

NEWSLETTER

The

Fall 2003

Resource

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

ERDC **IVISRC**

Major Shared Resource Center





From the Director's chair

The past several months have been – as always – busy ones for the Engineer Research and Development Center Major Shared Resource Center (ERDC MSRC) and our users. In addition to bringing an additional 2.2 TFLOPS to the Center as part of TI-03, we've also added more than 250 TB of storage capacity to keep pace with the expanding demand for archival services. Those that know me, however, know that I believe deeply that the only thing that makes a center like this one worth

keeping open is the degree to which our resources can be brought to bear on the mission of the agencies and services we serve. At the ERDC MSRC, the key to making that transfer possible is our technical staff.

Professionals in a field like high-performance computing (HPC) are lifelong learners, partly because they have to be, but also because niche fields like HPC attract a certain kind of person. Architectures, paradigms, and infrastructures for computing change continuously, and so therefore do we. Change is really about learning; in this issue, along with the usual features and tips to help you get the most out of your HPC center, you'll find several articles that highlight how the ERDC MSRC is involved with education, teaching, and training programs at a variety of levels.

The focus of this issue is not on the steps we take to make sure that we are able to stay abreast of the technologies needed to help you serve your mission, however. That's part of what we do and who we are, and you expect that from us. These articles focus on another part of our mission – external outreach and education in the HPC and science and technology communities. Last year MSRC team members devoted 6,000 hours to presenting papers, tutorials, and seminars at professional gatherings; lecturing to high school and college students; teaching graduate-level courses at accredited universities; serving as mentors in summer and cooperative education programs; volunteering as science fair judges; chairing technical working groups, and so on.

The portion of our education and outreach directed at the peer level is an expected part of our mission to provide service – educating users and disseminating the results of our efforts to press HPC into the service of the defense mission is central to what we do. We believe strongly, however, that we have a broader mission to serve as well. If our Nation is to continue to maintain – and in some cases regain – its technological leadership and turn that to tactical and strategic advantage, then it must attract our best and brightest students to science and engineering and keep them there. By exposing students to the exciting things we do for the country with the technology we provide through the High Performance Computing Modernization Program (HPCMP), we feel we are playing a small part in that mission.

And sometimes, the interest we are serving turns out to be our own. This issue features two articles by full-time staffers who are graduates of this MSRC's student programs: the article by Rikk Anderson, who is currently heading up our measurement and monitoring program, and the article you are reading now. I spent 2 years as a student in the MSRC. Those years shaped my decision to go on to graduate school and make HPC my career.

John West
Director, ERDC MSRC

About the Cover:

Large-Eddy Simulation..... (see article, page 2). Cover design by the ERDC MSRC Scientific Visualization Center.

Features

Simulation of Steep Breaking Waves and Spray Sheets Around a Ship: The Last Frontier in Computational Ship Hydrodynamics	2
From Altair to X1...It's the experience that makes the difference!	12
SAME Engineering and Construction Camp a Learning Experience	14
JSU Summer Institute Fosters Interest in HPC Careers	15
National Science Foundation Scholarship for Service Cyber Corps Student Interns at ERDC MSRC	16
Students Succeed in 2003 HPCMP PET Summer Intern Program	17
Watch Out for "Bites" When Writing About Bytes	18
Monitoring Produces Insight About Services for the Future	20
Almost "Sleepless in Seattle"—Users Group Conference 2003	21
PET Training Moves to New Facility	24

Departments

announcements	8
HPCMPO Presents Award to ERDC MSRC Team Member	8
ERDC MSRC High Performance Computers on Top 500 Supercomputers List	8
Cray X1 Added to ERDC MSRC Configuration	9
off-campus	9
Dr. Ruth Cheng, ERDC MSRC, Makes Presentations at Two International Conferences	9
Benchmarking and Performance Characterization	9
Computational Science Australian Style	10
happenings column	25
ERDC MSRC Hosts Four Groups During Same Week	25
ERDC MSRC Team Members Serves at ERDC Graduate Institute	26
technology update	27
ERDC MSRC Mass Storage Archival System Expansion	27
Software Profiling: Focusing on Interprocessor Communications	27
Using SHMEM for Low-Latency Communication on the Compaq SC40 and SC45	30
upcoming conferences	31
visitors	32
acronyms	36
training schedule	36

Simulation of Steep Breaking Waves and Spray Sheets Around a Ship: The Last Frontier in Computational Ship Hydrodynamics

*By Kelli Hendrikson, Lian Shen, and Dick K.P. Yue, Massachusetts Institute of Technology (MIT);
Dr. Douglas G. Dommermuth, Science Applications International Corporation (SAIC);
and Paul Adams, ERDC MSRC*

The simulation of complex free-surface phenomena such as breaking waves, spray sheets, and air entrainment play a key role in the design and operation of naval combatants. These phenomena are among the most challenging problems in computational fluid dynamics today, and the current capabilities to model them are far from satisfactory. Within the realm of hydrodynamics, these phenomena involve complex processes such as strong free-surface turbulent interactions, jet and spray formation, bubbly flows, and postbreaking turbulent dissipation for which researchers' physical understanding (much less how to model these processes) is insufficient. The study of complex free-surface flows involves macroscale phenomena such as wave-body interactions and wave overturning, and microscale phenomena such as structure of the free-surface boundary layer. The research at SAIC is focused on developing the macroscale solver that incorporates the microscale physics and models that are being developed at MIT. For both the microscale and macroscale flows, physical understanding of the phenomenon is obtained via simulation of canonical problems. These studies are used to guide development of the base-flow solver and statistical models. The resulting numerical capabilities, transport models, and source terms are then incorporated into the analysis capabilities of the Navy for use in the design of naval combatants.

In this project, a suite of unique, scalable, parallel-computing codes is being applied. The computational methods employed are direct numerical simulation (DNS), large-eddy simulations (LES), and boundary-fitted and Cartesian-grid-based methods that are capable of modeling the turbulent, highly mixed flows around ships. The project's ultimate goal is to establish DNS, LES, and modeling capabilities for complex free-surface hydrodynamics and provide the Navy with a complementary suite of analysis capabilities for multiphase computational ship hydrodynamics through large-scale simulations. This will form the basis for the development of the next generation of LES and RANS tools to answer the call for the Navy's need of better design tools for the next generation of ships.

Problem and Methodology

The ShipLES suite of codes developed at MIT includes both DNS and LES capabilities that work for both the

single-phase and multiphase free-surface hydrodynamics. In the multiphase environment, interface capturing is accomplished using a second-order accurate level-set method, which is a Eulerian Interface-Capturing Method (EICM). Both water and air are treated as an incompressible fluid governed by a multifluid Navier-Stokes equation. The interface between the two fluids is smoothed over a few grid points for numerical stability. The pressure is solved for using a variable-coefficient Poisson equation that projects the velocity onto a divergence-free field, which ensures continuity. The Cartesian grid uses a staggered MAC-type formulation to ensure continuity and a compact stencil for the Poisson equation, and thus no upwinding of the convective terms is necessary. This feature is highly desirable for developing turbulence models of the microscale physics. Communication between processors is accomplished via MPI making the code portable to any platform. The CPU requirements are linearly proportional to the number of grid points and inversely proportional to the number of processors.

The Numerical Flow Analysis (NFA) code provides turnkey capabilities to model breaking waves around a ship, including both plunging and spilling breaking waves, the formation of spray, and the entrainment of air. Cartesian-grid methods are used to model the ship hull and the free surface. The equations are solved on a MAC grid. Following Goldstein et al. (1993), Colella et al. (1999), and Sussman and Dommermuth (2001), body-force and finite-volume methods are used to enforce the boundary condition on the hull. The body-force technique uses a source term in the Navier-Stokes equation to impose no-slip conditions on the body surface. The finite-volume technique modifies the Pressure-Poisson equation to account for fractional cells to impose free-slip boundary conditions. A surface representation of the ship hull is used as input to construct the body force and the metrics for the fractional cells. This greatly simplifies the geometry input that is required to simulate three-dimensional (3-D) wave breaking. The interface capturing of the free surface uses a second-order accurate, volume-of-fluid technique. At each time-step, the position of the free surface is reconstructed using piece-wise planar surfaces as outlined in Rider et al. (1994). A second-order, variable-coefficient Poisson equation is used to project the velocity onto a solenoidal field, thereby ensuring mass conservation. A preconditioned

conjugate-gradient method is used to solve the Poisson equation. The preconditioner uses an incomplete Cholesky factorization with overlapping blocks of grid points. Details of a similar projection operator are provided in Puckett et al. (1997). The convective terms in the momentum equations are accounted for using a slope-limited, third-order QUICK scheme as discussed in Leonard (1997). The momentum equations are integrated in time using a second-order Runge-Kutta algorithm. The governing equations are solved using a domain decomposition method based on the PARAMESH algorithm as described by MacNeice et al. (2000). Communication between processors on the Cray T3E is performed using Cray's shared-memory access library (SHMEM). The CPU requirements are linearly proportional to the number of grid points and inversely proportional to the number of processors.

Results

To date, NFA simulations have been performed for a vertical strut, a wedge-like geometry, and a half model of the DDG 51. These numerical simulations are used to guide the development of the base-flow solver. ShipLES simulations have been performed of spilling breaking waves in order to improve researchers' understanding and to develop models of free-surface turbulence.

The DNS results for an impulsively started, overly steep plane progressive wave using ShipLES-EICM are compared with the laboratory experiments of Qiao and Duncan (2001) in Figure 1. The spilling breaking wave generated in the laboratory was generated using a wave-focusing technique, and images were taken over a successive number of runs. The experimental images represent one of these snapshots, while the DNS result is time-averaged over a portion of the breaking event. Even though the numerical results are for a lower Reynolds number ($Re_w \sim O(10^3)$) compared with $Re_w \sim O(10^6)$) and larger boundary layer to amplitude ratio ($\delta/A \sim 0.3$ compared with $\delta/A \sim 0.01$), the numerical and laboratory results agree qualitatively for the free-surface shape and the vorticity in the face of the wave. In particular, being able to predict the vorticity in the free-surface boundary layer is the first step toward developing microscale models of free-surface turbulence as is discussed later.

Figure 2 illustrates a comparison between numerical predictions and experimental measurements for a vertical strut moving with constant forward speed. The numerical predictions are shown on the left side of the figure, and the experimental measurements are shown on the right side of the figure. The Froude number is $Fr = U/(gL)^{1/2} = 0.55$, where U is the speed

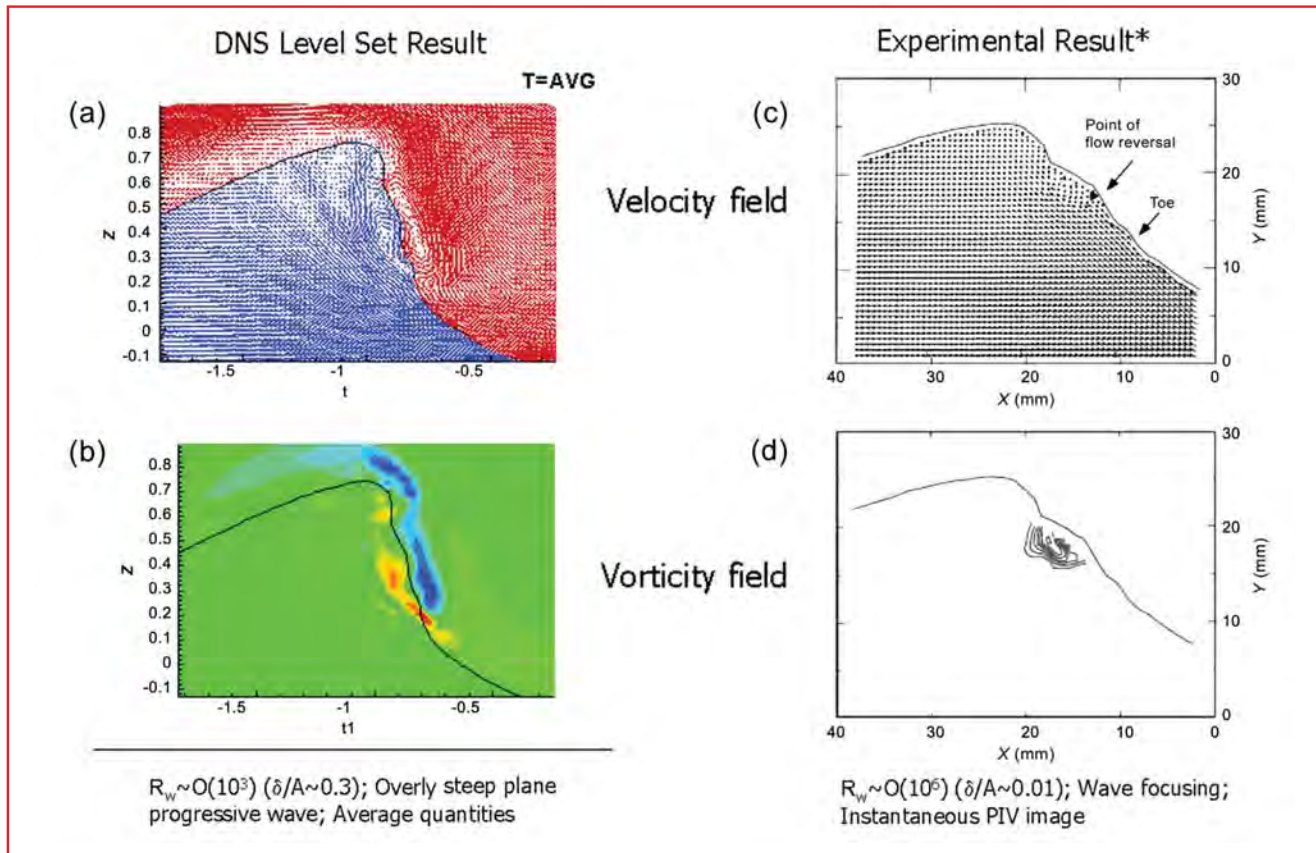


Figure 1. Comparison of ShipLES-EICM results (a-b) with laboratory experiments (c-d). Velocity vectors (a and c) and vorticity contours (b and d) (after Qiao and Duncan 2001)




Figure 2. The wave field around a vertical strut. Red contours denote wave crests; blue contours denote wave troughs

of the strut, g is the acceleration of gravity, and L is the chord length of the strut. The experimental measurements are described in Zhang and Stern (1991). The 3-D numerical simulations used $512 \times 128 \times 128 = 8,388,608$ grid points, 512 processors, and approximately 50 hours of CPU time.

In general, the agreement between the numerical simulations and the experimental measurements is very good. However, there are some notable differences. For example, the numerical simulations show more fine-scale detail than the experimental measurements. This is because the experimental measurements are time-averaged and the numerical simulations show an instantaneous snapshot of the free surface. The authors also note that unlike the numerical simulations, the measuring device that had been used in the experiments was only capable of measuring single-valued free-surface elevations. Another difference between numerical simulations and experimental measurements occurs away from the strut where the numerical simulations show edge effects because of the smaller

domain size that is used relative to the actual experiments. Figure 2 illustrates that the authors are able to model the macroscale features of the flow associated with the body interacting with the free surface.

Figures 3 and 4 show details of the numerical simulations and experimental measurements for two different views. The numerical results are emphasized on the left sides of the figures, and the experimental measurements are emphasized on the right sides of the figures. The red and green dots along the sides of the strut denote experimental measurements of the free-surface profile. Toward the rear of the foil, the dots indicate the upper and lower bounds of the unsteady rise and fall of the free surface. Note that in Figure 3 and to lesser degree in Figure 4, the experimental contours off the body do not appear to agree with the experimental profiles on the body. This is because the contour measurements off the body could not be performed too close to the body because of limitations associated with the measuring device. In Figure 3, the numerical simulations show the formation of a spilling breaker and spray near the leading edge of the strut. Toward the rear of the strut, flow separation is evident. The numerically predicted free-surface elevations agree well with the profile measurements in both Figures 3 and 4. The numerical simulations in Figure 4 illustrate that air is entrained along the sides of the strut and in the flow separation zone in the rear.

The numerical simulations of the strut have been animated using novel techniques generated at the ERDC MSRC Scientific Visualization Center (SVC) that are described in Adams and Dommermuth (2003). The animations illustrate physical phenomena that are not possible to depict in an article. For example, the animations show a necklace vortex is formed along the sides of the foil. The rotating flow is observed in the animations because of the entrained air acting like a




Figure 3. Side views looking down on the strut




Figure 4. Side views looking up on the strut

passive tracer. In addition, jets along the side of the foil tend to inject air bubbles deep into the fluid, which is a phenomenon that had been conjectured but never observed in a numerical simulation. The animations are viewable at <http://www.erd.c.hpc.mil/>.

The instantaneous vortical structure inside a spilling breaking wave (SBW) is shown in Figure 5. At this point in the breaking event, the toe of the breaking wave has moved down the forward face of the wave crest and resides at the lower circled flux point. There is a significant influx of positive vorticity into the water at the “kinks” in the surface, and the front face of the wave is, in general, a significant source of negative vorticity for the air. Analysis reveals two main sources of vorticity: (a) the work done by the surface tension forces peaks at the same points of intense vorticity flux, particularly at the bulge (top circled flux point); and (b) surface-parallel velocity undergoes sharp changes in curvature that is concentrated at the upper and lower “kinks” in the surface. Here, the small details associated with the flow in the air are interesting because they affect the transport of droplets.

Figure 6 illustrates the free-surface elevation around a wedge like geometry moving with constant forward speed. The Froude number is $Fr=0.302$. Numerical simulations are shown on the left side of the figure, and experimental measurements are shown in the small region on the right side of the figure. The experiments are described in Karion et al. (2003). Once again, spilling breaking waves and spray formation are observed near the bow. At the corner where two facets of the hull geometry intersect, there is massive flow separation. The agreement between numerical simulations and experimental measurements is good.

The transfer of energy at the surface between two immiscible fluids can be derived from the multifluid

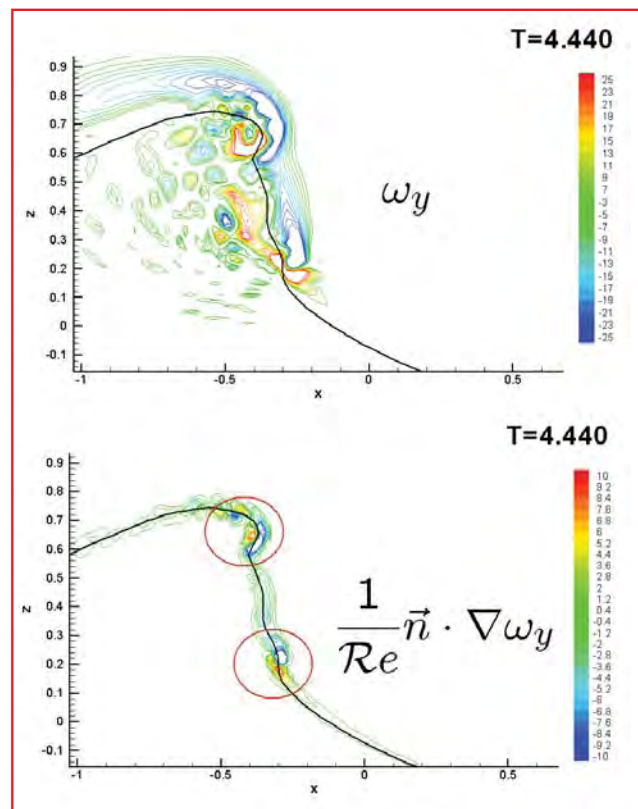


Figure 5. Instantaneous vorticity contours (top) of an SBW and vorticity flux (bottom) simulated with SHIPLES-EICM. Positive vorticity flux represents inflow of positive (clockwise) vorticity into water

energy equation. This term consists of inviscid and viscous energy transport and is shown in Figure 7. For this SBW, the transport term is dominated by the inviscid energy transport. This is consistent with the sources of vorticity highlighted previously in that the surface tension forces (pressure) dominate the work done that generates vorticity. Also note that, while the viscous energy transport is small, there is a concentration of this term where the toe has moved down the




Figure 6. The free-surface elevation around a wedge-like geometry

wave face. Analysis of the average and RMS quantities of the total energy transport term reveals that the peak RMS values occur in the regions of vorticity influx and are of the same order as the average. This latter point is important for modeling considerations.

Considerable effort was invested this year in developing LES formulations for mixed-phase flows and of the level set governing equation, analyzing two types of filtering approaches, and *a-priori* evaluation of local and global models for SGS terms. In general it was found that, for nondensity weighted LES filters, scale-similar global SGS models performed significantly better than local eddy-diffusivity type models. This can be seen in Figure 8. It is clear that the mixed model, which is a scale-similarity model combined with a

Smagorinsky model, performs significantly better with a correlation coefficient of 0.88. Now the goal is to incorporate this SGS model into RANS and the macroscale solver.

Figure 9 illustrates the free-surface elevation around a DDG 51. The Froude number is $Fr=0.276$. The results of preliminary numerical simulations are shown on the left side of the figure, and experimental measurements are shown on the right side of the figure. The experiments have been performed at the David Taylor Model Basin. For the numerical simulations, only the forward half of the geometry has been modeled. The numerical simulations have a numerical artifact near the leading edge of the domain where the solution is periodically wrapped. This problem has since been remedied using true inflow boundary conditions. Overall, the agreement between numerical simulations and experimental measurements is satisfactory.

Conclusions

Numerical methods have been developed to simulate the macroscale generation of breaking waves near a ship and the microscale features within the breaking wave itself. The numerical methods are the first step toward developing high-fidelity models of breaking waves. These capabilities are required to help the Navy meet stringent performance requirements for the next generation of ships. Future design concepts such as the DDX and LCS include features such as tumblehome bows, shallow transom sterns, planing hulls, and hydrofoils that are difficult to model with conventional

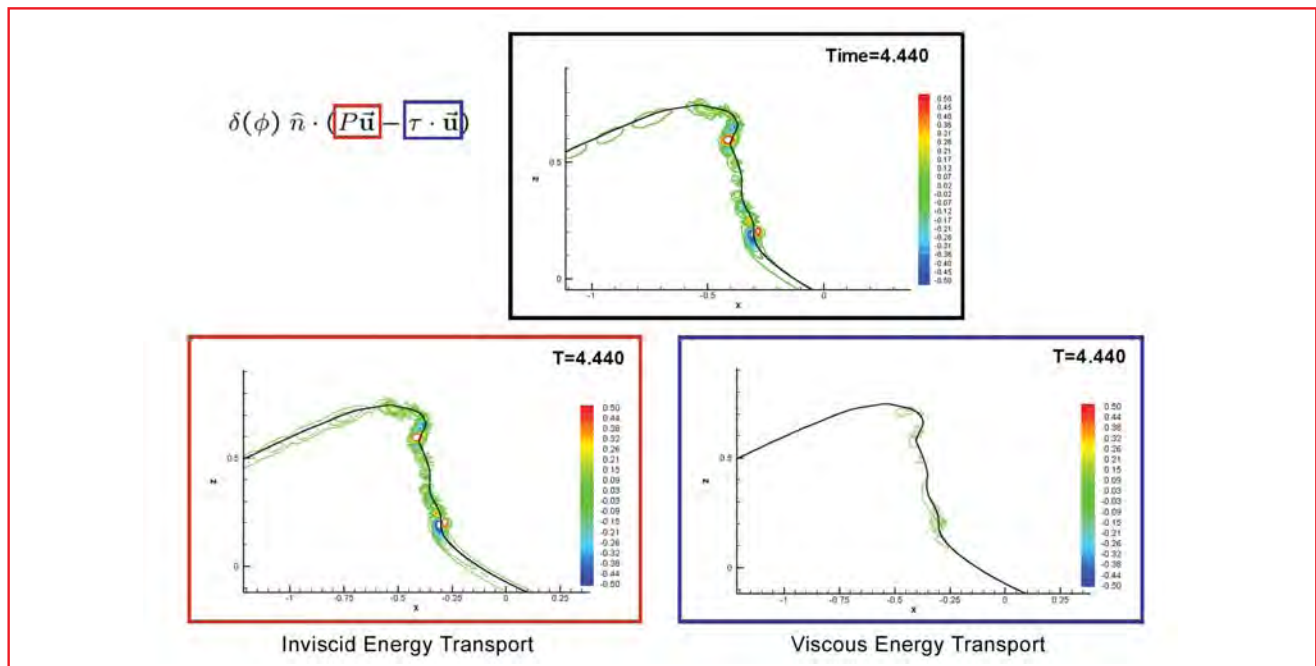


Figure 7. Transfer of energy at the free surface of SBW simulated with ShipLES-EICM. Total energy transfer at the interface is equal to the sum of the inviscid and viscous energy transport

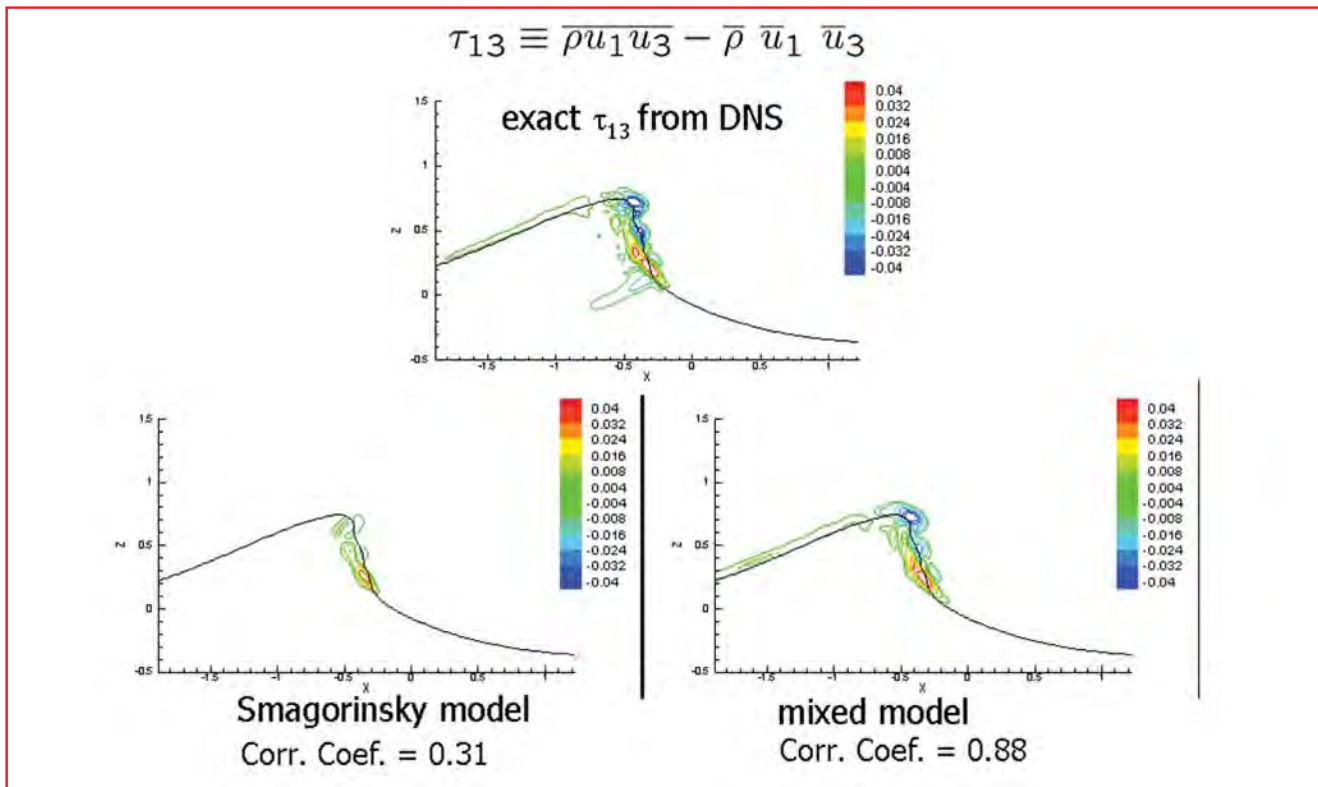


Figure 8. Comparison of SGS shear stress between exact DNS values and a local (Smagorinsky) model and a global, scale-similarity-based mixed model




Figure 9. The free-surface elevation around a half model of the DDG 51

approaches. The use of Cartesian-grid methods and interface-capturing methods as illustrated in this article will provide the necessary capabilities to the Navy to model unusual hull forms and complex physics such as breaking waves, the formation of spray, and the entrainment of air.

Acknowledgments

This research is sponsored by the Office of Naval Research. The contract numbers are N00014-02-C-0017 and N00014-01-1-0124. The research program is managed by Dr. Patrick Purtell (purtelp@onr.navy.mil). The DoD Challenge project is managed by Dr. Ki-Han Kim (kimk@onr.navy.mil).

References

- Adams, P. and Dommermuth, D. (2003) Visualization of Steep Breaking Waves And Thin Spray Sheets Around A Ship, *Submitted for publication in VIS 2003, Seattle, WA*.
- Colella, P., Graves, D.T., Modiano, D., Puckett, E.G., and Sussman, M. (1999) An embedded boundary/volume of fluid method for free-surface flows in irregular geometries. *In Proc. FEDSM99 3rd ASME/JSME Joint Fluids Engineering Conference, San Francisco, CA*.
- Goldstein, D., Handler, R., and Sirovich, L. (1993) Modeling a no-slip boundary with an external force field. *J. Comp. Phys.*, 105, 354-366.
- Karion, A., Sur, T., Fu, T.C., Furey, D., Rice, J.R., and Walker, D. (2003) Experimental study of the bow wave of a large towed wedge. *Submitted for publication in the 8th International Conference Numerical Ship Hydrodynamics*.
- Leonard, B.P. (1997) Bounded higher-order upwind multidimensional finite-volume convection-diffusion algorithms, *Advances in Numerical Heat Transfer*, edited by W.J. Minkowycz, W.J. and E.M. Sparrow, Taylor & Francis, Washington, D.C., 1-57.
- MacNeice, P., Olson, K.M., Mobarry, C., deFainchtein, R., and Packer, C. (2000) PARAMESH: A parallel adaptive mesh refinement community toolkit, *Computer Physics Communications*, 126, 330-354.
- Qiao, H. and Duncan, J.H. (2001) "Gentle Spilling Breakers: Crest Flow-Field Evolution," *J. Fluid Mech.*, 439, 57-85.
- Puckett, E.G., Almgren, A.S., Bell, J.B., Marcus, D.L., Rider, W.J. (1997) A second-order projection method for tracking fluid interfaces in variable density incompressible flows. *J. Comp. Physics*, 130, 269-282.
- Rider, W.J., Kothe, D.B., Mosso, S.J., Cerutti, J.H., and Hochstein, J.I. (1994) Accurate solution algorithms for incompressible multiphase flows. *AIAA paper 95-0699*.
- Sussman, M. and Dommermuth, D.G. (2001) The numerical simulation of ship waves using Cartesian-grid methods. *In Proc. Of the 23rd Symp. on Naval Hydro., Nantes, France*, 762-779.
- Zhang, Z. and Stern, F., (1996) Free-Surface Wave-Induced Separation. *ASME J. Fluids Eng.*, 118, 546-554.

HPCMPO Presents Award to ERDC MSRC Team Member



The High Performance Computing Modernization Program Office (HPCMPO) presented a trophy to Dr. Bill Ward, leader of the Computational Science and Engineering (CS&E) group, ERDC MSRC, at the 2003 Users Group Conference conducted in Seattle, Washington, June 9-13. The award was in recognition of his significant contributions to the success of the ERDC MSRC and the entire HPCMP. The CS&E group has ported, tuned, and optimized codes for users, and the group’s benchmarking and performance characterization activities have underpinned all major HPCMPO hardware acquisitions since 2001. Consequently, the CS&E group’s efforts have touched each MSRC and every user in a significant way. Dr. Ward is a public advocate for the art and practice of HPC and is widely recognized in the HPC technical community as an accomplished master of his craft.

Dr. Bill Ward, Lead for Computational Science and Engineering, ERDC MSRC

ERDC MSRC High Performance Computers on Top 500 Supercomputers List

All HPC platforms at the ERDC MSRC have been ranked among the Top 500 supercomputers in the world.

The latest Top 500 list was released in June 2003 by the University of Mannheim, the University of Tennessee, and the National Energy Research Scientific Computing Center/Lawrence Berkeley National Laboratory. Compiled twice a year since 1993 and based on the LINPACK benchmark, the list provides a statistical component in the overall picture of comparative supercomputing performance.



The highest-ranking ERDC MSRC supercomputer on the list is the Cray T3E at number 46. Rounding out the list from the ERDC MSRC are the Compaq SC45 at number 75, the Compaq SC40 at number 106, and components of the SGI Origin 3000 (O3K) Complex at positions 155, 156, and 384. The Earth Simulator Center in Japan is ranked number 1.

With 1,904 processing elements, the ERDC MSRC’s Cray T3E is the largest unclassified Cray T3E supercomputer in the world. The Compaq SC40 and SC45 each have 512 processors, as do the two SGI Origin 3900s and the SGI Origin 3800 that make up the ERDC MSRC’s O3K Complex. With the addition of the 64-processor Cray X1, the ERDC MSRC now has 6.4 teraFLOPS of computing capability.

The aggregate capability of the ERDC MSRC, which serves the HPC needs of thousands of engineers and scientists throughout the Department of Defense (DoD), continues to place the ERDC MSRC among the top 20 percent of supercomputing centers in the world.

To view the Top 500 list in its entirety, go to <http://www.top500.org/>.

The ERDC MSRC welcomes comments and suggestions regarding *The Resource* and invites article submissions. Please send submissions to the following e-mail address:
msrchelp@erdc.hpc.mil

Cray X1 Added to ERDC MSRC Configuration

With the delivery and installation of a liquid-cooled 64-processor Cray X1 this summer, the ERDC MSRC now has what Cray Inc. bills as “the world’s most powerful supercomputer product.”

The system, named diamond, provides 800 gigaFLOPS (peak) of computational power to ERDC MSRC users through 64 800-MHz CPUs, each containing four multistreaming processors (MSPs) and 4 GB of RAM, for a total of 256 GB of memory. The system is configured with 8 TB (raw) of Fibre Channel disk storage for home directories and user workspace.

The Cray X1 was acquired during the Technology Insertion 2003 (TI-03) procurement process.

The Cray X1, a multistreaming vector-computing supercomputer, creates a distinctly new parallel

programming environment for ERDC MSRC users. “This new architecture allows the ERDC MSRC to continue its tradition of innovation and technology leadership,” Center Director John E. West said.

With the addition of the Cray X1, the ERDC MSRC has an aggregate computing capability of 6.4 teraFLOPS. The Center’s other HPC platforms include the Cray T3E with 1,904 processing elements; the 512-processor Compaq SC40 and SC45; and the two SGI Origin 3900s and the SGI Origin 3800, each with 512 processors, that make up the ERDC MSRC’s O3K Complex. The two Origin 3900s, named silicon and sand, were also TI-03 hardware installations.

The Cray X1 was delivered on June 30, 2003, and entered production use in October.

off-campus

Dr. Ruth Cheng, ERDC MSRC, Makes Presentations at Two International Conferences

Dr. Ruth Cheng presented a paper entitled “Implementation of a Parallel Particle Tracking Algorithm in the FEMWATER Chemical Transport Code” on June 23 at the 2003 World Water and Environmental Resources Congress in Philadelphia, Pennsylvania. This paper demonstrated the successful incorporation of parallel particle tracking software into the parallel FEMWATER code. The performance was evaluated on the ERDC MSRC Compaq SC45 computer. Dr. Cheng also attended the 2003 International Conference on Parallel and Distributed Processing Techniques and Applications on June 25 in Las Vegas, Nevada, and presented a paper entitled “A Software Architecture for Parallel Particle Tracking Algorithms.” This presentation described the architecture of new parallel particle tracking software and how its design goal of interfacing legacy software libraries and application codes can be achieved. The software is available for users on the MSRC HPC systems at /usr/local/usp/pt.

Benchmarking and Performance Characterization

Dr. Bill Ward, ERDC MSRC, gave a presentation on the DoD HPCMO’s recent efforts in the areas of benchmarking and instrumentation of scientific application codes at the Workshop on Performance Characterization, Modeling and Benchmarking for HPC Systems, Emeryville, California, May 5-7. The workshop’s motivating issue was the growing gap between theoretical peak performances and observed sustained performance of HPC systems. Representatives from academia, DoD, DOE, and HPC hardware vendors discussed performance characterization, modeling, and benchmarking of both HPC hardware and software that are necessary to understand this problem and to develop remedies. Significant attendees included Cray Henry, Director, HPCMP, DoD; Adolffy Hoisie, Leader of the Parallel Architecture and Performance Team, Modeling, Algorithms and Informatics Group, Computer and Computational Science Division, Los Alamos National Laboratory; John McCalpin, IBM (author of STREAM benchmark); Horst Simon, Director of the National Energy Research Scientific Computing (NERSC) Center Division of the Lawrence Berkeley National Laboratory (LBNL), DOE; Burton Smith, Chief Technical Officer, Cray, Inc.; and Jeff Vetter, Lawrence Livermore National Laboratory, DOE.



1



14

Dr. Fred Tracy works as a research computer scientist in the Computational Science and Engineering group of the ERDC MSRC. He recently attended the International Conference in Computational Science (ICCS 2003) June 2-4 in Melbourne, Australia, and the article below “journals” his once-in-a-lifetime experience.

Computational Science Australian Style

I have never fretted about a trip more than this one, but despite threats of terrorists and the SARS virus, I found myself presenting the paper “Application of the Multi-Level Parallelism (MLP) Software to a Finite Element Groundwater Program Using Iterative Solvers with Comparison to MPR” (Figure 1) at the Grand Hyatt Hotel (Photo 1) in Melbourne, Australia. The purpose of my paper was to give the results of the performance evaluation of MLP versus MPI using the groundwater program FEMWATER where iterative solvers are employed to solve the system of linear equations generated from the finite element formulation of the equations. The test problem used was the remediation of a military site. The major findings were that the MLP times were generally better than the MPI results for Processing Elements (PEs) 8-64, and the MPI times were as good or better when 128 PEs were used. The percentage differences were always less than 10 percent, which means that no significant advantage was achieved with either parallel paradigm for this application. Approximately 300 people attended the conference where approximately 200 papers were given.

The first keynote speaker was Professor Steven Praver, School of Physics, University of Melbourne, and he gave the talk “Quantum Computing: What’s All the Fuss About?” This was especially interesting and timely because of the emphasis that Dr. John A. Parmentola, Director for Research and Laboratory Management, Office of the Assistant Secretary of the Army for Acquisition, Logistics and Technology, Washington, DC, placed on the importance of basic research in quantum computing for the U.S. Army when he recently briefed ERDC research scientists and engineers. This importance is because quantum computing can potentially dramatically increase the ability to work with coded messages (cryptology). A normal computer with say three bits can compute with numbers between 0 and 7. The bits are either on or off. In a quantum computer with three molecules, however, the spin of these molecules is determined by quantum mechanics. Thus, you only know whether the spin is “up” or “down” when you measure it. Fifty percent of the time you will get “up,” and fifty percent of the time you will get “down”; but, like flipping a coin, you do not know its outcome in advance. This concept, when applied to quantum computing, results in the ability to represent in our three-molecule model the numbers 0-7 simultaneously! Thus, many scenarios can be considered simultaneously. If built, the quantum computer certainly holds exciting new frontiers in computational science.

The inhabitants of Australia were very friendly (Photo 2); however, when I approached some of them with the subject of computational science, they soon lost interest (Photo 3). The blue ribbon for “reaching out” (Photo 4), therefore, goes to my wife, Barbara, who does ocean modeling in the ERDC Coastal and Hydraulics Laboratory.

Actually, I made some excellent contacts at the conference. Of special help was Dr. Mark Kremenetsky (Figure 2), Principal Scientist, SGI, who is located at Mountain View, California. He not only had significant knowledge about my talk but also presented an enlightening talk on using OpenMP that helped me in my MLP study.

Melbourne is a European/Canadian style city (Photos 5 and 6) with wide streets, beautiful parks and gardens (Photo 7), and excellent cultural events such as those at the Victorian Arts Center. Within driving distance of Melbourne is a remarkable diversity of landscape. Examples include flat country good for sheep (Photo 8), Philip Island (Photo 9), where the penguins come up out of the water and waddle to their nests after sunset, the Grampion Mountains (Photos 10 and 11), marshland (Photo 12), jungle (Photo 13), the Great Ocean Road (Photos 14 and 15), waterfalls (Photo 16), beautiful mountain forests (Photos 17 and 18), the sometimes mighty Murray River (Photo 19), and beautiful wineries (Photo 20).

Besides the kangaroos, other interesting wildlife that I was able to see consisted of koalas (Photo 21), snakes with head and tail the same (Photo 22), beautiful birds (Photo 23), and wombats (Photo 24).



2



3



4



5



6






Figure 1. Dr. Tracy speaking at ICCS 2003



Figure 2. Dr. Mark Kremenetsky speaking at the conference



From Altair to X1 ...

It's the experience that makes the difference!


By John D. Mauldin

After only 2 months of facility improvements and system tests, the Cray X1 is ready for service. From the outside, it might appear that preparations for this milestone have been confined to the last few months. However, in addition to the spate of recent activity, one might suggest that staff members at the ERDC MSRC have spent the entirety of their professional careers preparing for this exciting event!

While the X1 is new and innovative, effective utilization of this resource will depend on a broad and solid foundation of experience in supercomputing and information technology (IT). With knowledge gained over careers that span the spectrum of IT from the General Electric (GE) ERMA in the late 1950s to the Altair personal computer in the 1960s to the X1 in 2003, the ERDC MSRC's team of engineers, scientists, technicians, administrators, and managers uses the experiences of a lifetime to solve the challenging problems that face today's researchers.

This experience is as much a hallmark of the ERDC MSRC as is our state-of-the-art technology, and this experience manifests itself in the enthusiasm and competence with which our staff serves our users. On a given day, ERDC MSRC team members are fully engaged supporting users: solving problems, porting code, upgrading machines, tuning operating systems, monitoring job queues, testing system performance, analyzing data with state-of-the-art visualization and data discovery tools, and analyzing future technology trends and needs in preparation for the next advance in high-performance computing (HPC) capability. In addition to this routine of outstanding user support, MSRC team members are also spreading the word about the challenge and reward of technical careers and building enthusiasm in the next generation of HPC professionals.

Last year, MSRC team members approached 6,000 hours in professional education, community outreach, and related activities. This year, team members have accumulated nearly 3,000 hours of outreach service and expect to reach and exceed last year's impressive total. Team members write technical papers, chair technical panels, host tours of the MSRC, and share their knowledge with colleagues through brown bag lunch presentations. During the past 2 years, team



John D. Mauldin
Education Portfolio Manager
ERDC MSRC

members have presented more than 50 technical papers in national and international forums including conferences in Australia and Scotland. Tours of the MSRC have been hosted for numerous groups and individuals: from congressional staff members to representatives of the Nigerian Inland Waterways Authority; from high school students to university professors; and from military officials to the mayor of Carlsbad, California. And in their spare time, MSRC team members visit high schools, teach classes at local community colleges, mentor college students, and sponsor local robotics teams.

So why do they do it? The simple answer is that for most of us here at the ERDC MSRC, a career in HPC is more than a job. It is a great place to be for those with a lifelong curiosity, who thrill at the mention of questions that begin with "how" and "why" and "what if." ERDC MSRC team members add daily to more than 1,100 years of IT experience. And they join this experience with more than 575 years of HPC experience to work together with the Nation's top researchers and scientists, solving some of the most challenging technical questions facing our Nation. From experience with GE's ERMA in the 1950s to the X1 in 2003, from the first personal computer, the Altair, to today's powerful desktop workstations, ERDC MSRC team members have also added to their resumes with accomplishments and technical breakthroughs. They have added to the foundation of knowledge that now prepares them for supporting one of the Nation's most advanced computing centers.

From the world's first personal computer to the world's most powerful supercomputer, ERDC MSRC team members continue their journey – using lessons learned yesterday to solve today's problems and preparing today for the challenges of tomorrow.

From the File of Esoteric Computer Stories

Part of a group that developed a computer-based data acquisition and analysis tool for applying the principles of pneumatics and fluid flow to determine the location of constrictions and blockages in the male urinary tract.

— Submitted by: Bob Tipton, ERDC Cray
and Mass Storage Staff

ERDC Staff memories and experiences

1959 – General Electric produces the GE ERMA to process checks in a banking application via magnetic ink character recognition.

— Bill Bennett

1970 – UNIX is developed at Bell Labs by Dennis Richie and Kenneth Thomson. The floppy disk and daisy wheel printer make their debut.

— We are all indebted!

1977 – Apple is incorporated and the Apple II is introduced and becomes the benchmark for personal computers. PCs from Tandy and Commodore come with built-in monitors and require no television hookup. The Cromemco Z2 system, an evolution of the Z1 model, was introduced with a new CPU board using a 4-MHz Z80 microprocessor.

— Greg Rottman, Dave Stinson,
Kevin Williams, Jay Cliburn

1984 – IBM introduces the PC/AT, the third generation of the PC series. Sony and Phillips introduce the CD-ROM that provides significantly greater storage capacity for digital data.

— Paula Lindsey

1990 – Microsoft introduces Windows 3. Berners-Lee writes the initial prototype for the World Wide Web, URLs, HTML, and HTTP.

— We are all indebted!

1968 – The Seymour Cray-designed CDC 7600 supercomputer achieves 40 megaFLOPS performance.

— Fred Tracy

1975 – The first PC, the Altair, appears on the cover of Popular Electronics magazine.

— Bob Tipton

1976 – The Cray 1, Cray's first supercomputer, has a peak performance of 133 megaFLOPS.

— Jeff Hensley, Bill Ward,
Bob Tipton

1980 – Timex Sinclair introduces the ZX-80 and inaugurates the transition between the hobbyist world and consumer electronics by proposing a true computer in its case for less than \$100.

— John West, Bobby Hunter

1985 – Cray introduces the XMP nearly a decade after the Cray 1.

— Jay Cliburn, Jeff Hensley

2003 – The Cray X1, designed to be the world's most powerful supercomputer product, features ultra-fast (12.8 gigaFLOPS) individual processors, up to 819 gigaFLOPS of peak computing power in a single chassis.

— Our journey to the sustained petaFLOP continues...

ERDC Facts

Aggregate HPC Experience
of MSRC staff: 578 years

Aggregate IT Experience: 1,130 years

% Staff with advanced degrees: 23%

% Staff with Ph.D.s: 10% Teaching experience:
colleges and universities: 89 years

Oldest, unwashed coffee cup
still in active service by a
staff member at the MSRC:
Commission date, 1982

SAME Engineering and Construction Camp a Learning Experience

By Ginny Miller

Twenty-seven high school juniors and seniors interested in careers as engineers took part in a weeklong Engineering and Construction Camp during the summer, with every ERDC laboratory at the WES site, including the Information Technology Laboratory (ITL) and the ERDC MSRC, participating.

The camp, sponsored by the Vicksburg and Louisiana posts of the Society of American Military Engineers (SAME), was designed to provide students with hands-on experience in engineering and construction skills, according to SAME member Leo Phillips, a civil engineer and Chief of the Construction Division at the U.S. Army Corps of Engineers Vicksburg District.

“The idea is to introduce high school students to engineering so as to attract the best and the brightest to the profession.”

— Leo Phillips

“The idea is to introduce high school students to engineering so as to attract the best and the brightest to the profession,” Phillips said, adding that the camp is modeled after a similar program at the Air Force Academy in Colorado Springs, Colorado. “Each day was intended to focus on an area of engineering.”

Professional engineers and volunteers from local engineering organizations supervised the camp, held June 8-14, 2003, and served as mentors to the students.

At ERDC, campers participated in a construction activity at the Geotechnical and Structures Laboratory and a river engineering activity coordinated by the Coastal and Hydraulics Laboratory. They also learned about soil and wetlands classifications during a visit to the Environmental Laboratory.

At the ERDC ITL, students visited the CADD-GIS Technology Center, where they participated in the hands-on analysis and design of a structural component of a building system. In the ERDC MSRC Scientific Visualization Center (SVC), SVC Lead Paul Adams taught students how to perform scientific visualization using the visualization package EnSight. Students then performed interactive applications on the SVC’s ImmersaDesk using space shuttle data.

After the EnSight tutorial, camper Kinney Dandridge, 16, of Scotlandville Magnet High School in Baker, Louisiana, said he liked “being able to take a problem



High school students participating in the SAME Engineering and Construction Camp laugh as they watch an immersive application in the svc

and solve it and find out why things work the way they do. This was fun.”

“All of it was fun,” said camper Amanda Huskey, 17, a student at Warren Central High School in Vicksburg. “I learned a lot.”

“I learned a lot about things that are everyday life,” said camper Cody Hebert, 16, of North Vermilion High School in Maurice, Louisiana. “You take for granted the slab that’s under your feet, but how much goes in to designing the concrete?”

The camp, which also included a cookout, swim party, and leadership and team-building activities, concluded with a graduation ceremony attended by Lt. Gen. Robert B. Flowers, Chief of Engineers. Flowers encouraged the students to pursue engineering as a career.



Richard Walters (far left) and svc Lead Paul Adams (second from left) assist campers Tim Clapp (second from right) and Stephen Mehn (far right) during a tutorial using the scientific visualization package EnSight

JSU Summer Institute Fosters Interest in HPC Careers

By Ginny Miller

Nine students representing seven minority institutions participated in the 7th annual Jackson State University (JSU) Summer Institute sponsored by the HPCMP Programming Environment and Training (PET) program at the ERDC MSRC. The 2-week Summer Institute was held June 16-27, 2003.

JSU serves as the host institution for the Summer Institute program, with the Engineering Research Center (ERC) at Mississippi State University also playing an instrumental role. Activities this year, held at JSU and the ERC, included workshops, seminars, and tutorials in the areas of high-performance computing (HPC), scientific visualization, Geographic Information Systems, Distributed Network Storage Management, meteorology, OpenGL programming, UNIX operating systems, and Access Grid Node operations.

“The Summer Institute gives underrepresented segments of the U.S. student population a first-hand look at research activities and the use of HPC while fostering interest in HPC careers,” said Reginald Liddell, PET Education Outreach and Training Coordinator/Technologist at the ERDC MSRC.

Selected participants were Fredrick Harris of Tougaloo College, Wanda Haynes and Jennifer Middleton of JSU, Takita Johnson of Benedict College, Tijan Kanteh and

Omar Kebbeh of Rust College, Deidra Artis of Winston Salem State University, Stacy Woodard of Paine College, and Jessica McElroy of Texas Southern University.

During a June 23 visit to the ERDC, Summer Institute students toured the Coastal and Hydraulics Laboratory, the Environmental Laboratory, the Geotechnical and Structures Laboratory, and the Information Technology Laboratory. Their visit to the ERDC MSRC included a demonstration in the Scientific Visualization Center and a tour of the Joint Computing Facility, where MSRC Director John West shared fun facts about the Center.

When she heard West say that the amount of cooling needed for the MSRC’s high-performance computers would provide cool air for 100 homes, Woodard was surprised. “This was something new and different for me,” the 21-year-old math major said. “It’s something that I think everybody should know a little bit about.”

Woodard and her peers completed the program by writing a report and giving a presentation on one of the technical areas they were exposed to during the institute. Certificates, stipends, and awards were presented to the students during a closing ceremony at the JSU e-Center.



ERDC MSRC Director John E. West (right) talks to PET JSU Summer Institute students before a tour of the Joint Computing Facility

National Science Foundation Scholarship for Service Cyber Corps Student Interns at ERDC MSRC

By Charles Ray

Mississippi State University (MSU) graduate student Trent Townsend recently completed an Information Security internship at the ERDC MSRC. Trent's work at the MSRC consisted of researching of a Sunscreen firewall product; patching Red Hat (7.1 and 7.3) Linux workstations; researching building RPM for Kerberos distribution; developing Perl scripts for parsing NMAP output/log files; and developing a CSA Windows script for the MSRC.

"My time at ITL (MSRC) was very beneficial to me and hopefully to the Center. I lacked 'real-world' experience in the computer security arena, and this summer remedied that along with confirmed that this is what I want to do when I finish with graduate school," said Trent.

The National Science Foundation (NSF) funds scholarship programs at 13 universities, one of which is MSU. The NSF does this to train individuals in cyber-security through its Scholarship for Service program.

Cyber Corps students receive 2 years of support that includes tuition payments and a cash stipend while enrolled in the program. If the opportunity is available, the students get to serve in Government internships. Graduates are then required to serve 2 years with the Federal Government upon completion of the program.

For the universities to participate, they must be reviewed and certified by the National Security Agency (NSA) as Centers of Academic Excellence in Information Assurance. In the shadow of Y2K, this program was initiated to safeguard the Nation's information infrastructure. The Bush administration has increased funding to the program since the terrorist attacks of September 11, 2001. Prospective Cyber Corps students make their application to participate through one of the approved universities.

For more information readers can go to the following Web site: <http://www.ciao.gov/education/>.



Trent Townsend delivers his final presentation to the ERDC MSRC and to others off station via the ITL Access Grid (photo by Richard Walters)



Trent (third from left) with ERDC MSRC team members (left to right) George Moncrief, Chris Fortenberry, Tim Dunaway, George Jackson, and Charles Ray (photo by Richard Walters)

Students Succeed in 2003 HPCMP PET Summer Intern Program

By Ginny Miller

College students who participated in the 2003 HPCMP PET Summer Intern Program say the experience has solidified their career choices and will help them achieve success.

“After this internship I know exactly the direction I want to go in,” said Omar Rodriguez, one of five participants in the program held June 2-August 8 at the ERDC MSRC.

After completing work toward his bachelor’s degree in computer science at Arizona State University in Tempe, Rodriguez, 20, plans to become a network engineer. “This PET internship has given me the opportunity to see people in the positions I wish to one day be in,” said the college junior, whose projects at the ERDC MSRC included working with the Repository in a Box Toolkit and the PET On-line Knowledge Center.

“This was an exciting summer filled with learning opportunities,” said PET Education Outreach and Training Coordinator/Technologist Reginald Liddell, who organized enrichment activities including lectures from each of the PET Onsite Leads as well as discussions about research ethics, technical writing, and graduate school strategies.

“This experience has given me the ability to know and understand what I want to do as far as a career.”

— Anthony Smith

“Over a period of 10 weeks, the PET Summer Intern Program provides students from across the Nation with hands-on research experiences,” Liddell said. “Students work independently but spend much of their time with mentors in an assigned research group.”

Joseph Cate, 20, a junior at Mississippi State University majoring in electrical engineering, worked with Dr. Bill Ward of the Computational Science and Engineering group managing/graphing data from HPC performance studies. He also created code for benchmarking tests via C and MPI and learned Fortran



The 2003 HPCMP PET summer interns are (from left) Anthony Smith, Leelinda Parker, Omar Rodriguez, Joseph Cate, and Brittany Owens

90. Cate said he most enjoyed “being able to see what other engineers actually do. This was a good experience and job training.”

Nineteen-year-old Brittany Owens attends Hampton University in Hampton, Virginia. While at the ERDC MSRC, the junior mathematics major was mentored by Dr. Gary Howell and also worked with Paul Adams of the Scientific Visualization Center (SVC). Her assignments included writing an easy-to-use Web interface to run a visualization program from a command line. “The people working here were very willing to help me when I had problems,” she said. “It was a great experience and it will definitely benefit my career.”

Leelinda Parker, 22, is a recent graduate of Jackson State University with a degree in computer science and mathematics. At the ERDC MSRC, Parker worked with Dr. Nathan Prewitt, the PET Computational Fluid Dynamics Onsite Lead, to create a GTK graphical user interface for truss structure analysis. “I liked the working environment,” Parker said. “This experience as a PET intern has given me more hands-on experience and knowledge of the computer science field, which will definitely be beneficial to my career.”

Anthony Smith, 29, is a graduating senior studying computer engineering at Central State University. “This experience has given me the ability to know and understand what I want to do as far as a career,” said Smith, who worked with the visualization tool EnSight while assigned to the SVC.

Watch Out for “Bites” When Writing About Bytes

By Rose J. Dykes

During the 10 weeks that the Programming Environment and Training (PET) interns were at the Engineer Research and Development Center Major Shared Resource Center (ERDC MSRC) this summer, they attended a weekly lecture that Reginald Liddell, the PET Training Coordinator, had set up to add variety to their learning experience. Technical writing was the subject for one of the lectures. Reginald had asked if I would present this one since I am a writer/editor for the MSRC. I was excited about the opportunity because I had enjoyed my earlier days as a high school English teacher. But then I decided that the lecture would probably be more impressive to these students in technical fields if given by a technical person that is also a good technical writer. (Besides, English teachers give some folks nightmares.) Also, someone came to my mind immediately that I met while working as an editor in the ERDC Coastal and Hydraulics Laboratory (CHL) that would be perfect if he were in the country. I gave Nick Kraus a call, and he agreed to do the presentation.

Dr. Nicholas C. Kraus, a Senior Scientist in the ERDC CHL, is an internationally known coastal scientist who has published more than 300 journal articles, technical reports, conference papers, and book chapters. In addition to his Ph.D. in physics that he earned in 1972, he was awarded an Honorary Doctorate from Lund University in Sweden in 2001. He was cited for his contribution to coastal engineering and for numerical models of shoreline change and erosion, which he developed in collaboration with three doctoral students from Lund. (Talk about having nightmares—I did when I was informed that I would be leaving my position as the editor of the ERDC Environmental Laboratory and going to be the CHL editor just down the hall from Dr. Kraus.) I realized immediately that he was an excellent writer, but always wanted everything reviewed for improvements. I have never known anyone else quite like Nick. Intense, intelligent, friendly, helpful, and appreciative with an insatiable appetite for knowledge is how I see him. He takes more pride in his work than most and will probably take his last breath while finishing a big project. As busy as he is, he still finds time to mentor young, aspiring minds. The PET interns were intrigued with him.

Dr. Kraus told the interns that his notes for his presentation to them were inspired by Strunk and White’s *The Elements of Style*, which he had read about 30 years ago. He recognized a need for a short “elements of style” applicable to technical writing after reviewing many engineering and science papers through the years. “We have scientific objectivity, mathematics symbols, units of measurement, and other specialized topics to deal with that do not enter non-technical writing,” said Dr. Kraus. He suggested that the copy of his notes that he presented to the students be used along with a copy of *The Elements of Style*.

“Enlightenment about my poor understanding of writing came one day while I was working in Japan, upon receiving proofreading results on one of my papers made by a non-native English-speaking colleague. I had written ‘...the curve rapidly changes...’ in reference to a variable plotted as a function of distance. My colleague said something like, ‘Doesn’t the word *rapid* relate to time, so that the phrase *steep gradient* would be more appropriate?’ At that moment I realized that every word must be written with care, and that manuscripts must be checked and rechecked before being submitted for publication.

“Many of those who have received similar stimulation (if not annoyance) from my proofreading have remarked on how they now see poor writing all around them. Indeed, once you read Strunk and White, much technical written matter will become disappointing. I call this the ‘bite of the Word Maven.’ Once bitten, you too will go forth with new eyes and sharpened pencil, recognizing the long path toward the goal of perfect writing. May these notes bite you!

“The reader can become distracted and even annoyed by sloppy writing. The originality, ideas, and conclusions of the document, and the corresponding credit, glory, fame, and fortune owed to the author might be lost if obscured by poor writing. Worse, reviewers may respond antagonistically to careless format and poor writing, basing decision more on the writing than on the content of the manuscript. Sloppy format and writing imply sloppy thinking. It is not the reviewer’s job or the editor’s job to rewrite the paper,” said Dr. Kraus.

**“Writing should not be stiff.
Word Mavens can bend the
rules, if they do it as a
conscious decision.”**

— Dr. Kraus



Dr. Nicholas Kraus, ERDC CHL, presents a workshop on technical writing to PET interns

Before discussing his rules of technical writing, Dr. Kraus talked about his “Zen Path of Writing.” This path included three main points:

“1. Strive to write well and never lower your standards. Continual practice will enable you to write better and more efficiently.

2. Seek good writing and able writers. I read the *New Yorker* and the *Atlantic Monthly* magazines in part to be exposed to quality professional writing.

3. Establish or join a proofreading group and associate with colleagues interested in improving their writing.”

“The Less Ink the Better,” “Select the Apt Word,” “Conventions for Mathematics Notation,” “Common Punctuation Pitfalls,” and “Do Not Mix Opinion with Objectivity” were just a few of the topics that Dr. Kraus included in his discussion of the rules of

technical writing. My favorite example that he used was in regard to “Serial Commas” under the “Common Punctuation Pitfalls” topic. He indicated that it was taken from *The Elements of Style*: “The will of the recently deceased rich uncle was opened and it said ‘I give all my money in equal parts to John, Fred and William.’ – Who would you like to be—John, Fred, or William? I recommend serial commas always as A, B, and C, not as A, B and C.”

In his “Final Note” on rules Dr. Kraus said, “Writing should not be stiff. Word Mavens can bend the rules, if they do it as a conscious decision.”

The PET interns left the lecture realizing that good technical writing skills are important and necessary if they want to become credible, productive professionals. They also realized that technical writing is more precise and objective than non-technical writing.


Monitoring Produces Insight About Service for the Future

By Rikk Anderson

Big Brother (BB) is a system-performance monitoring client-server application available as freeware on the Internet. It is not a real-time monitoring service but can operate in small time windows. The architecture is simple—with a monitoring part and a Web part. Monitoring is handled primarily by a BB client installed on each HPC and a few other systems important to the “Unbreakable Center.” In 5- to 10-minute intervals, each client executes tests on its local host. As testing occurs the clients determine the event and color status based on the thresholds set in its configuration files. An event occurs when the test result exceeds the threshold set for that specific test. The color status of the event is determined by how severe surpassing the set threshold really is. Most tests run by the client have two thresholds that determine severity. If no event occurs, the color status is defaulted a green “all is well” status. Once color status is determined, the client sends the results directly to the server with event results sent upon detection and with nonevent status results sent last.

The Web part of BB is handled by the BB server. The server receives the reports from the clients and logs them. In addition, BB has been modified to archive the logs in a comma delimited monthly file for insertion into a database. Big Brother then proceeds to read each log and produces a Web page with the name for the monitored system, the test executed on the system, and the color status of the test. In the case of an event, BB notifies the system administrator of the affected HPC by use of e-mail and text messaging. The BB Web page is meant to be interactive and designed to show the color status in small buttons of each test per system. To view more detail of the event, clicking on the color button will present a brief description of the event selected. The Web page of BB is currently being changed to present a more visual interactive effect. The change will present the details of an event along with color status, eliminating the need for a user to click on buttons.

Big Brother is not the “ultimate” monitoring system, and it does not sit among the expensive high-end monitoring (or status reporting) services like Site Assure. But for a freeware product, BB is a good,



Rikk Anderson sits in front of Big Brother

reliable, efficient service that can handle several clients reporting to it at once. It is not absolutely restricted to the default services and tests, such as connectivity, disk capacity, and CPU utilization. It allows change or addition to its functionality with the inclusions of one’s own personal shell and/or Perl scripts, for example, the “sar” test implemented for the Cray T3E. In addition to the inclusion of external scripts, changes can be made to the BB source to improve its performance.

Big Brother is used by management as well as systems staff members to maintain high levels of service availabilities and spot potential problems early enough to intervene. Since BB was included into routine operations, it has slowly become the constant source for performance monitoring by system administrators as well as program managers. As time goes on, with the addition of more tests and tweaking of performance, BB could become the primary source for retrieving system-performance statistics onsite and for the user community at large.

Almost “Sleepless in Seattle”— Users Group Conference 2003

By Rose J. Dykes



The HPCMPO hosted Users Group Conference (UGC) 2003 on June 9-13, in Bellevue, Washington, located in the Eastside economic and cultural hub of Seattle. Conference plenary sessions including a keynote speaker and other invited speakers for two mornings in addition to two all-day sessions and two other afternoon sessions of tutorials, Challenge Project presentations, and technical paper presentations filled each day for attendees with a variety of opportunities for hearing about the latest in HPC technology. A poster session one night and 20 birds-of-a-feather sessions on other nights were also offered. For the annual conference social, one night the attendees cruised around the seaport – one of the world’s busiest – and across Puget Sound to Tillicum Village on an island that is part of Blake Island State Park, where everyone enjoyed fresh Pacific salmon prepared in the traditional manner of the Northwest Coast Indians. Of course, most attendees found time on another night or two to take in the Seattle nightlife, including dinner at the 605-ft Space Needle, where they could view the city of Seattle, Puget Sound, Mount Rainier, and the Cascade and Olympic mountain ranges. With all of the intellectual and cultural opportunities and exciting entertainment available, time for sleeping was limited.

The UGCs bring together personnel from all of the HPCMP computing centers and the users of these resources and provide a forum for communication, training, and discussion of HPC and its impact on science and technology. The ERDC MSRC contributed to the 2003 conference by presenting two tutorials and five technical papers. It also hosted a booth that included a poster depicting the ERDC MSRC infrastructure and services, copies of the latest edition of *The Resource*, and copies of its fact sheets. “These contributions by the MSRC staff emphasize the important role our team plays in ensuring the DoD community has access not only to state-of-the-art systems but also to the tools and expertise to make the best use of those systems,” said ERDC MSRC Director John E. West.



The Conference Keynote Address was given by Dr. Charles Holland (right), Deputy Under Secretary of Defense (Science and Technology), standing with HPCMO Chief Scientist Dr. Robert Peterkin (left)



ERDC Participants



ERDC ITL Director Dr. Jeffery Holland addresses a plenary session as a Conference Invited Speaker



ERDC Geotechnical and Structures Laboratory scientist Tommy Bevins presented a technical paper entitled "Multiple Building Simulations and Effect of Berms Behind Blast Barrier Walls"



ERDC MSRC Participants - Tutorials



Tom Biddlecome (standing) and Paul Adams (far right) present an intermediate and advanced scientific visualization tutorial



Carrie Mahood (seated left) and Dr. Tom Oppe (standing) present a tutorial on "Dual-Level Parallelism"

Poster Session

(Left to right) Dr. Jeffery Holland, Bradley Comes, John E. West, and Carrie Mahood attend the Conference Poster Session

ERDC MSRC Participants - Technical Papers




"Parallel Particle Tracking Algorithms in Computing Overland Flow" was presented by Dr. Ruth Cheng




Dr. Jeff Hensley discusses the "Use of Adaptive Meshes in ADH for Flow Problems"

ERDC MSRC Participants - Technical Papers



Dr. Paul Bennett talks about "Creation of a Synthetic Application of Benchmarking"



Bobby Hunter presents "Algorithm Study of N-Layer Interface Synchronization."



"A Survey of the Algorithms in the TI-03 Application Benchmark Suite with Emphasis on Linear System Solvers," was presented by Dr. Fred Tracy



Others from the ERDC MSRC



Drs. Nathan Prewitt (left) and Rick Weed



Lesa Nelson



Dr. Gary Howell



Dr. Wayne Mastin



Dave Sanders



(Left to right) Carrie Mahood, Rose Dykes, Dean Hampton, and Bob Alter



Reginald Liddell

PET Training Moves to New Facility

By Lesa Nelson and Dr. Wayne Mastin

The PET training program took a big step forward with the recent move from the old PET Training and Education Facility (TEF) to the newly renovated ITL Collaboratorium. The new room gives more space for larger classes, and a flexible room configuration permits seating to be reconfigured for specific audiences. The ERDC Access Grid node has also been moved to the training room. The Access Grid is a tool for group-to-group collaboration over the Internet and is being used both to deliver training to ERDC from remote sites and to deliver ERDC-produced training events to remote sites with an Access Grid node. The distance training efforts of the PET program have also been enhanced with the installation of RealNetworks products for webcasting seminars and training courses to remote users at any location with Internet access. The RealNetworks system is also used to archive training events on the PET On-line Knowledge Center (OKC) and produce CDs for distribution to users.

The inaugural PET training event in the new facility was a tutorial on Arbitrary Lagrangian-Eulerian methods delivered on May 20 by Dr. David Littlefield, University of Texas at Austin, the Computational Structural Mechanics (CSM) Functional Area Point of Contact (FAPOC). Since then, the PET program has made frequent use of the facility and the new equipment. The Access Grid has been used for several PET training events. Seminars were delivered by Dr. Rhonda Vickery, Enabling Technologies (ET) Onsite at ASC, from Mississippi State University, and Dr. Geoffrey Fox, OKC FAPOC, from Indiana University. The event that best utilized the capabilities of the Access Grid was the summer intern presentations.

Interns from both ERDC and NAVO delivered their final presentations from their local node. The audience for the presentations included the interns and mentors from ERDC and NAVO, Dr. Sue Brown, Education, Outreach, and Training Coordination (EOTC) FAPOC, who participated from the University of Hawaii, and PET personnel from Jackson State University, the home institution for one of the ERDC interns. Of course the interns were extremely nervous, but performed well under pressure and enjoyed chatting with other interns and participants at remote sites.

The move to the new PET training room coincided with the acquisition of new laptop computers for hands-on training sessions. The laptops use the Windows OS with the open source package Cygwin installed to deliver a UNIX/Linux environment for running applications on the laptops or connecting to the MSRC HPC machines. This system was first used for a training course on Zapotec delivered on August 20-21. Zapotec is a system for coupling the Eulerian code CTH with the Lagrangian code Pronto and is applied extensively by DoD users for simulation of weapon-target interaction. During the training course, users were able to run Zapotec on the ERDC MSRC SGI Origin systems and view graphical output on their laptops.

The transition to the new facility was not without occasional delays and glitches, as is any effort involving new technology; however, the move raises PET training capability to a higher level. Both instructors and users, who experienced training in quarters the PET training program occupied for the past 7 years, have highly praised the new facilities.



ERDC summer intern Leelinda Parker from Jackson State University delivers her final presentation on "GTK GUI Truss Structure Analysis" over the Access Grid

ERDC MSRC Hosts Four Groups During Same Week

By Rose J. Dykes

The ERDC MSRC hosted three important meetings that involved approximately 25 out-of-town folks along with approximately 20 ERDC folks during the week of September 8.

On Monday, a Benchmarking Session was held involving people from the HPCMPO; the User Advocacy Group (UAG); the NAVO, ARL, ASC, and ERDC MSRCs; and Instrumental, Inc. The purpose for meeting was to discuss the benchmark scoring of proposed systems from vendors for TI-04.

An Information Environment (IE) meeting was held on Wednesday that involved “a small core group of IE stakeholders and users... structured in a similar fashion to an integrated product team (IPT) to ensure that the Government and developers are working as a team to deliver a successful product in a streamlined and efficient manner,” said Larry Davis, HPCMP Deputy Director. The purpose of this IE-IPT gathering was to “identify and prioritize what the (IE) follow-on development activities should be and the milestones for each.”

On Thursday, the UAG convened for one of its two to three meetings a year to discuss issues that need to be

brought before the entire Program for the good of the whole user community. This group influences policies and practices of the Program; facilitates the exchange of information between the user community and the HPCMP; serves as an advocacy group for all HPCMP users; and advises the HPC Program Office on policy and operational matters related to the HPCMP.

In addition to these three meetings the same week was the first of 2 weeks that the ERDC MSRC security posture was being evaluated by the Comprehensive Security Assessment Team. This annual accreditation and assessment process results in certification that enables the ERDC MSRC to connect to the Defense Research and Engineering Network (DREN).

With many out-of-town folks here at the same time, the MSRC decided to arrange for a little Southern Hospitality with a social on Wednesday night at Magnolia Bluff, a lovely Southern mansion on the banks of the Mississippi River. All hopefully enjoyed the delicious food, great fellowship, and outstanding entertainment by David Stinson, ERDC MSRC, and Larry Davis, HPCMPO, on the saxophone.



Benchmarking attendees



Information Environment-Integrated Product Team



Larry Davis and David Stinson entertain on saxophone at social



User Advocacy Group



Comprehensive Security Assessment Team



Out-of-town guests and ERDC MSRC team members attend social

ERDC MSRC Team Member Serves at ERDC Graduate Institute



Dr. Oppe instructs Applied Mathematics class

Dr. Thomas Oppe, a member of the CS&E group at the ERDC MSRC, is serving as an instructor for Fundamentals of Applied Mathematics I this fall at the ERDC Graduate Institute located in Vicksburg, Mississippi. Dr. Oppe is an adjunct professor in mathematics for Mississippi State University (MSU).

The ERDC Graduate Institute is an association of universities and ERDC through which academic credit and graduate degrees can be earned from member universities. Louisiana State University, MSU, and Texas A&M University are the member universities affiliated with the Institute. Dr. C. H. Pennington is the Director of the Graduate Institute.

ERDC MSRC Mass Storage Archival System Expansion

By Jay Cliburn

The ERDC MSRC is pleased to announce an expansion to its mass storage archival system (MSAS), adding more than 200 TB of magnetic tape near-line robotic storage for a new total capacity of over 400 TB available for user data.

The MSAS consists of two major elements: (1) a data management system (DMS) consisting of two high-availability Sun Microsystems servers, a 4-TB RAID 5 disk array in a Fibre Channel storage area network (SAN), and three StorageTek robotic tape silos; and (2) a remote management system (RMS) consisting of a single Sun server, a modest disk cache, and a single StorageTek silo. The RMS serves primarily in a disaster recovery role and is located in a building some distance away from the MSRC computing facility. Two copies of every user file are stored—one in the local DMS, and a second copy in the RMS. File fetches are always attempted first on the local DMS, but if that fails the file is fetched from the RMS.

The most recent upgrade includes replacing all four tape drives in the remote facility with new StorageTek 9940B tape drives, each drive providing a nominal capacity of 200 GB per media cartridge, bringing the capacity of the remote facility from 198 to 460 TB. The upgrade also adds four 9940B drives to the local DMS facility and its existing eight 9840 drives, bringing its capacity up from 200 to 400 TB. StorageTek 9840 drives are faster than the new 9940B drives, but the capacity – at 20 GB per cartridge – is much lower. The storage strategy will have larger, older files written to 9940B (slower, denser) media, and smaller, newer files written to 9840 (faster, less dense) media.

As always, users should continue to move data to and from the MSAS using the *msfput/msfget* family of commands available on each HPC system. The upgrade described in this article has no effect on the mechanics of moving data and should be practically invisible to users, except for making more archive space available. Contact the Customer Assistance Center for additional information.



Jay Cliburn
Technical Operations
Manager
ERDC MSRC

Software Profiling: Focusing on Interprocessor Communications

By Dr. Sam Cable

The software profiling effort of the ERDC MSRC Computational Science and Engineering (CS&E) group, reported previously in *The Resource*, continues. “Profiling” means examining a piece of software and determining how it spends its time and resources as it calculates whatever its developers designed it to calculate. Profiling might entail, for example, determining how many floating-point operations a program performs in a typical run, or how many times it must store numbers in memory or retrieve them to perform a calculation. Software developers and users find such information useful and even critical to their tasks of creating efficient programs and fine-tuning them to run on specific machines.

The CS&E group is enhancing its profiling capabilities with the introduction of a library of tools for specifi-

cally profiling those parts of a code dedicated to inter-processor communications. Much of the “super” in a contemporary supercomputer originates in the fact that the supercomputer can bring to bear the combined strength of scores of processors (CPUs) on a computational problem. So, on a supercomputer, a large problem can be divided among 100 processors, and then can be solved in one-hundredth of the time one computer would take.

Or maybe not. Here, as in many of life’s endeavors, too many cooks spoil the broth. There are very few problems of interest that can be broken up into truly



Dr. Sam Cable
Computational Scientist
ERDC MSRC

isolated pieces. In a fluid flow problem, for example, fluid will constantly flow from one processor to another (“virtually” speaking, of course). Therefore, the processors will have to constantly “check up” on the progress of one another to make sure their individual calculations are mutually consistent. Therefore, interprocessor communication is needed. This communication requires time and computational resources on top of the independent computations that the processors perform. So, the independent processors have fewer computations to perform, but now they have to perform communications operations that a single computer does not need. For this reason, even though 100 processors might solve the problem much faster, 1,000 processors (if that many could actually be found anywhere!) might actually be slower.

Hence the need for CS&E’s new profiling tools specific to processor communication. These tools will allow software developers and users to profile the work their programs do specifically in processor communication. On the one hand, such profiling would allow users to find inefficiencies in the communications portions of their programs. On the other hand, it would help users to fine-tune efficient codes to different computational platforms, since different platforms present different communication challenges, or would even help them decide which platforms to run on and which to avoid altogether.

The new profiling tools library runs with codes that perform processor communications via the commonly used Message Passing Interface (MPI). The library is based on the Performance Application Programming Interface (PAPI), a general profiling library developed by researchers at the University of Tennessee, Knox-

ville. PAPI can interact directly with event counters in the computer’s CPUs, starting the counters, retrieving the number of events they have recorded, and stopping them when no longer needed. The specific events that can be counted vary from machine to machine. Most machines will provide basic counts such as CPU cycles and floating-point instructions performed. More detailed counts such as the number of times a particular type of floating-point operation is performed, or the number of occurrences of cache hits or misses are also often available. Users will make use of the library by instrumenting their codes in areas they want to particularly scrutinize (see Listings 1a and 1b), and then linking with the library during code compilation. The code will run as usual, but will give the users extra output detailing how much time the code took to perform the computational tasks specified by the user. The CS&E library has been used to profile pure FORTRAN codes for some time now. This summer, the capability of profiling pure C codes and FORTRAN/C hybrids was added.

Tests of the new library have produced interesting preliminary results. A simple communication-intensive test code was instrumented and linked with the new library and was run on one of ERDC’s SGI Origin 3900 (O3K) machines. The number of computational cycles the code spent in MPI communications was found to be surprisingly variable. (See Figure 1.) Several consecutive runs of the code were made in four different batch jobs, each batch job being run on a different date. In each job, the code was run five times each through 50 iterations, 100 iterations, 500 iterations, and so on up to 100,000 iterations. As seen in Figure 1, the number of

<pre>PROGRAM CALCULATION IMPLICIT NONE REAL (KIND=8) :: PI INTEGER (KIND=8) :: N . . . CALL CALCULATE_LAST_DIGIT_OF_PI (PI, N) . . END</pre>	<pre>PROGRAM CALCULATION IMPLICIT NONE REAL (KIND=8) :: PI INTEGER (KIND=8) :: N . INCLUDE "PAPI_GO.h" . INCLUDE "PAPI_VIEW.h" CALL CALCULATE_LAST_DIGIT_OF_PI (PI, N) INCLUDE "PAPI_VIEW.h" . INCLUDE "PMPI_PAPI_END.h" END</pre>
(a)	(b)


Listing 1. Instrumenting a code to use PAPI to perform profiling of a FORTRAN code, with special scrutiny given to a particular subroutine. The code in PAPI_GO.h will specify which hardware counters to record, and will begin their operation. The code in PAPI_VIEW.h will take snapshots of the hardware counter values, which can then be used to determine how many operations of various types occurred in the execution of the subroutine. PMPI_PAPI_END.h will stop the hardware counters and report their final values, along with the work done specifically in MPI. In this particular example, PMPI_PAPI_END.h will have some rather large numbers to report.

cycles spent in MPI communication varies proportionally overall with the number of iterations the code completes, as expected. However, the number of cycles within a single batch job at any single iteration value varies typically by 20 percent of the average, and by as much as 60 percent. Further, the variation in number of cycles from one job to the next is far larger, particularly when the code performs higher numbers of iterations. For instance, the program spent about four times as many cycles performing MPI communications on August 1 as it did on August 26 and 27. This result needs further investigating, but it indicates that a variation in load on the machine from other programs can affect MPI efficiency much more than expected. If this result holds true on further investigation, and on other platforms, it means that users will have to make allowance for a great variability in the efficiency of their codes, particularly if the codes spend a great deal of time in communication.

The same tests have revealed a surprising difference in how MPI is implemented on different machines. Tests on ERDC's SGI O3K show that the MPI routine `MPI_Recv()` makes about 50 floating-point stores in dealing with a single 8-byte floating-point number. The IBM Power 3 and Power 4 machines at the NAVO MSRC make somewhat less stores on average. However, on the ERDC MSRC's Cray T3E, `MPI_Recv()` makes no use of floating-point operations whatsoever. Judicious use of the new profiling library, then, may help users uncover differences in MPI implementations that will affect the efficiency of their code.

It has turned out then, that the CS&E profiling library, designed to help users profile their individual codes and the MPI communication within their codes, may be able to give further insight into the workings of MPI itself. Further testing will continue, and will soon expand to other platforms such as the NAVO MSRC IBM Power 3.

Figure 1. The number of CPU cycles a test program spends in MPI communication as a function of the number of iterations the program performs. Each data point represents one run of the code. Jobs consisting of several runs at iterations of 50, 100, 500, 1,000, 5,000, 10,000, 50,000, and 100,000 were performed on one of ERDC's SGI O3K machines on August 1, August 26, the morning of August 27, and the evening of August 27. Variability in cycles within a single job is typically 20 percent and as great as 60 percent. Variability between different jobs is even greater. For instance, about four times as many cycles were spent in communication on August 1 as were spent on August 26 and 27



Using SHMEM for Low-Latency Communication on the Compaq SC40 and SC45

By Dr. Gary Howell

FORTRAN and C codes on Crays, SGIs, and AlphaServer SCs can use the shared-memory (SHMEM) communication library. SHMEM performs one-sided communication by call active messaging. In SHMEM, one processor gets or puts data directly into the memory of another processor, without the participation of the second processor. Compared with message passing, active messaging requires less cooperation between two processors that exchange information, and requires less communication overhead. For the process to work, variables must be symmetric in the sense of having the same address on all processors.

The following are guaranteed to be symmetric variables:

- Local static variables
- Variables in common blocks
- FORTRAN variables modified by !DIR\$ SYMMETRIC
- C variables modified by #pragma symmetric

The following are not guaranteed to be symmetric variables and should not be used with SHMEM:

- Dynamically allocated variables
- Local variables
- Other stack-allocated variables

The SC version of SHMEM (in contrast to the Cray version) requires that the first SHMEM call on each processor be

```
call shmem_init() ! for Fortran
shmem_init(); /* for C */
```

Some supported SHMEM operations are shmem_put, shmem_get, shmem_barrier_all, and shmem_broadcast, as well as many others. The concept of active sets allows a subset of allocated processors to participate in a broadcast or reduce operation. The SHMEM library routines that provide multiprocess programs with access to a contiguous region of virtual address space are not supported in the Quadrics AlphaServer SC implementation.

A SHMEM user guide, which includes a sample C program, can be found at <http://www.quadrics.com>. The sample C program sping.c can be copied from /usr/local/sc45/usp/VAMPIRexampl on the SC45 (emerald).



Dr. Gary Howell
HP Applications Scientist

The Web page http://www.sdsc.edu/SDSCwire/v3.15/shmem_07_30_97.html has a link to a sample FORTRAN program.

To use SHMEM on an AlphaServer SC, one of the following include statements should be used:

```
include 'shmem.fh' !in FORTRAN
#include <shmem.h> /* in C */
```

To compile and link SHMEM programs with the FORTRAN SHMEM libraries,

```
f90 -o myprog myprog.f -lshmem -lslan
```

Or in C,

```
cc -o myprog myprog.c -lshmem -lslan
```


SHMEM offers low-latency communications. For example, the sping program referenced above “pings” by the shmem_put command in an average (internodal) time of 3.5 microseconds on emerald. In contrast, a Message Passing Interface (MPI) program using nonblocking sends and receives required 52.73 microseconds per “ping.” Active messaging by MPI required 28 microseconds per “ping,” while “pings” by paired synchronous MPI sends and receives required 7.1 microseconds. The following figure shows off-node bandwidth for nonblocking MPI sends and receives (MPI_IRecv and MPI_Isend) compared with transfer of data by SHMEM. For SHMEM, the maximal bandwidth is around 67 Mbytes/sec. For the nonblocking MPI program, the maximal internode bandwidth is higher at about 122 Mbytes/sec. The nonblocking MPI program (in data not plotted) obtained intranode communication of 390 Mbytes/sec, but for one-word messages (such as the internode MPI) required 52 microseconds per “ping.”

Legacy code being ported from a Cray or SGI machine may already be written in SHMEM. If code will be reused on newer Crays such as the X1, using SHMEM will continue to be appropriate.

Code developers migrating from SHMEM to MPI (or vice versa) may find it convenient to link to both libraries, using both MPI and SHMEM calls in the same code. One code optimization would be “hybrid code” that performs small messages by SHMEM and longer messages by MPI. GAMESS is a commonly used code that runs much faster in its SHMEM version than its MPI versions.

SHMEM low-latency communication may help many codes run much faster, but there are some reasons for preferring MPI.

As seen in the results above, MPI has a higher peak bandwidth for long messages. MPI has a wider functionality and is more portable. Since SHMEM is not a standard, some of its features are not supported on all platforms. SHMEM bindings to supported functions differ from platform to platform and may not be well documented. On some parallel platforms, e.g., IBM SPs and Linux boxes built with Myrinet connections, SHMEM does not exist. As MPI libraries are updated and refined, it may be that the “active messaging” parts of the MPI-2 standard will attain comparably low latencies to SHMEM.



Off-node bandwidth for nonblocking MPI sends and receives (MPI_Isend and MPI_Irecv) compared with transfer of data by SHMEM

upcoming conferences

The 2004 International Symposium on Collaborative Technologies and Systems (CTS'04)

January 18-23, 2004
 Catamaran Resort Hotel
 San Diego, California
<http://www.engr.udayton.edu/faculty/wsmari/cts04/cfp.html>

Advanced Simulation Technologies Conference 2004 (ASTC'04)

April 18 - 22, 2004
 Hyatt Regency Crystal City
 Arlington, Virginia
<http://www.scs.org/confernc/astc/astc04/cfp/astc04.htm>

High Performance Computing Symposium 2004 (HPC 2004)

April 18-22, 2004
 Hyatt Regency Crystal City
 Arlington, Virginia
<http://www.eng.uci.edu/~jmeyer/hpc2004/>

visitors



(Left to right) Peter N. Whitehead, Army Research and Technology Protection, U.S. Army Corps of Engineers (USACE), Washington, D.C., and David Stinson, ERDC MSRC, September 10



Judith Blake (center), Chief, Small Business Office, Headquarters, USACE, with Timothy Ables (left), ERDC Assistant to the Director, and Dennis Gilman (right), ERDC MSRC, September 3



(Left to right) Dr. Hugh Thornburg, ASC MSRC, Tom Biddlecome, ERDC MSRC, Brian Schafer and Dr. Rhonda Vickery, ASC MSRC, and Dr. Robert Moorhead, Mississippi State University, August 19



Bradley Comes (far left), HPCMPO, Dr. James Houston (2nd from left), ERDC Director, Republican Congressman from Ohio David L. Hobson (2nd from right), Chairman, Subcommittee on Energy and Water Development, House Appropriations Committee, Dennis Kern (left of Congressman Hobson), Appropriations Staffer, and Robert Vining (far right), Chief, Programs Management Division, USACE, July 29



Richard Walters (left), ERDC MSRC, and Lt. Col. Frank Green, United Kingdom, July 29



University of Puerto Rico-Mayaguez students working at ERDC this summer and Greg Rottman (far right), ERDC MSRC, July 17



(Left to right) Dr. Jeffery Holland, ERDC ITL Director, and John E. West, ERDC MSRC Director, shown with Dr. John A. Parmentola, Director for Research and Laboratory Management, Carolyn J. Nash, Deputy Director for Research, and Jim Wisnewski, Special Assistant to Dr. Parmentola, all from the Office of the Assistant Secretary of the Army for Acquisition, Logistics and Technology, Washington, D.C., July 1



COL Thomas C. Supler (center), Deputy Chief of Staff, Reserve Affairs, USACE, with David Stinson (left), ERDC MSRC, and Kent Turner (right), Chief of the Central Processing Center in CEEIS, June 12

acronyms

Below is a list of acronyms commonly used among the DoD HPC community. You will find these acronyms throughout the articles in this newsletter.

ARL	Army Research Laboratory	LES	Large-Eddy Simulations
ASC	Aeronautical Systems Center	MAC	Marker and Cell
BB	Big Brother	MIT	Massachusetts Institute of Technology
CADD-GIS	Computer-Aided Design and Drafting-Geographic Information Systems	MLP	Multi-Level Parallelism
CEEIS	Corps of Engineers Enterprise Infrastructure Services	MPI	Message Passing Interface
CHL	Coastal and Hydraulics Laboratory	MSAS	Mass Storage Archival System
CPU	Central Processing Unit	MSP	Multistreaming Processor
CS&E	Computational Science and Engineering	MSRC	Major Shared Resource Center
CSM	Computational Structural Mechanics	MSU	Mississippi State University
DDG	Guided Missile Destroyer	NAVO	Naval Oceanographic Office
DDX	Guided Missile Destroyer X Generation	NERSC	National Energy Research Scientific Computing
DMS	Data Management System	NFA	Numerical Flow Analysis
DNS	Direct Numerical Simulation	NSA	National Security Agency
DoD	Department of Defense	NSF	National Science Foundation
DOE	Department of Energy	OKC	On-Line Knowledge Center
DREN	Defense Research and Engineering Network	PAPI	Performance Application Programming Interface
EICM	Eulerian Interface-Capturing Method	PEs	Processing Elements
EOTC	Education, Outreach, and Training Coordination	PET	Programming Environment and Training
ERDC	Engineer Research and Development Center	RANS	Reynolds Average Navier Stokes
ET	Enabling Technologies	RMS	Remote Management System
FAPOC	Functional Area Point of Contact	SAIC	Science Applications International Corporation
GE	General Electric	SAME	Society of American Military Engineers
GTK	Graphical Toolkit	SAN	Storage Area Network
HPC	High-Performance Computing	SBW	Spilling Breaking Wave
HPCMP	High Performance Computing Modernization Program	SGS	Subgrid-Scale
HPCMPO	HPCMP Office	SHMEM	Shared Memory
IE	Information Environment	SVC	Scientific Visualization Center
IPT	Integrated Product Team	TEF	Training and Education Facility
IT	Information Technology	TI-03	Technology Insertion 2003
ITL	Information Technology Laboratory	UAG	User Advocacy Group
JSU	Jackson State University	UGC	Users Group Conference
LCS	Littoral Combat Ship	USACE	U.S. Army Corps of Engineers

training schedule

For the latest on PET training and on-line registration, please go to the On-line Knowledge Center Web site:

<https://okc.erd.c.hpc.mil>

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



ERDC MSRC Newsletter Editorial Staff

Interdisciplinary Coordinator

Mary L. Hampton

Chief Editors

Rose J. Dykes

Virginia "Ginny" Miller

Visual Information Specialist

Betty Watson

ERDC MSRC Customer Assistance Center

Web site: www.erdchpc.mil

E-mail: msrchelp@erdchpc.mil

Telephone: 1-800-500-4722

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

Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the DoD. Your comments, ideas, and contributions are welcome.

*Design and layout provided by the Visual Production Center, Information Technology Laboratory,
U.S. Army Engineer Research and Development Center.*

Approved for public release; distribution is unlimited.

Dr. Douglas Dommermuth, Kelli Hendrikson, Lian Shen,
Dick K. P. Yue, and Paul Adams

“Large-Eddy Simulation of Steep Breaking Waves and Thin
Spray Sheets Around a Ship”

