost 8m. The significant wave height can be broken down into various components of energy at different frequencies coming from different directions. The vectors in Figure 9 show these energy components by color-coding and length. Direction of the vector indicates where the wave is coming from.

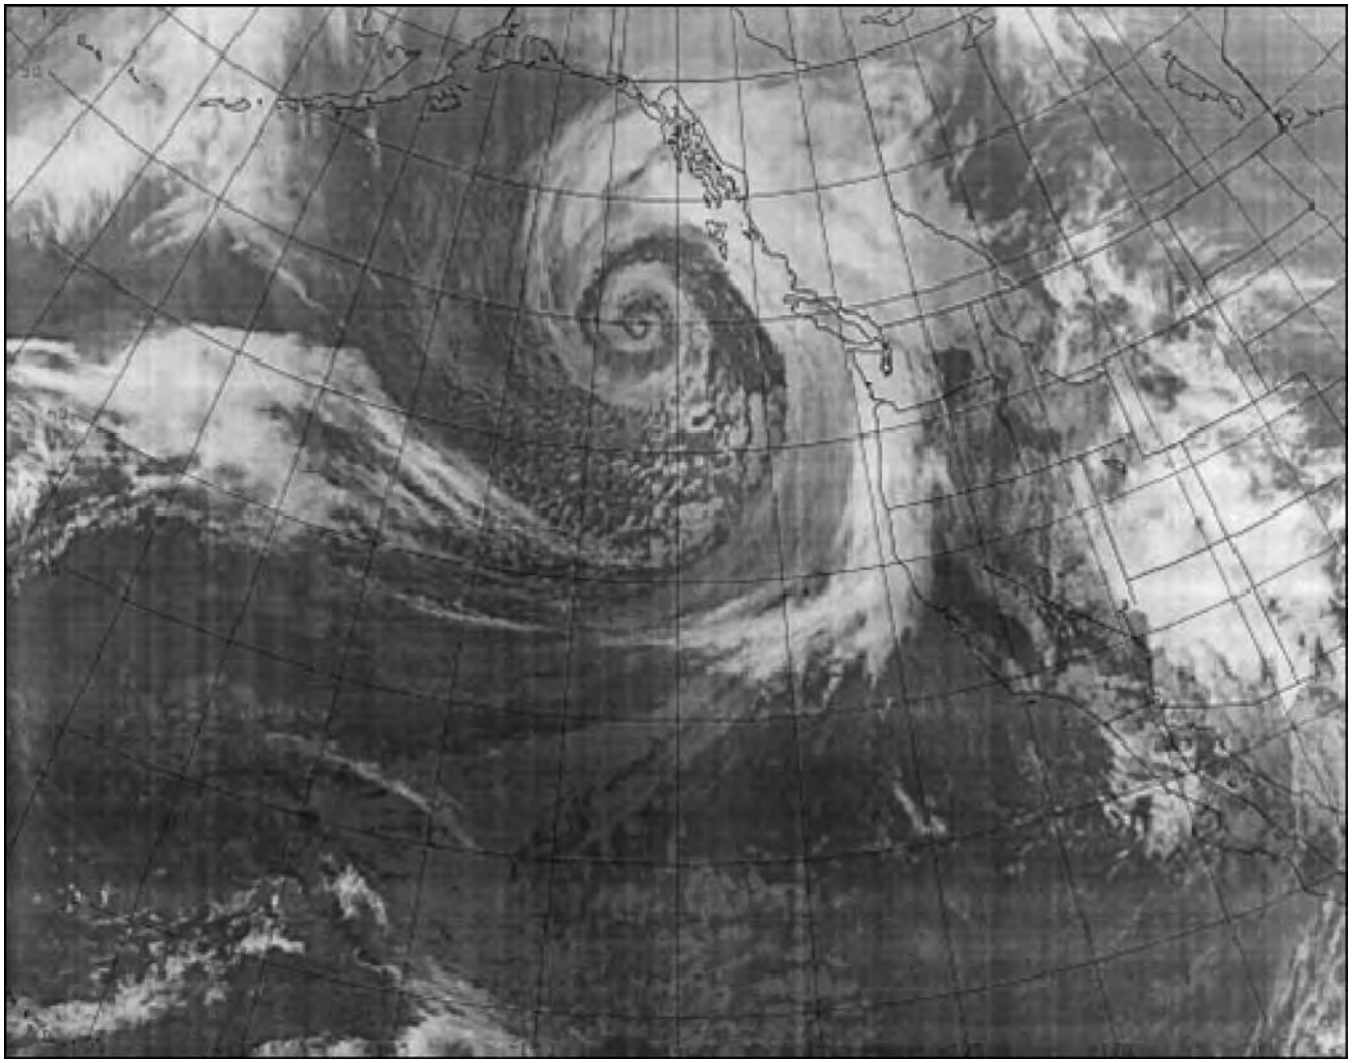

Figure 4. October 27, 2000, GEOS10 infrared satellite, 2300 UTC (Bancroft 2001) from Mariners Weather Log

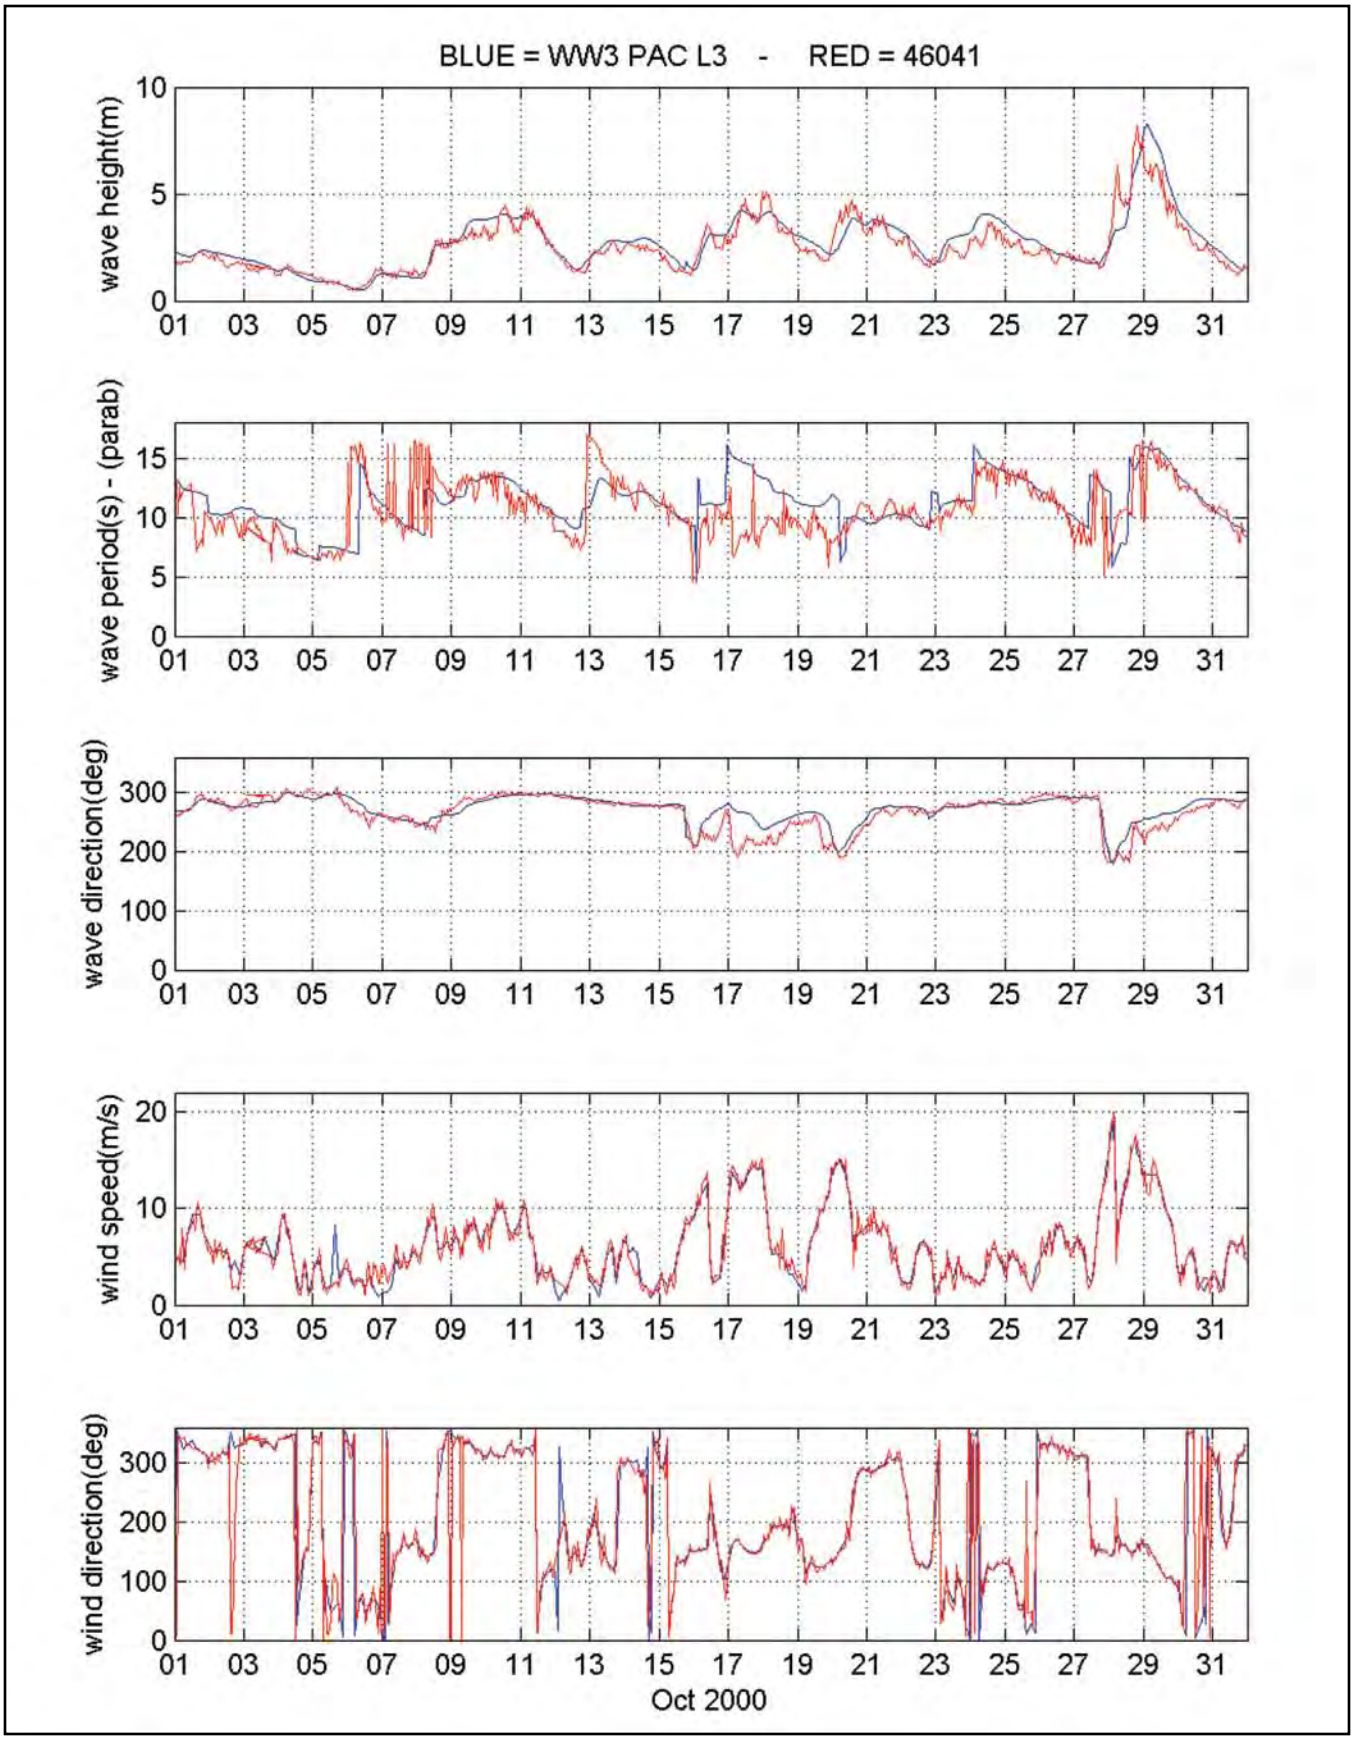

Figure 5. Comparison plot of hindcast versus measurement for NDBC 46042 (36.75N, 122.42W, Cape Elizabeth, WA). Plots show significant wave height, peak wave period, overall mean wave direction, wind speed, and wind direction for October 2000

The x-axis scale is days of October 2000, and the y-axis scale is the wave period of the spectral component vector. The vector wave component plot uses spectral partitioning output from the WAVEWATCH III version 3.12 model (Tracy et al. 2007). Figure 10 shows an example three-dimensional, frequency-direction wave spectrum plot that has been broken up into wave partitions. Evaluating the wave height, period, and direction of each of these partitions creates wave component information to plot vectors for 1 hour in a plot like Figure 9. Figure 9 clearly shows the maximum 16-sec wave energy at station 11 on October 29.

The authors would like to give special thanks to the rest of the WIS staff, Alan Cialone and Jeffrey Hanson.

References:

Bancroft, G. P. (2001): Marine weather review, North Pacific area September through December 2000. *Mariners Weather Log*, Vol. 45, #1, April 2001, p. 32-27.

Hanson, J. L., B. A. Tracy, H. L. Tolman, and R. D. Scott (2008, draft): Pacific hindcast performance of three numerical wave models. Journal paper submission.

Spindler, D. M. and H. Tolman (2007): Example of WAVEWATCH III for the NE Pacific. Technical Note 259. U.S. Department of Commerce, NOAA, NWS, NCEP, Washington, D.C.

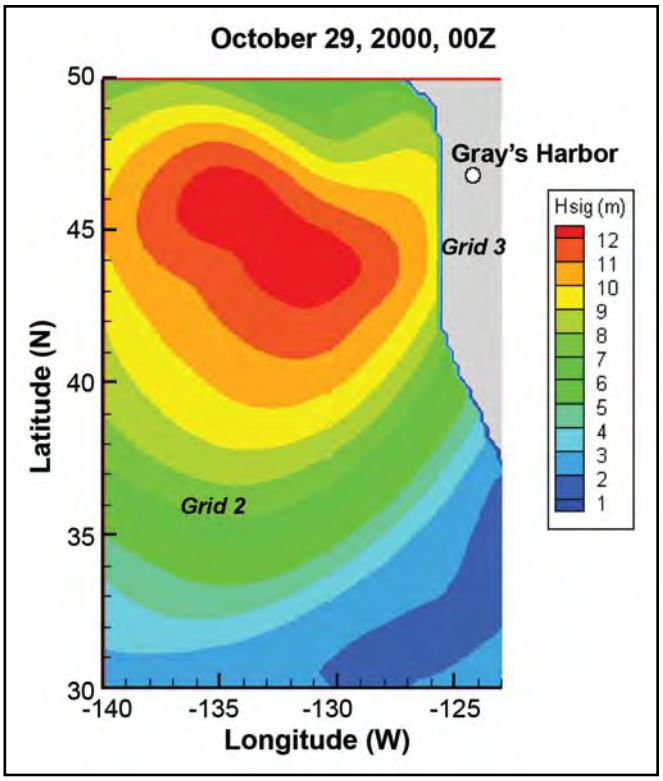

Figure 6. Contoured significant wave height results for the 0.25-deg resolution GRID 2

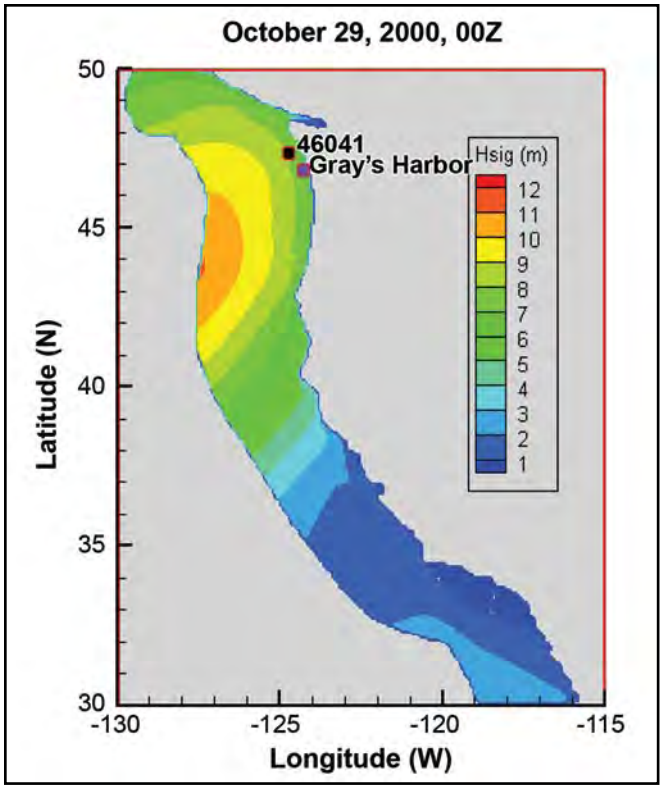

Figure 7. Contoured significant wave height results for the 1/12-deg resolution GRID 3

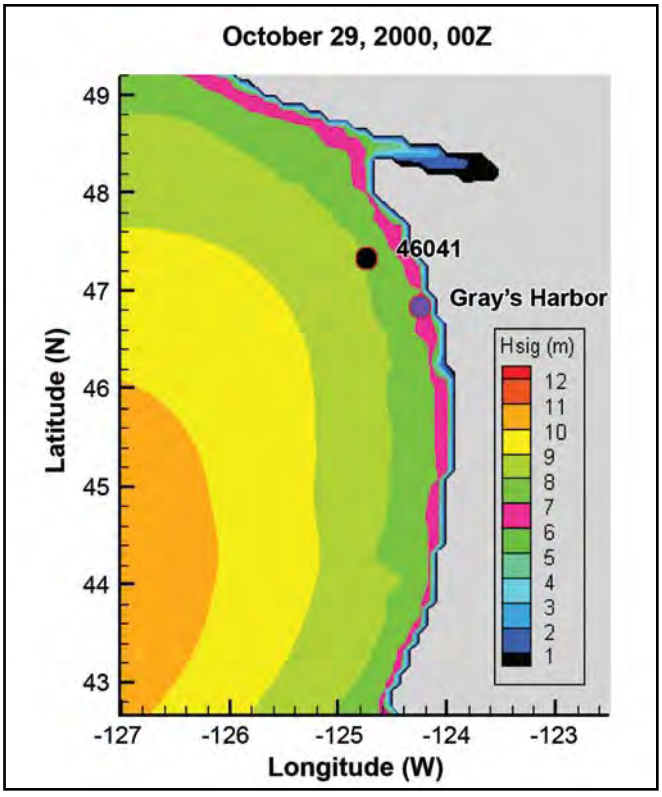

Figure 8. Contoured significant wave height results for the Washington and Oregon coast in the 1/12-deg resolution GRID 3

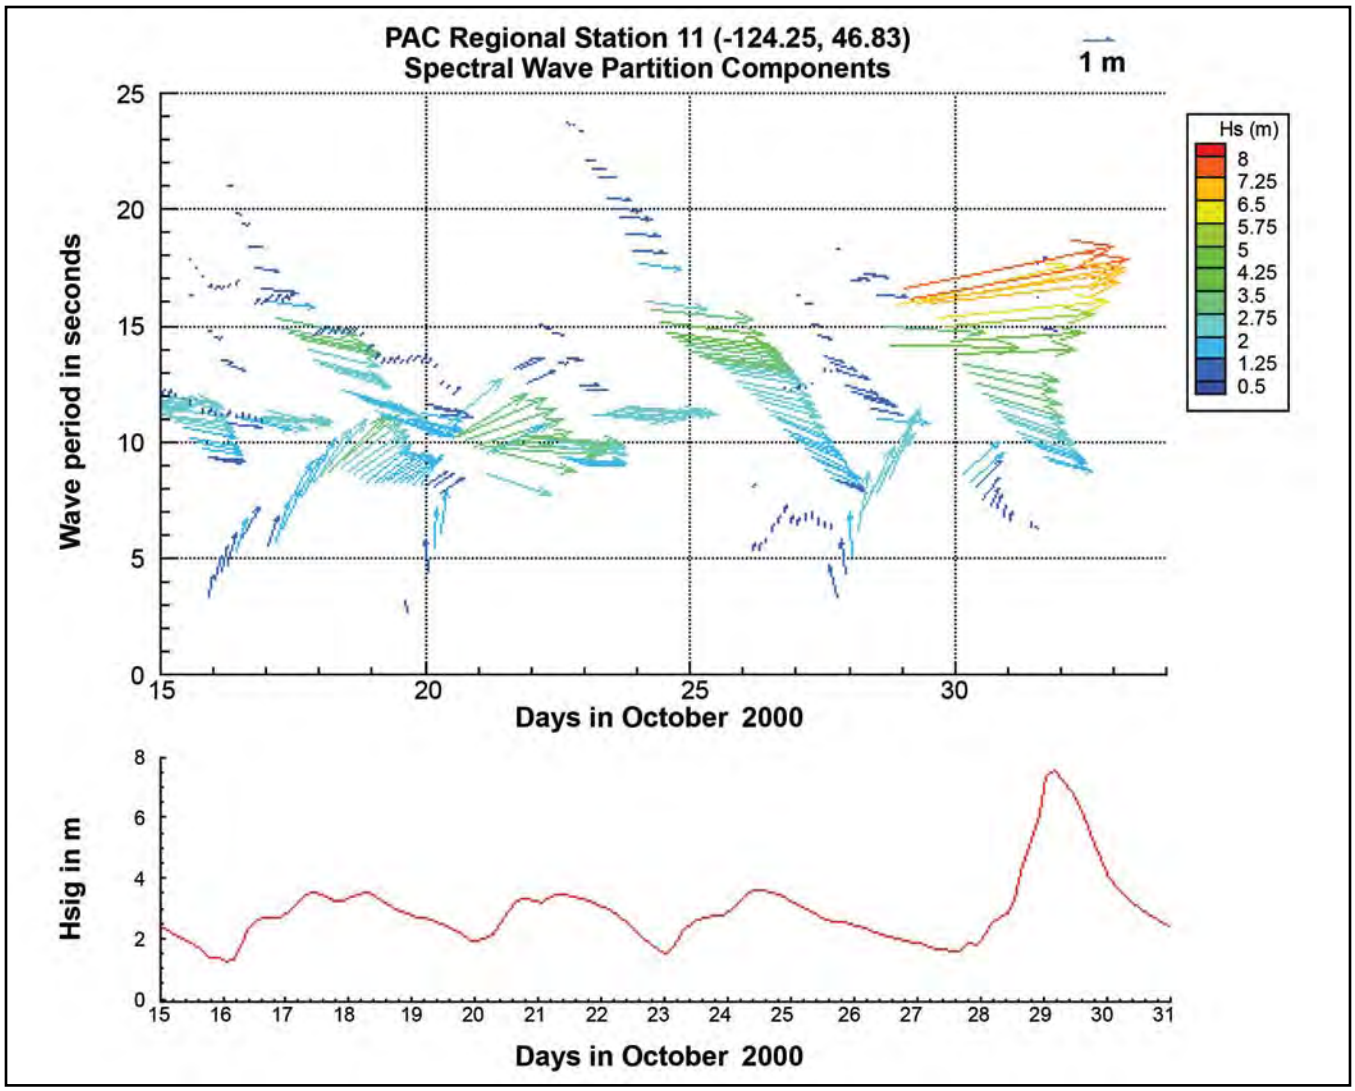

Figure 9. October 15-30, 2000, plot of spectral wave partition components for PAC regional station 11 near Gray's Harbor, WA. A significant wave height plot is included below the vector plot for reference

Tolman, H. L. (2002a): User manual and system documentation of WAVEWATCH-III version 2.22. Technical Note 222, U.S. Department of Commerce, NOAA, NWS, NCEP, Washington, D.C.

Tolman, H. L. (2007): User manual and system documentation of WAVEWATCH-III version 3.14. In preparation, U.S. Department of Commerce, NOAA, NWS, NCEP, Washington, D.C.

Tolman, H. L. (2003): Treatment of unresolved islands and sea ice in wind wave models, *Ocean Modeling*, 5, 219-231.

Tracy, B.A. and D. Spindler (2008): WIS Pacific Regional Hindcast. *Proceedings of HPC Users Group Conference, July 2008, Seattle, WA.*

Tracy, B.A., E. M. Devaliere, T. Nicolini, H. L. Tolman, and J. L. Hanson (2007): Wind sea and swell delineation for numerical wave modeling, 10th International Workshop on Wave Hindcasting and Forecasting, November 11-16, North Shore, Oahu, HI.

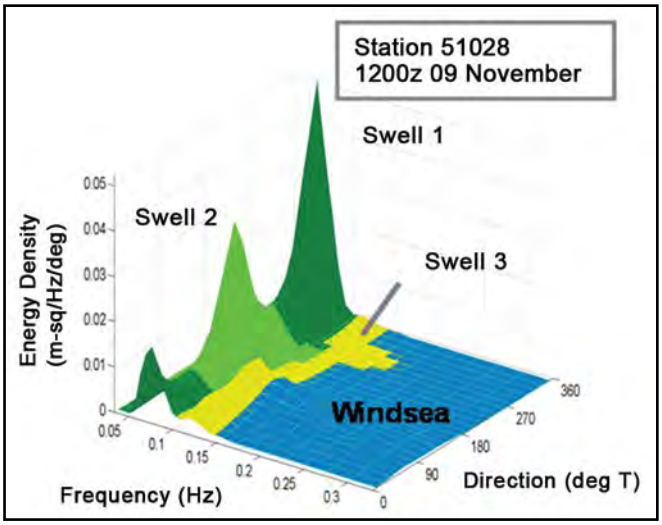

Figure 10. Example of a partitioned wave energy spectrum showing one area of wind sea and three swell partitions

Higher Order Frequency and Time-Domain Seismic/ Acoustic Modeling for UGS Applications

By Saikat Dey, Naval Research Laboratory/SFA Inc.; Charbel Farhat, Stanford University; Michael W. Parker and Stephen A. Ketcham, ERDC Cold Regions Research and Engineering Laboratory; Christopher Kun, User Productivity Enhancement and Technology Transfer (PET), Army Research Laboratory; and Steven Wong, PET, Air Force Research Laboratory

Introduction

Protecting Army personnel and facilities from ground attacks using unattended ground sensor (UGS) systems relies on the system's ability to identify acoustic and seismic signatures of potential threats. Effective time- and frequency-domain computer simulations of the wide seismic signature space are required because of the high cost of experimental characterization. This problem of acoustic and seismic wave propagation is similar in nature to underwater acoustic radiation and scattering encountered in many Navy applications. Recognizing this commonality, under the auspices of the DoD High Performance Computing Modernization Program (HPCMP) User Productivity Enhancement and Technology Transfer (PET) initiative, researchers at the Naval Research Laboratory (NRL), ERDC Cold Regions Research and Engineering Laboratory (CRREL), and Stanford University have been collaborating to extend and apply the Navy's structural acoustics code STARS3D to UGS applications.

STARS3D originated as a frequency-domain finite element simulation code for computing the response of coupled and uncoupled underwater fluid-elastic systems to acoustic, mechanical, and flow-induced excitations. Its algorithms use adaptable hp-refinement elements to achieve a high level of accuracy, meaning that both the element size and basis function order are adaptable. For code parallelization, STARS3D utilizes a technique called the Finite Element Tearing and Interconnecting (FETI) method to decompose the domain so that calculations in the frequency domain can be separated and calculated independently over numerous processors on HPC systems. The frequency-domain code provides the ability to compute frequency responses accurately for geometrically complex large-scale models. For the study of certain UGS systems

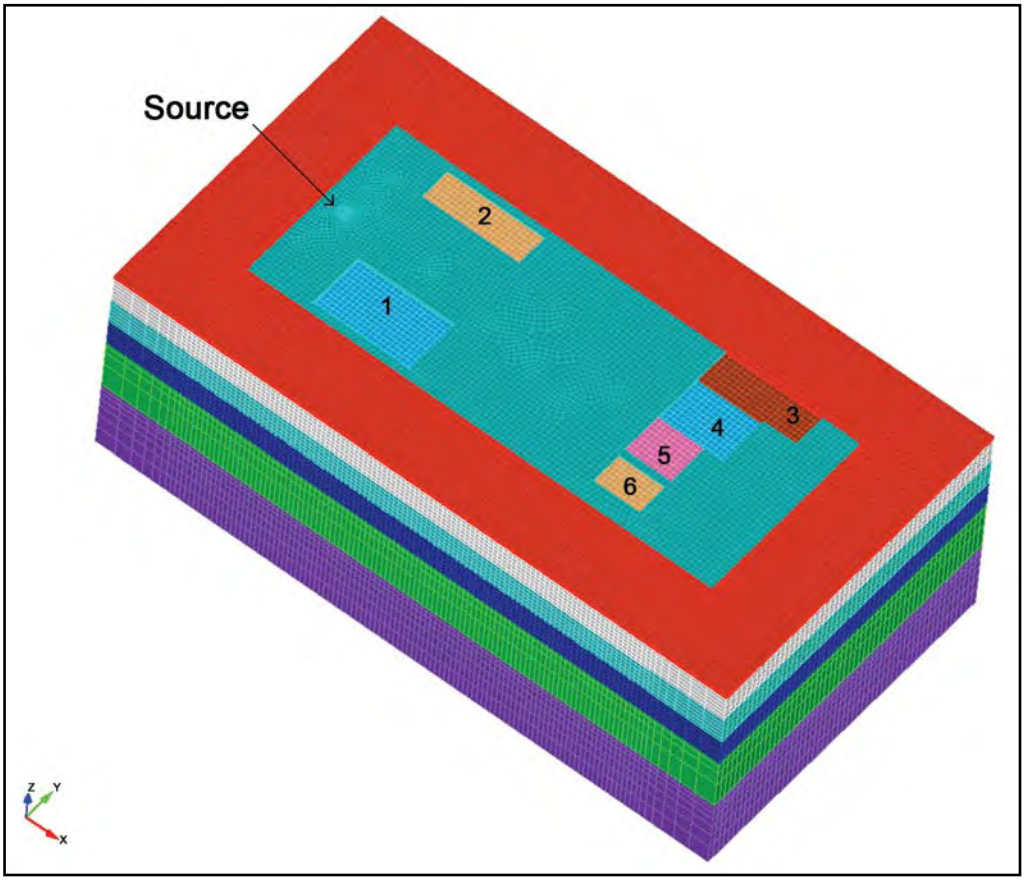

Figure 1. Soil domain with foundation and source location

Table 1. Soil Layer Properties				
Layer	Depth (m)	Density (Kg/m^3)	Compression (Shear) wave speed (m/s)	Damping Parameters
1	2	1600	420 (210)	0.067
2	12	1600	460 (230)	0.04
3	12	2000	490 (280)	0.025
4	12	2000	525 (300)	0.02
5	∞	2400	3800 (2000)	0.029

involving acoustically large problems of interest to the Army, time-domain techniques are more efficient than frequency-domain techniques because of the elimination of repeated matrix-factoring, a numerically expensive calculation. The simulation of time-dependent loads and excitations are easily implemented in the time domain as well. In addition, a time-domain simulation can provide visual verification that is not as easily recognized in the frequency domain.

In a collaborative effort among NRL, PET, and CRREL, a time-domain version of STARS3D has been produced for use in UGS applications. Explicit and implicit time-domain update schemes are incorporated into STARS3D using concepts based upon the Newmark family of time integration, which is widely used in time-dependent elastodynamics and electromagnetics. This methodology allows the user to adjust several parameters that lead to different time-updating schemes with different properties with respect to accuracy and stability. A fourth-order explicit update scheme was also implemented that showed increased accuracy over the second-order method in certain metrics. Also, to truncate the (semi) infinite computational domain, a second-order time-dependent absorbing boundary condition (ABC) is incorporated, which has been tested at Stanford University and proven to work with acoustic and elastodynamic applications.

This article first presents the application of the frequency-domain capability to study the scattering of seismic waves because of shallow building foundations. Then it presents an application of the new time-domain STARS3D for the computation of acoustic wave propagation in an urban environment setting. These studies were performed on the SGI Altix system (Eagle) at the Air Force Research Laboratory (AFRL) MSRC.

Scattering of Seismic Waves by Shallow Building Foundations

The goal of this example is to use STARS3D to study scattering of seismic waves by urban structures, in particular shallow building foundations. In this numerical experiment, a five-layer semi-infinite soil model with six building foundations was considered in a domain that extends 200 meters in the x-direction and 80 meters in the y-direction. Figure 1 shows the computational domain, and Table 1 describes the properties of each soil layer starting from the top.

The domain was excited by a time-harmonic normal traction source applied to a small square patch with a side length of 1 meter on the ground surface. The study first computed the frequency response of the soil domain without any foundations and then compared it with that when a set of six building foundations were added. Additionally, two sets of depths for the six foundations are considered given by {2, 1, 1.5, 2, 0.5, 1} and {3, 4, 3, 1, 2, 4} meters, respectively.

In Figures 2-5, the difference in the frequency response because of the presence of the foundations as compared with the baseline case without the foundations is plotted. The difference is plotted on the dB-scale as $20 \log_{10}(x/x_0)$ where x is the response in the presence of the foundations and x_0 is the response in the absence of the foundations. The top plot in Figures 2-5 shows the displacement in the x-direction, and the bottom plot shows the displacement in the z-direction.

One of the goals of such studies is to study the various constructive and destructive interferences that occur because of the presence of buildings and foundations to guide proper locations for sensor placements.

Convergence Study of an Urban Acoustic Model

To verify the capability of STARS3D in simulating acoustically large problems, a time-domain acoustic scattering problem is examined with a large urban acoustic model. The domain consists of the above-ground air surrounding three buildings (Figure 6). The radius of the hemispherical volume is 170 meters. The height of the cylindrical tower is 120 meters; the height

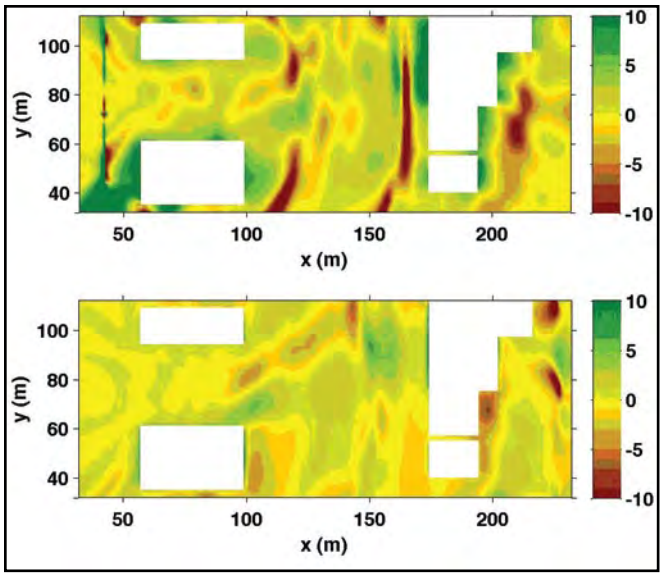

Figure 2. Comparison of relative power in dB, 20 $\log_{10}(|U_i/U_0|)$ and 20 $\log_{10}(|W_i/W_0|)$ at ground level, f = 12 Hz, and foundation depths of 0.5-2 meters

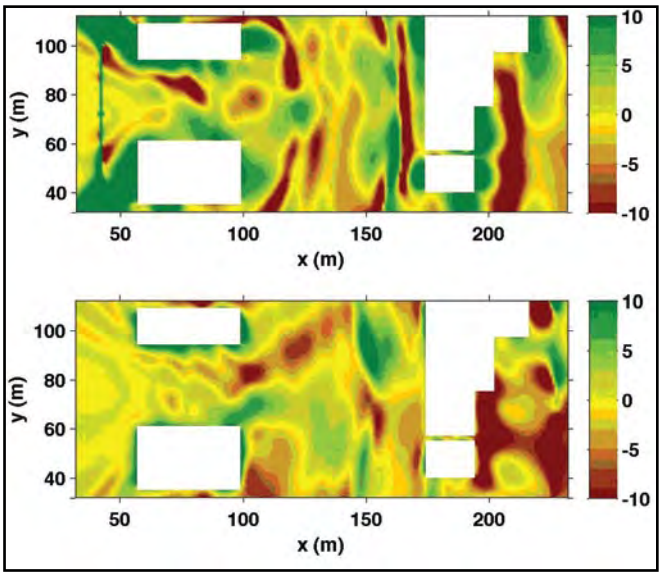

Figure 3. Comparison of relative power in dB, 20 $\log_{10}(|U_i/U_0|)$ and 20 $\log_{10}(|W_i/W_0|)$ at ground level, f = 12 Hz, and foundation depths of 1-4 meters

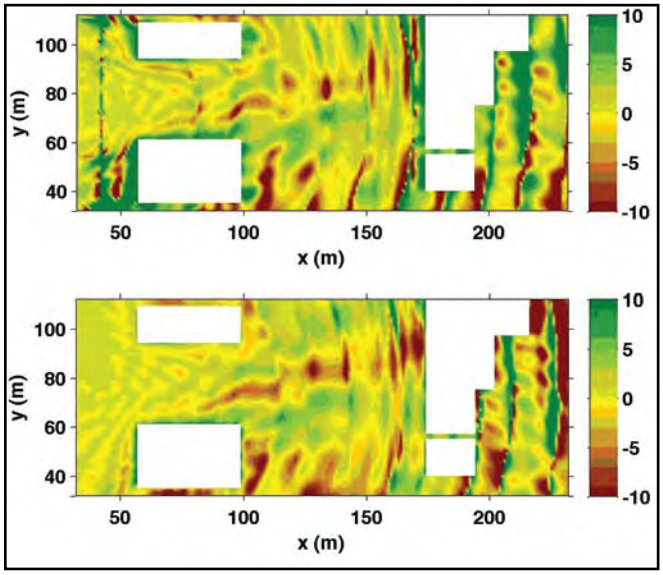

Figure 4. Comparison of relative power in dB, 20 $\log_{10}(|U_i/U_0|)$ and 20 $\log_{10}(|W_i/W_0|)$ at ground level, f = 24 Hz, and foundation depths of 0.5-2 meters

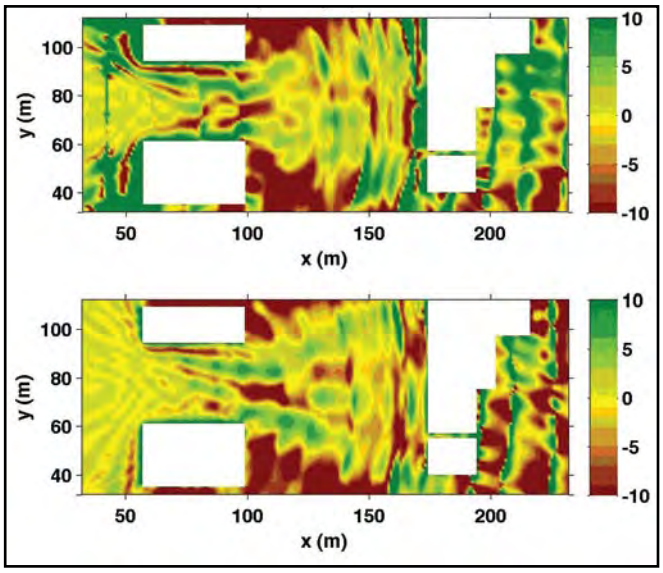

Figure 5. Comparison of relative power in dB, 20 $\log_{10}(|U_i/U_0|)$ and 20 $\log_{10}(|W_i/W_0|)$ at ground level, f = 24 Hz, and foundation depths of 1-4 meters

of the two other buildings is 100 and 120 meters; and the spacing between the buildings is of the order of 20 meters. An acoustic excitation $\frac{\partial p}{\partial n} = \delta(\omega)$ is prescribed over a spherical surface radius of 2 meters with the center located 5 meters above the ground surface.

Since an analytical solution is not available for comparison, a convergence study was done using the urban domain. Figures 7 and 8 plot the pressure for the mesh shown in Figure 6 with $p = 2$ and $p = 3$, where p refers to the approximation order of the finite element basis function being used, at time-steps 20 and 1000, respectively. One observes that the convergence is much better early in the simulation, and this is expected because the time integrators accumulate more errors with increasing time-steps.

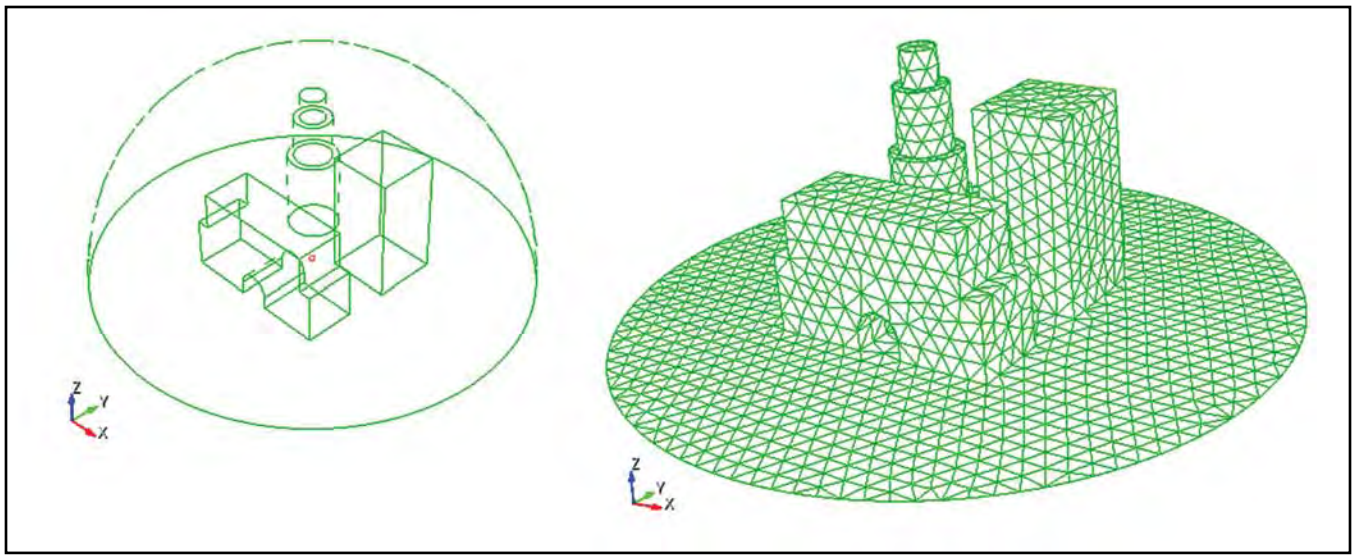

Figure 6. Urban acoustic domain (left) and meshed urban domain (right)

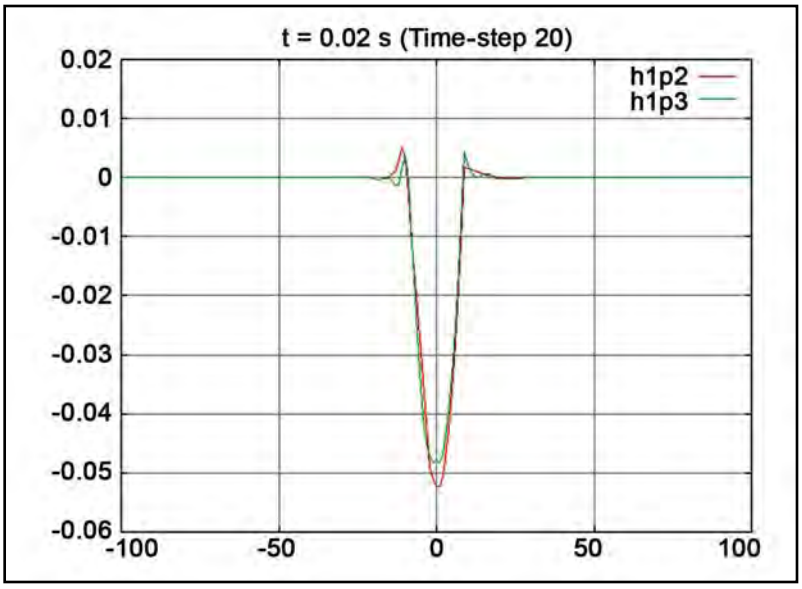

Figure 7. Pressure on groundline for p = 2, 3 at time-step = 20

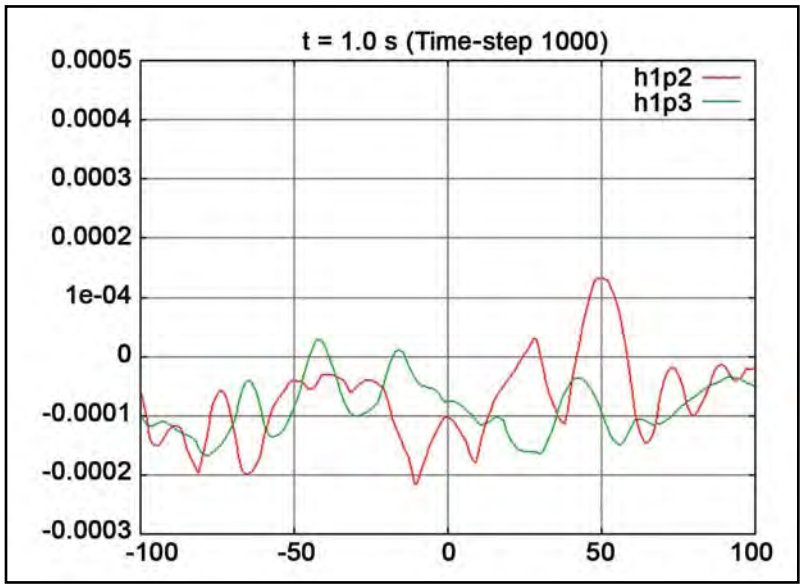

Figure 8. Pressure on groundline, p = 2, 3 at time-step = 1000

Figure 9 plots the pressure isosurfaces at time-steps 20, 100, 500, and 1000 with $p = 3$, and Figure 10 plots the pressure isovalues on planes at ground level. The diffraction of the acoustic wave fronts around the buildings can be seen from these plots.

Concluding Remarks

STARS3D has now been expanded to incorporate time-domain capabilities alongside its frequency-domain capabilities. Future work to be done includes parallelizing the time-domain capability so that it can take better advantage of HPC computing resources. The parallelization of the time-domain code will allow for the simulation of even larger acoustically large problems.

Acknowledgment

This work was supported by the DoD HPCMP under the PET projects CEA-KY6-001 and CEA-KY7-001.

References

Dey, S., et al. (2001). "Mid-frequency structural acoustic and vibration analysis in arbitrary, curved three-dimensional domains," *Computers and Structures*, 79, 617-629.

Dey, S. and Datta, D. K. (2006). "A parallel *hp*-FEM infrastructure for three-dimensional structural acoustics," *Int. J. Num. Meth. Eng*, 68, 583-603.

Parker, M. W., Ketcham S. A. and Dey, S. Scattering of seismic waves by shallow building foundations using high-order FEM. *Proceedings of HPCMP Users Group Conference 2008*.

Hughes, T. J. R. (1987). "The finite element method: Linear static and dynamic finite element analysis," Prentice-Hall, New Jersey.

Newmark, N. M., (1959). "A Method of Computation for Structural Dynamics," *ASCE Journal of the Engineering Mechanics Division*, 85, 67-94.

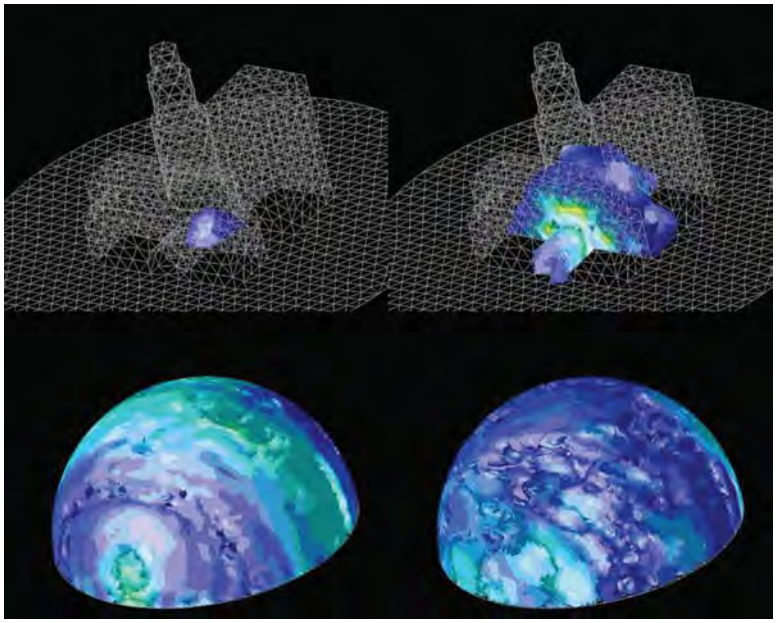

Figure 9. Pressure isosurfaces at time-steps 20, 100, 500, and 1000

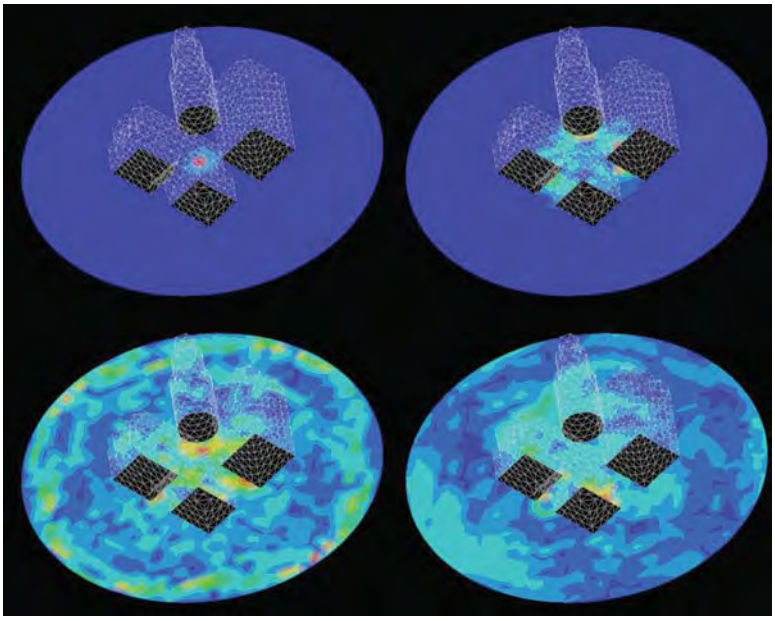

Figure 10. Pressure isovalues at ground level (z = 0) at time-steps 20, 100, 500, and 1000

Jade Becoming Preferred Gem—Users Praise Cray XT4

By David Dumas, Phillip Bucci, and Kenneth Matthews, ERDC MSRC

The Cray XT4 at the ERDC MSRC officially entered production at 0001 hours, Tuesday, July 2. Known as Jade, the system contains 2,152 compute nodes, each containing one 2.1-GHz AMD Opteron 64-bit quad-core processor and 8 GBytes of memory. When Jade's compute pool processor count of 8,608 is combined with Sapphire's 8,320 processors, the computing power at the ERDC MSRC is truly impressive. With uptime since entering production exceeding 99 percent, Jade is quickly becoming a preferred HPCMP resource among our users.

Comments from users about the system are positive. "Overall, Jade has been a very nice machine to work on," said Joseph Metzger of the Naval Research Laboratory. "Earlier this summer, we had ~2 million hours under a CAP [Capabilities Application Project] that we used in its entirety. The main goal was to implement tidal forcing in our 1/12 deg global HYbrid Coordinate Ocean Model (HYCOM)—a first for such a high-resolution global ocean model. The model by itself outputs a lot of data because of the array sizes (4500×3298×32), and we typically output once per model day. The inclusion of tides requires us to output information hourly. It was still not feasible to do this for the whole domain for the duration of the experiment, but we did create hourly whole domain output for a single month (~7 TB) and hourly subregion output for the entire simulation."

"Jade is a much more pleasant machine to run on than Sapphire."
— Joseph Metzger
Naval Research Laboratory

Metzger explained that his experiment produced enormous amounts of output, almost 40 TBytes, which the /work filesystem on Jade ultimately processed successfully and supported a relatively painless transfer of said data to his archive storage at the Naval Oceanographic

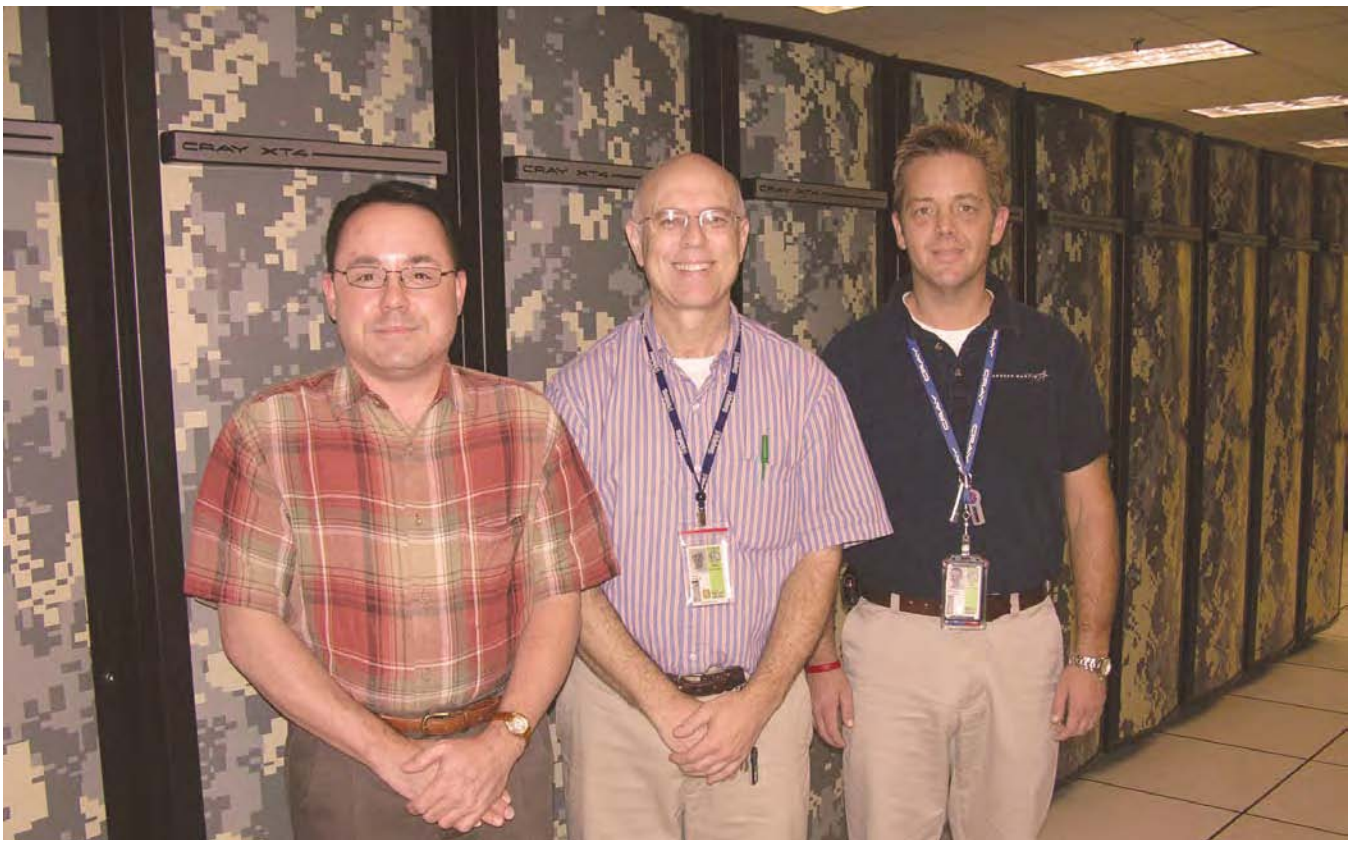

(Left to right) David Dumas, ERDC MSRC HPC Service Center Specialist; Phillip Bucci, ERDC Service/Agency Approval Authority and ERDC HPC Allocations Officer Support; and Kenneth Matthews, ERDC MSRC Senior Database Engineer, Systems Administration Group

Office (NAVO) MSRC. "I had been running a similar experiment on Sapphire (although it didn't output nearly as much data), and it seemed the /work filesystem would often get corrupted and I'd have to restart my simulation," he said in a comparison to the XT3. "Jade is a much more pleasant machine to run on than Sapphire."

"I found the machine easy to work with and required very little code modification to get everything to run stably."

— Dr. James Freericks
Department of Physics
Georgetown University

Under his CAP Phase II project, Dr. James Freericks, Department of Physics, Georgetown University, said, "I found the machine easy to work with and required very little code modification to get everything to run stably. I was able to achieve a performance on the order of 40 percent of theoretical peak efficiency with better than 90 percent of linear scaling of the code." Dr. Freericks' codes are written in Fortran 77 using standard LAPACK and BLAS routines and run under the MPI environment. According to Freericks, his code is quite transportable, and he rarely has porting issues. "My main objective is to try to sort out appropriate compiler flags to get the best efficiency and the best libraries to use for the codes," he said. "In the end, I achieved excellent performance, and I will be using Jade significantly over the coming year with a challenge project I have and am enthusiastic about the opportunity to do so."

Dr. Jim Pfaendtner, Department of Chemistry, University of Utah, contacted the ERDC MSRC HPC Service Center for help in compiling his customized NAMD code for Jade. "I requested support from the ERDC [MSRC] help desk, and within 1 week I had the code compiled efficiently and running well. My code now scales well up to 1,000 processors, and I have very efficient and fast computation. I particularly like using Jade because of its large size. My research group has developed a technique for running multiple simulations simultaneously which exchange a limited amount of information. These simulations are ideal for Jade, as on most days I am able to launch a large number of medium-sized jobs and get results quickly. I have used Jade for these types of medium-sized simulations as well as large simulations on up to 1,000 cores, and I am very pleased with the results. The file systems work very well and the archival system is fast and easy to use. Overall, having Jade available for my research improves the quality of my work, and I look forward to using it in the next FY."

Jade is also playing a vital role in support of the U.S. Army Corps of Engineers (USACE) efforts in hurricane storm surge modeling for flood plain mapping, design of flood protection structures, and exploration of the role of natural impediments, such as marshes and barrier islands, in the role of surge attenuation. According to Dr. John Atkinson, Computational Fluid Dynamicist with the ARCADIS consulting firm, "the New Orleans District in particular is very concerned with designing flood protection levees to shelter people and property from hurricane storm surge; thus, accurate hurricane storm surge prediction is critical to that aspect of USACE's mission."

"the New Orleans District in particular is very concerned with designing flood protection levees to shelter people and property from hurricane storm surge; thus, accurate hurricane storm surge prediction is critical to that aspect of USACE's mission."

— Dr. John Atkinson
Computational Fluid Dynamicist
ARCADIS consulting firm

Atkinson further explained that the ADCIRC finite element hydrodynamic computer model is used on Jade to accurately simulate the complex tidal, wind, and wave-driven circulation as well as the hurricane-generated storm surge on the coastal flood plain. "With the large number of equations to solve for these simulations, Jade has become an important platform for our work," he said. "Many of our simulations are needed by designers and contractors with tight schedules for constructing vital flood protection barriers in advance of next season's hurricane season. Timely delivery of the many hundreds of simulations has been made possible with platforms like Jade and Sapphire, which allow us to use up to a thousand processors per simulation."

Without a doubt, Jade is serving the DoD research and USACE civil works communities well. The ease of porting code to the architecture, the system's consistent uptime, and the short queue waits are attracting new users every day. The ERDC MSRC is proud to be home for this Cray XT4. For more information on Jade's configuration, guidelines for using the system, and other tips and tricks, please visit the ERDC MSRC Web site at *http://www.erdc.hpc.mil*, and click on Documentation.

Large-Scale Performance on the Cray XT3 and XT4 Systems

By Dr. Kent T. Danielson, ERDC Geotechnical and Structures Laboratory

Introduction

This article describes a large-scale benchmark and performance study on the Cray XT3 and XT4 (Sapphire and Jade) at the ERDC MSRC with regards to large single-, dual-, and quad-core counts as well as to compiler type. Three important observations were made from this study with runs using up to the full capacity of each machine:

- The complete 8,000+ cores of the XT3 and XT4 systems can perform well in a single analysis.
- Provided memory limits are not exceeded, separate processes on individual cores of a node can be much more efficient (nearly perfect speedup) than in single-core mode (e.g., `-SN` option).
- Executables using the PathScale compiler consistently outperformed the other compilers by as much as 70 percent.

Benchmark Description

The benchmarks were made with an analysis of an explosive detonation in a reinforced concrete bridge pier using the parallel explicit dynamic finite element code, ParaAble, developed by the author and described in (Danielson et al. 2007, 2008). In all cases, the problem size is fixed with only computational options as variables. ParaAble is a three-dimensional transient Lagrangian solid dynamics program written in FORTRAN 95. All interprocessor communication can be made with explicit Message Passing Interface (MPI) calls or with a hybrid MPI/SHMEM option using a Single Program Multiple Data (SPMD) paradigm. The default Portland Group, the PathScale, and the GNU compilers were considered with comparably aggressive but relatively simple optimization flags of `-O3 -fastsse`, `-Ofast`, and `-O3` for each, respectively. Execution was made with processes both on only a single core and on multiple cores of a node, `-SN` and `-VN` options on the parallel execution command. Note that the runs were real-world application analyses complete with all input and output of model information and visualization plot data, time-history data, and restart files. The MPI-only version was used for all cases, and the top speed was performed at nearly 10 teraFLOPS.

System Limitation, Compiler Bug, and Workarounds

During the application of ParaAble to thousands of processors on the Cray XT3, it was noticed that All-To-All type calls, e.g., `MPI_ALLTOALL` or `MPI_ALLGATHER`, would overwhelm the communication channels. These types of collective communications work quite well with small processor counts, but the number of interprocessor messages produced by these calls is the number of processors squared. Previous MPI implementations originally intended for hundreds of processors thus experience several orders of magnitudes more messages for these types of communications on thousands of processors. Whereas the MPI call is the correct syntax, its use can exceed the system's capacity on the number of posted messages. Cray has modified its implementation for some of these operations, and simple workarounds were made for the others by replacing, for example, an `MPI_ALLGATHER` with an `MPI_GATHER` nested within a DO loop

A bug in the parallel default Portland Group FORTRAN compiler was also found early in the XT3 installation and apparently fixed, but unfortunately this problem persists on the XT4. In FORTRAN 90 and later, the system must name files defined as scratch, but the FORTRAN standard does not consider independent parallel processes. If the message passing, e.g., MPI or SHMEM, implementation of the compiler does not coordinate separate scratch file names among processes, they can write over each other with disastrous results. That is, each process is defining the file in the same way in parallel and thus is giving the same name to the scratch file of each process. The PathScale and GNU compilers work properly with parallel scratch files. The easy fix is to simply append the automated filename with a processor number, which is also the basis of the following XT4 workaround that names a file and then deletes it upon closing. The following code fragment demonstrates this easy modification that is essentially benign to performance in most cases. The drawbacks with this approach are that the files remain if the `CLOSE` statement is not executed (because of a system crash, segmentation fault, etc.), and additional care is needed if multiple runs are executed concurrently in the same directory.

```
LOGICAL, PARAMETER :: SCR_AUTO = .FALSE.
CHARACTER (LEN=10) :: CMYPE

WRITE(CMYPE,"(I10)")MYPE;
CMYPE = ADJUSTL(CMYPE)

IF(SCR_AUTO)THEN
    OPEN(UNIT=1,STATUS='SCRATCH')
ELSE
    OPEN(UNIT=1,STATUS='REPLACE',
    FILE='Scratch.'//TRIM(CMYPE))
END IF

        … other source coding …

CLOSE(1,STATUS='DELETE')
```

Performance Comparisons

Figure 1 compares scalability of the XT3 and XT4 in multicore mode (dual and quad core, respectively) using the default Portland Group compilers. Note that both scale well and that the differences between the XT3 and XT4 performances were consistently at about the differences in their processors speeds of 2.6-GHz and 2.1-GHz, respectively. Whereas the XT4 is slower per core, it has twice as many cores per node—thus performing at about 1.7 times faster per node than the XT3. Using the default Portland Group compilers, Figures 2 and 3 compare the scalabilities of the XT3 and XT4 in single-core versus multicore mode (dual and quad core, respectively). Note that all cases scale well and that sharing memory/cache on a node by going from single core to multicore had little effect on performance. Clearly, using individual cores is much more efficient here than utilizing a node's entire memory and cache on a single core. In addition, the scaling is so good that multithreading or other inner loop parallelism on nodes would not provide significant improvements on these scalar machines. Of course, this is just one example, but other applications may be more memory intensive and may even need the memory of the entire node. Indeed, application herein to large numbers of processors for the same size problem naturally breaks the analysis into much smaller partitions with less memory requirements. This obviously can improve scalability or even produce greater than linear speedups and is thus another advantage of large numbers of processors. The author's experiences are that the analysis needs to be large enough (e.g., ~500 elements) on each process so as to not be dominated by interprocessor communication but that can still fit on a single core.

Figure 4 compares different compilers on the XT4. Again, all cases scale well, and the PathScale compiler is clearly outperforming the others on the AMD Opteron processors. It is not surprising that GNU performed the worst; however, the PathScale compiler being nearly three times faster than the GNU one and still almost 70 percent faster than the commercial Portland Group compiler was certainly unexpected. This is especially astonishing since it only required simple changes in the compiler options and switching to the PathScale compiler with the simple module swap command:

```
jade01% module swap PrgEnv-pgi PrgEnv-path-
scale
```

The same Cray compiler commands, e.g., `ftn`, are used in the same way as with the Portland Group compiler. Figure 4 only considers a single application of fixed size, so such differences may not always result.

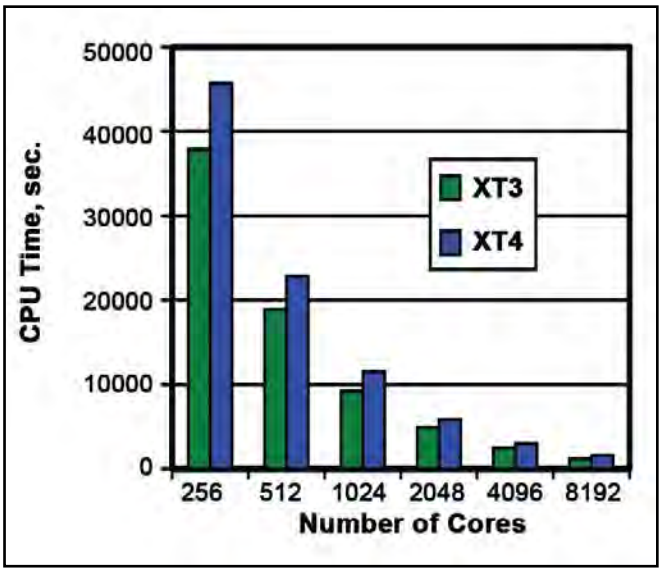

Figure 1. Scalability of Portland Group (default) compilers on the Cray XT3 and XT4

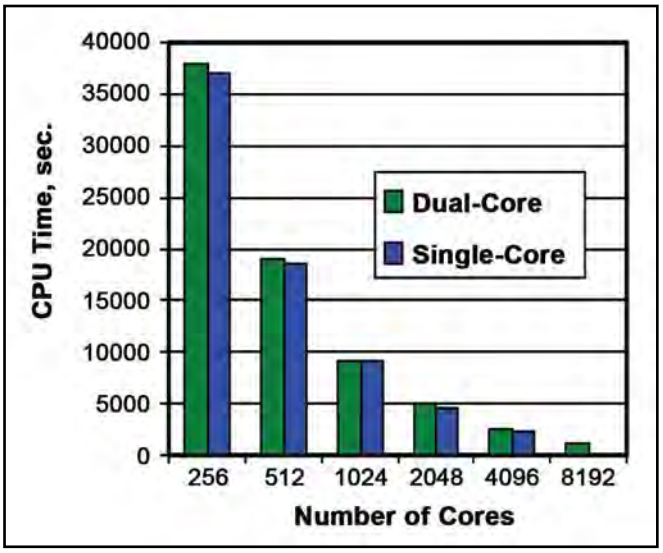

Figure 2. Scalability of Portland Group compiler on the Cray XT3 in single- and dual-core modes

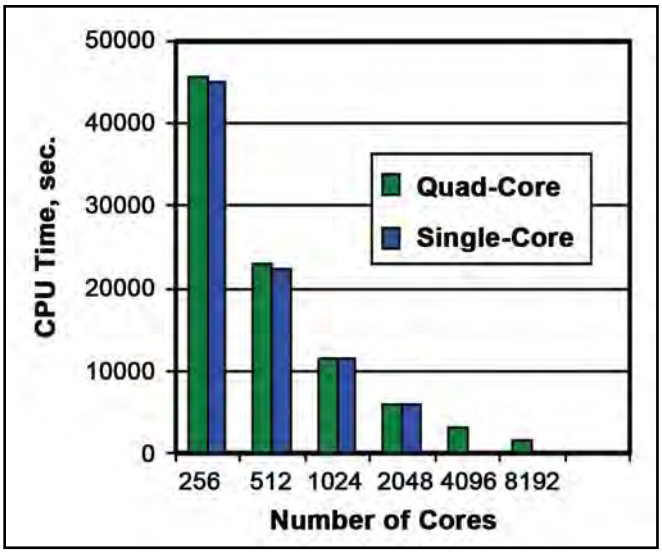

Figure 3. Scalability of Portland Group compiler on the Cray XT4 in single- and quad-core modes

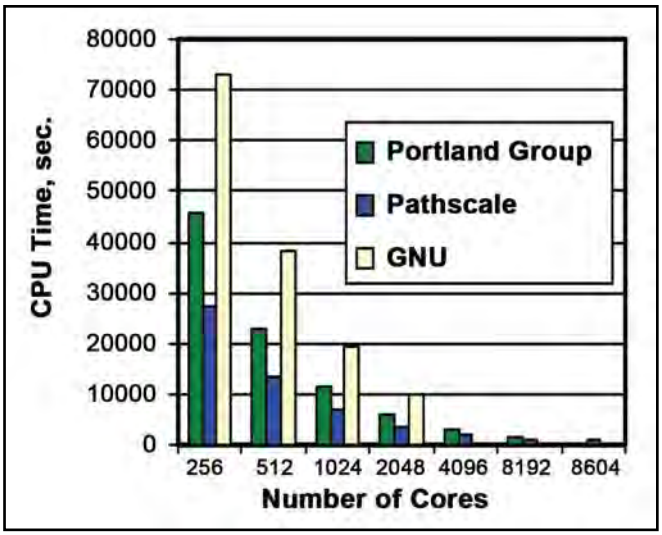

Figure 4. Scalability of different compilers on the Cray XT4 in quad-core mode

The performances reported herein also considered only a single optimization. The author has found, however, by running many other different analyses, trying different optimization flags and using completely different codes, some written by others and in other programming languages, that the PathScale compiler creates executables that typically outperform those created by the Portland Group compiler by at least 40-50 percent.

Concluding Remarks

Large-scale application of a nonlinear finite element analysis demonstrated that large numbers of cores, including the complete 8,000+ cores of the XT3 and XT4 systems at the ERDC MSRC, can perform well in a single analysis. If memory size is not a problem, separate processes on individual cores of a node were shown to be much more efficient (nearly perfect speedup) than in single-core mode. Executables created using the PathScale compiler consistently outperformed the other compilers by as much as 70 percent. This small study and limited other experiences certainly do not mean that the PathScale compiler is the best for all applications. It does mean, however, that tremendous benefits can sometimes be obtained by changing compilers. Especially since it is quite easy to do, code installers might find it beneficial to simply try the different available compilers.

Acknowledgments

Permission to publish this article was granted by the Director of the ERDC Geotechnical and Structures Laboratory.

References

Danielson, K. T., O'Daniel, J. L., Adley, M. D., Akers, S. A., and Garner. S. B. "Large Cray XT3 System Greatly Increases ERDC Explosive—Concrete Modeling Capabilities." ERDC MSRC *Resource*, Spring 2007.

Danielson, K. T., Akers, S.A., O'Daniel, J. L., Adley, M. D., and Garner, S. B. (2008). "Large-Scale Parallel Computation Methodologies for Highly Nonlinear Concrete and Soil Applications." *ASCE Journal of Computing in Civil Engineering* 22(2) 140-146.

ERDC Infrastructure – BUILDING STRONG!

By Chad Christophersen, ERDC MSRC

This past August, Lieutenant General Robert Van Antwerp, the Chief of the U.S. Army Corps of Engineers, announced the new Corps slogan – BUILDING STRONG! The steady hum of bulldozers in the backyard of the ERDC Information Technology Laboratory (ITL) leaves little doubt that the Laboratory and the ERDC MSRC are busy breathing life into those emblematic words. In November 2007, ITL and the MSRC began shaping the Lab's collective future when construction began on infrastructure (facility, power, and cooling) improvements intended to accommodate the next generations of high performance computers. That effort, managed by the U.S. Army Engineer District, Vicksburg, now continues in full force.

Project requirements are based on the ERDC MSRC analysis of past DoD High Performance Computing Modernization Program (HPCMP) Technology Insertion (TI) acquisitions of high performance computers as well as industry trends and forecasts. As a rule of thumb, system capability of TI-acquired systems approximately doubles every 2 years. More capable systems are generally also more demanding systems when it comes to infrastructure. Analysis revealed that ITL would need to add 4.0 megawatts (MW) of power and 1000 tons (T) of chilled water cooling to the existing Laboratory infrastructure (~8.0 MW of power and ~1200 T of chilled water cooling).

Once complete, the planned work will yield the additional space, power, and cooling capacity that the Laboratory and MSRC need in order to welcome future HPC systems with open arms.

Planned facility improvements include a 10,000 square foot Information Technology (IT) shelter intended to house the MSRC's anticipated TI-09 (TI-Fiscal Year 2009) HPC system(s) and other future HPC resources. The shelter will roughly double the Laboratory's current raised floor compute space.

Planned power modifications include two self-contained, 2 MW generator sets; battery backup modules; an Uninterruptible Power Supply (UPS) system scalable to 8 MW; a 2500 kilovolt-ampere (kVA) transformer; and associated infrastructure to deliver power to the new IT shelter. The UPS system will provide immediate, high-capacity, quick-discharge power protection for HPC resources in the event of power disturbances, voltage sags, surges, momentary disruptions, and complete outages.

Planned cooling enhancements include two 500 ton chillers; piping; and a 49,200 gallon chilled water reservoir composed of four upright thermal storage tanks, each 29 feet tall and 9 feet in diameter. The reservoir will provide 15 minutes of backup cooling to HPC systems in the event of a chiller failure. "Now that's pretty cool," comments Travis Vance, an important member of the ITL and MSRC facility team.

Today, construction continues, and progress is clear. The utility corridor that will allow the planned power and cooling pad to tie in to the existing laboratory building is nearly complete, and construction of the IT shelter is projected to begin this fall. When the dust settles in June of 2009, the Laboratory, MSRC, and Vicksburg District, will have collaboratively created an expanded ITL infrastructure capable of sustaining ERDC and DoD HPC efforts well into the foreseeable future. The Corps of Engineers is BUILDING STRONG infrastructure today, so that DoD engineers and scientists can build an even stronger tomorrow for the Nation.

Utility corridor begins to take shape

Fostering Next Generation of Scientists and Engineers

By Rose J. Dykes

The future of the United States requires that more and more students prepare themselves for math and science careers. In fact, the U.S. is desperate to improve its rankings in these two fields to remain competitive in the global market. The ERDC MSRC attempts to take every available opportunity to encourage students to enter these fields by offering internships, mentoring high school and undergraduate students by participating in HPC institutes, engineering camps, and shadowing programs, and encouraging team members to serve as part-time professors in local colleges and universities, including minority-serving institutions.

A West Point cadet information technology (IT) major, Matt Rogers, joined the Data Analysis and Assessment Center (DAAC) in the field of computer graphics and visualization at ERDC during this past summer as an intern through the Advanced Individual Academic Program at the United States Military Academy. Rogers remarked, "The time spent here has been very insightful for me. I have gained firsthand experience into what graphics and visualization are and how information technology professionals use their skills to benefit our country."

West Point cadet Matt Rogers works as summer intern in DAAC

David Stinson with JSU students participating in NSF Research Experiences for Undergraduates Program

Jackson State University (JSU) students participating in the National Science Foundation (NSF) Research Experiences for Undergraduates Program visited the ERDC MSRC this past June. David Stinson, MSRC Acting Director, talked with the group about HPC, explaining that many complex scientific problems can only be solved with the use of supercomputers. He then showed them the supercomputer infrastructure at ERDC. Afterwards, Paul Adams, Visualization Team Lead, took the group through the DAAC, telling them that HPC users often require high-quality images or animations of their computer data to show the data in context or natural environment.

Shaping the Future

The ERDC MSRC also participated in two programs for younger students this past June. Team members conducted classes on HPC and scientific visualization and took on tours 40 high school students from all over the U.S. who were chosen to participate in the Society of American Military Engineers (SAME) Engineering and Construction Camp. This annual, week-long camp held in Vicksburg is one of three held throughout the U.S. The Vicksburg camp provides high school students with hands-on experience in engineering and construction activities and exposes them to a wealth of knowledge about the various careers in these fields.

SAME Engineering and Construction Camp participants shown with instructors in Information Technology Laboratory

The Gains in Education of Math and Science (GEMS) Army Educational Outreach Program came to ERDC for the second year, bringing 60 teenagers from Mississippi, Louisiana, Tennessee, and Ohio. This program exposes students in grades 8-12 who have the potential for careers in the areas of math, sciences, and engineering but need motivation and more information to pursue them.

Dr. Jerry Morris, MSRC computer scientist, worked with the students when they came to the Information Technology Laboratory. He conducted a participatory tutorial on empirical and analytical investigation of objects falling under the force of gravity. The students participated enthusiastically as they dropped water balloons for the investigation.

GEMS team participating in experiment with Dr. Jerry Morris

The MSRC also sponsored two students from Mississippi State University (MSU) as they participate in the Scholarship for Service (SFS) program, which is partially administered by the National Security Agency and Department of Homeland Security and funded by the NSF. The SFS promotes the growing cadre of information assurance (IA) professionals in the Federal Government by both requiring instruction at accredited Centers of Excellence and summer internships at Federal entities. As such, Chris Boler, a senior at MSU, and Wes Riley, a graduate student at MSU, both worked for 10 weeks at the ERDC MSRC on various IA-related projects, allowing them to further their education while obtaining real-world experience.

A second West Point cadet IT major, John Kimtis, joined the MSRC in the field of IA. During his 4-week stay at ERDC, he was exposed to various topics that allowed him to see firsthand the types of issues and concerns that are common in IA. He was also given several hands-on projects, allowing him to hone skills not always available in classroom settings.

All three of the above students worked under the guidance of Trent Townsend, the ERDC MSRC Information Assurance Security Officer. Trent has given this kind of guidance to summer students ever since he joined the MSRC as a Government employee after participating in the SFS program from MSU.

"Solving the Hard Problems"

By Rose J. Dykes

Attendees of the 18th annual DoD High Performance Computing Modernization Program (HPCMP) Users Group Conference (UGC) gathered in Seattle, WA, July 14-17, to exchange ideas regarding HPC and see the kinds of progress being made across multiple disciplines in "Solving the Hard Problems," the theme of this year's meeting. In addition to presentations from invited speakers from industry and academia during two half-day plenary sessions, HPC users made 120, 30-minute presentations to their colleagues, attended two full-day and five half-day tutorials on insights into effective HPC techniques and utilization, and presented 57 scientific posters at a well-attended night session.

At the Awards Ceremony, held at the Seattle Museum of Art, HPCMP Director Cray Henry announced that Robert (Bob) Alter, ERDC MSRC team member, won the Hero Award for Technical Excellence in 2008 for distinguishing himself in both management of people and processes at the ERDC MSRC Service Center and in his personal technical contributions to the MSRC and the HPCMP. Henry also said Alter's expert technical assistance to HPCMP users and his understanding of the Program's high performance computers are unparalled.

For its Center poster at the Poster Session of the Conference, the ERDC MSRC chose to highlight programwide initiatives that it performs for the entire HPCMP. These include Benchmarking, the Unclassified Data Analysis and Assessment Center, ezHPC, and the User Interface Toolkit.

Bob Alter selected for HPCMP Hero Award for Technical Excellence for 2008

at UGC 2008 in Seattle

ERDC Information Technology Laboratory team members presented the following:

Data Analysis and Remote Visualization Techniques —Paul Adams

iME (integrated Modeling Environment): A Graphical Tool for Launching and Managing Jobs on HPC Platforms—Owen Eslinger

ezHPC Version 2.0: Simplifying Access to HPCMP Resources—Scotty Swillie

A New Approach to Streambed Modeling and Simulation Using CFD—Jeffrey Allen

Software Framework of Numerical Fish Surrogate Using Unstructured CFD Data —Ruth Cheng

Effects of Loop Blocking for Cache Reuse on AMD Opteron Systems —Tyler Simon

Testing Computational Algorithms for Unsaturated Flow Using the Adaptive Hydrology (ADH) Program—Fred Tracy

An Ontology-Based Warehouse for Scientific Modelers—Amanda Hines

New Features Coming to the OKC—Renee Mullinax, Abdul Mohamed, and John Mason

Extending Automated Meshing to Large Multi-Block Domains: Countermine Mesh Generation on HPC Platforms—Owen Eslinger and Amanda Hines

LAMMPS Simulations of the Effects of Chirality and Diameter on Pullout Force in a Carbon Nanotube Bundle—Dustin Majure

SceneGen: Defining Terrain-based Scenes for Engineering Analysis—Barry White

HPCMP Sustained Systems Performance Test—Paul Bennett

David Stinson, ERDC MSRC Acting Director, shown with Center poster

Consolidated Customer Assistance Center (CCAC) Takes the Next Step—Applying the Enhanced User Experience

By Michelle McDaniel, Tracey Smith, and Casey Bretti, CCAC

The HPCMP Consolidated Customer Assistance Center (CCAC) has taken the next step—the new support contract is applying the Enhanced User Experience (EUE), its innovation to transform CCAC into a more unified, consistent activity across all four MSRCs and two of the Allocated Distributed Centers (ADCs). CCAC began in June 2007 as a consolidation effort by the HPCMP. Its purpose was to provide technical service and account center support to the four MSRCs and two ADCs.

Prior to CCAC, six help desks provided support to the six centers, each center individually providing customer service under its own management. These separate help desks merged into one primary assistance center (CCAC), which provided a single point of contact (help@ccac.hpc.mil or 877-222-2039) for all unclassified HPCMP support needs. This united help desk provided many benefits beyond the single point of contact, including a centralized Kerberos accounts center, uniform support procedures, and consistent CSI software support; but there were still areas that could not be addressed because the separate sites all had their own requirements.

Currently, CCAC operates with seven phone systems and six ticketing systems from five vendors. The realities of working with multiple contracts and local site requirements restrict the amount of consolidation that CCAC can employ. With this in mind, CCAC management (with input from all centers served by CCAC) assembled consultants to analyze and discuss CCAC procedures, seeking ways to better serve the HPCMP users.

The new transformation places the operation of all four MSRCs under a single organizational structure that nurtures a cooperative atmosphere among the sites. This alone enables CCAC to operate more efficiently; but a more comprehensive review was still needed.

To that end, a facilitated process review was held at the AFRL MSRC during the week of 25–29 August. The team conducting the review included representatives from all six centers served by CCAC, CCAC help desk personnel, and a representative from the HPCMP user community. This team reviewed the results of the past study, analyzed the current process, and recommended 22 separate actions that would take CCAC into the next phase of its development. Most of these actions will be behind the scenes – but the result will be a more streamlined experience to those requiring assistance from CCAC.

The changes in CCAC are already taking place, but the more visible aspects of the transformation should start appearing in early 2009 as the improved Web site and self-help capabilities are deployed.

MSRC Web 2.0

Coming To a Browser Near You

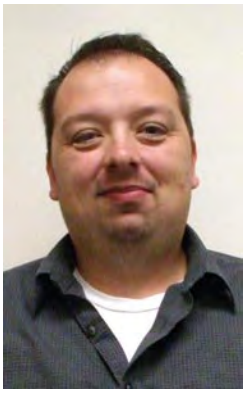

Greg Brewer

Amid the multitude of changes since the June 1 support contract transition, there is another change that, despite the enormous challenge it seeks to address, has gone virtually unnoticed. Greg Brewer and Glen Browning, webmasters of the AFRL and ERDC MSRCs, respectively, have been asked to take the reigns of the four MSRC and CCAC public Web sites, and help chart a path to a more positive user experience across all of the sites. With nearly two decades of Web development experience between them, Greg and Glen will work to create a synergy among all five Web sites that will result in greater consistency, quality, and functionality for users of all the sites.

The MSRC Web team goals are as follows:

- Standardize the technologies and methodologies that underlie the Web sites.
- Create a consistent and intuitive system of documentation so users can easily find the information they need when they need it.
- Identify new Web-based documentation, tools, or services to help users run more efficiently at each of the MSRCs.
- Create a common navigational system across all four MSRC Web sites, while maintaining each site's unique identity.

Glen Browning

In the coming months, much of the work will be unseen by most. Planning is currently underway, with an emphasis on first standardizing the code of each site. While this won't change the current look and feel of the Web sites, it will significantly ease the process of future development for Greg and Glen.

The first changes that users will see will be in the form and content of User Guides, Tutorials, FAQs, and Quick Reference Guides. Once the documentation has been standardized, emphasis will then be placed on providing a common navigational system, with plans of presenting a prototype (maybe more than one) at UGC 09.

Greg and Glen are looking forward to working together and developing the MSRC Web sites into an even better user experience. If you have any thoughts on what you think should be improved concerning all four sites, please head to *www.ccac.hpc.mil/web2.php* and fill out the questionnaire. Your input definitely matters.

Charles Ray (far right) and the Forward Engineering Support Team – Main (FEST-M), July 31, 2008

Daniel Dunmire (second from right) Director, DoD Corrosion Policy and Oversight, Office Secretary of Defense, Washington, D.C., along with others from his office, talks with Dr. Robert Maier (third from left), MSRC Assistant Director, July 29, 2008

CSM Michael L. Buxbaum (left), Command Sergeant Major, USACE, Washington, D.C., and Dr. Maier, July 11, 2008

Paul Adams (left), DAAC Visualization Team Lead, and Gabrielle Puz, The World Bank, Washington, D.C., July 2

Members of 11th Engineer Battalion, Fort Benning, Georgia, and Bobby Hunter (far right), MSRC Acting Assistant Director of Operations, June 26, 2008

Paul Adams (far left) and students with Metadata Internship Program, Jackson State University, June 25, 2008

Bobby Hunter (left) and Emerging Leaders, Mississippi Valley Division, USACE, June 12, 2008

acronyms

Below is a list of acronyms commonly used among the DoD HPC community. These acronyms are used throughout the articles in this newsletter.

AFRL	Air force Research Laboratory
API	ADCIRC Advanced Circulation
ADH	Adaptive Hydrology
AFRL	Air Force Research Laboratory
BLAS	Basic Linear Algebra Subprograms
CAP	Capabilities Application Project
CCAC	Consolidated Customer Assistance Center
CERL	Construction Engineering Research Laboratory
CFD	Computational Fluid Dynamics
CNT	Carbon Nanotube
CPU	Central Processing Unit
CRADA	Cooperative Research and Development Agreement
CRREL	Cold Regions Research and Engineering Laboratory
DAAC	Data Analysis and Assessment Center
DoD	Department of Defense
ERDC	Engineer Research and Development Center
FETI	Finite Element Tearing and Interconnecting
GEMS	Gains in Education of Math and Science
GSL	Geotechnical and Structures Laboratory
HCP	Hexagonal Closet Packed
HPC	High Performance Computing
HPCMP	HPC Modernization Program
HS	Energy-Based Significant Wave Height
HYCOM	Hybrid Coordinate Ocean Model
IA	Information Assurance
IT	Information Technology
ITL	Information Technology Laboratory
JSU	Jackson State University
LAMMPS	Large-scale Atomic/Molecular Massively Parallel Simulator
MPI	Message Passing Interface
MSRC	Major Shared Resource Center
MSU	Mississippi State University
NAVO	Naval Oceanographic Office
NCEP	National Centers for Environmental Prediction
NDBC	National Data Buoy Center
NOAA	National Oceanic and Atmospheric Administration
NRL	Naval Research Laboratory
NSF	National Science Foundation
OKC	Online Knowledge Center
OWI	Oceanweather, Inc.
PET	User Productivity Enhancement and Technology Transfer
RMS	Root Mean Square
SAIC	Science Applications International Corporation
SAME	Society of American Military Engineers
SFS	Scholarship for Service
SHEM	Shared MEMory
SI	Scatter Index
SS	Skill Score
TB	Terabyte
TM	Mean Wave Period
TP	Peak Wave Period
UGC	Users Group Conference
USACE	U.S. Army Corps of Engineers
VBIED	Vehicle-Borne Improvised Explosive Device
WaveMEDS	Wave Model Evaluation Diagnostics System
WIS	Wave Information Studies
WS	Wind Speed

training schedule

For the latest on training and on-line registration, one can go to the User Productivity Enhancement and Technology Transfer (PET) Online Knowledge Center Web site:

https://okc.erdc.hpc.mil

Questions and comments may be directed to PET at (601) 634-3131, (601) 634-4024, or *PET-Training@erdc.usace.army.mil*

ERDC MSRC *Resource* Editorial Staff

Chief Editor/Technology Transfer Specialist
Rose J. Dykes

Visual Information Specialist
Betty Watson

ERDC MSRC HPC Service Center
Web site: *www.erdc.hpc.mil*
E-mail: *msrchelp@erdc.hpc.mil*
Telephone: 1-800-500-4722

Design and layout provided by the ACE-IT, Headquarters, U.S. Army Corps of Engineers.

Approved for public release; distribution is unlimited.

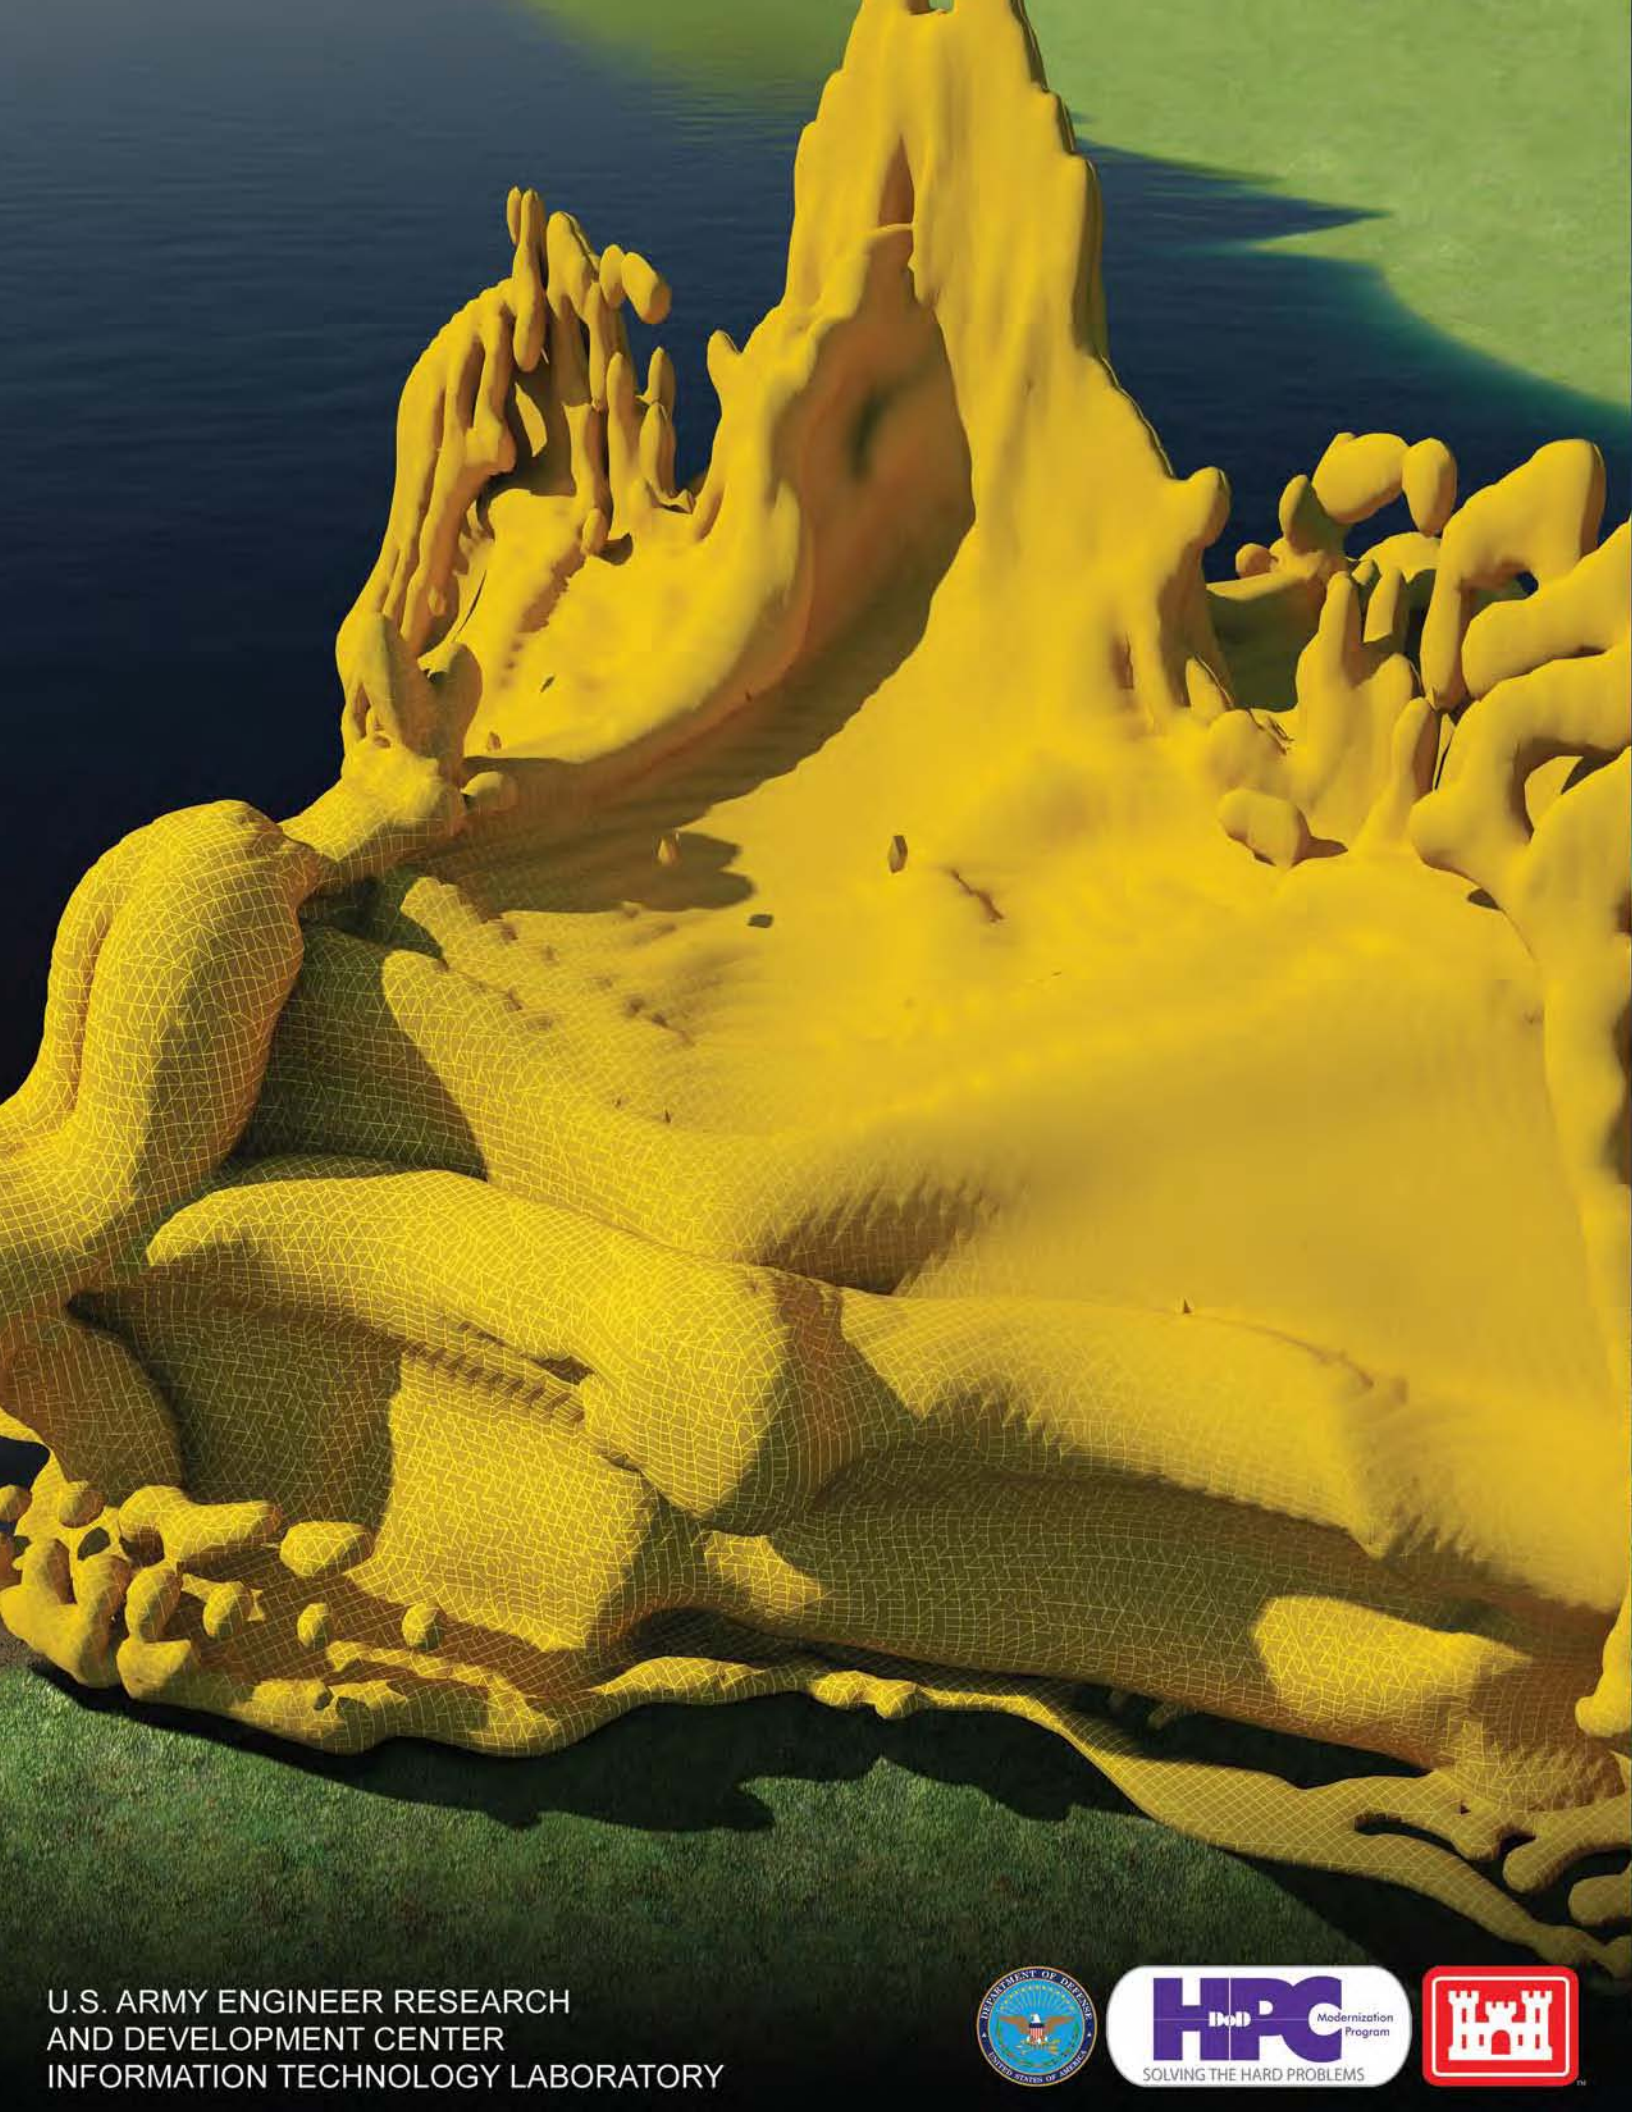

U.S. ARMY ENGINEER RESEARCH
AND DEVELOPMENT CENTER
INFORMATION TECHNOLOGY LABORATORY

DoD HPC Modernization Program
SOLVING THE HARD PROBLEMS