

New computers, new students, new connections ...

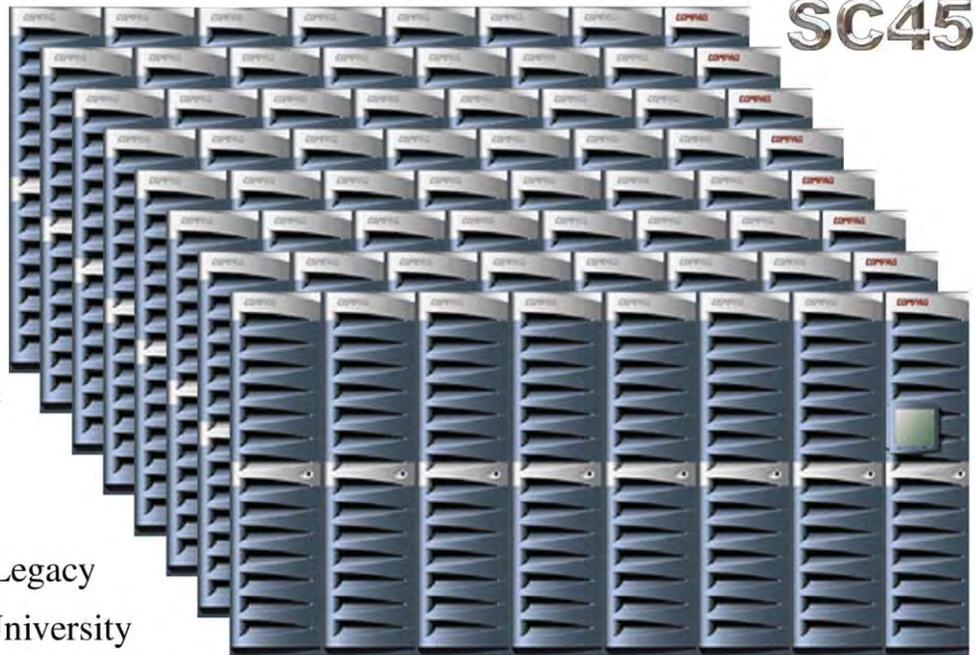


Keeping Pace with Technology

Compaq SC40 and SC45

Inside this issue:

- PET HPC Summer Institute
- Scientific Visualization
- Performance Benchmarking
- Adaptive Mesh Refinement
- HPCMP Users Group Conference
- HPC Tool Tips
- Access Grid Node
- A Reflection on Seymour Cray's Legacy
- HPC Activities at Jackson State University



ERDC MSRC

Major Shared Resource Center





In concert with the Programming Environment and Training (PET) program and the Common High Performance Scalable Software Initiative (CHSSI), in 1996, the U.S. Army Engineer Research and Development Center Major Shared Resource Center (ERDC MSRC) led the scalable systems initiative with the installation of a large T3E and two large IBM SPs. Along with deleting the vector-processor-based CRAY C90 and choosing not to follow on with a T90, we subsequently installed and continue to support a large SGI Origin. We focused for several years on these architectures successfully optimizing and porting users' applications. In January of 2000, we installed a large eight-way IBM SMP. This presented yet another opportunity to leverage "hybrid" programming techniques coupling MPI and OpenMP. However, we found that taking advantage of the full eight threads on large numbers of distributed SMP nodes was a rare and difficult task. We were faced with yet another fork in

the road. Do we scale the nodes to 16 or 32 way and focus on running shared-memory applications on one or at best a hand full of nodes, or do we downsize the size of the node and continue with the highly distributed scalable computing model? An intense study of the applications executing in the MSRC and some extensive benchmark testing answered the question—downsize the node; hence, the installation of the Compaq SC40 and SC45 systems introduced into the MSRC this fall.

I was a little skeptical about making a commitment to Compaq. I can remember the first time they showed up on the floor at the annual SuperComputing Conference. Everyone was snickering and making jokes about a Compaq supercomputer. Well, looking under the covers reveals that it is nothing more than a Digital Equipment Corporation (DEC) engineered system. In fact, most of the people I am communicating with today are the same folks as from the DEC days. The system uses the Alpha chip—the same chip the supercomputer leaders, Cray Research Inc., used in the T3E. It supports the four-way SMP node, which from a hybrid-programming standpoint, we have found to be a manageable SMP node. And, it has a high-speed switch as an interconnect. But next to the central processing unit (CPU), node architecture, and the switch is the family of software. It turns out that all of the heavy-duty software is brought forward from DEC. We (those of us around in the 80s) remember DEC as being a leader in server-based scientific computing. The intellect of the DEC team has produced an impressive set of highly optimized math libraries, programming tools, and compilers. Putting it all together has proven to result in a true supercomputer. The results of our benchmark test have proven this. It was not just the chip that was fast. We scaled applications into the 100s of processors, and it continued to perform very well.

So, we are pretty excited about our new additions and look forward to the new challenge of moving users and their applications to this system. The matching of compute architecture to applications continues to be a focus for the MSRC, and we look forward to the new opportunities that our newest addition brings us.


Bradley M. Comes
Director, ERDC MSRC

About the Cover:

New Technology. ERDC MSRC welcomes Compaq systems (see articles, pages 20 and 21).

Summer Interns. Five college students spend summer at ERDC MSRC (see article, page 10).

Features

Adaptive Mesh Refinement for Computational Structural Mechanics	6
A Brief History of Vector Supercomputers	8
Five Students Participate in ERDC MSRC HPC Summer Intern Program	10
ERDC MSRC Participation at Users Group Conference 2001	12
COL John W. Morris III, New Commander of ERDC	14
Impact Expands for Users with New PET Program	15
Dr. Fred T. Tracy – An Impressive Past and an Enthusiastic Future	18
Scientific Visualization Workshop – Bring Your Own Data	19

Departments

off-campus

PET Summer Institute at JSU on High Performance Computing	2
Scientific Visualization Course at Jackson State University	4

benchmarking initiatives

TI-02 Benchmark Project	5
-------------------------------	---

tool tips

Converting from PBS to LSF	16
----------------------------------	----

technology update

Upgrades to the ERDC MSRC's Computational Environment	20
A Technical Overview of the Compaq Supercomputer at the ERDC MSRC	21
Access Grid Multimedia Conferencing System	25

<i>visitors</i>	26
-----------------------	----

<i>ERDC MSRC publications</i>	27
-------------------------------------	----

<i>acronyms</i>	28
-----------------------	----

PET Summer Institute at JSU on High Performance Computing

By Ginny Miller

To 22-year-old Kylie Nash, the fifth annual Jackson State University (JSU) Summer Institute on High Performance Computing (HPC) was better than excellent. "It was 10 out of 10," said Kylie, one of 15 college students who participated June 4-15 in this summer's program.

Sponsored by the ERDC MSRC Programming Environment and Training (PET) program, the summer institute is intended to expose college students to research activities and the use of HPC. The summer institute also fosters interest in careers in high-performance computing and the Department of Defense (DoD) in particular.

"Oh, goodness. I've learned so much," said Kylie, a junior computer science major at JSU. "I've become aware of many aspects of computers. They can be used for more than just e-mail and typing papers. The way that technology is going, everything in a couple of years will be based on computers," she said. "If we don't learn now, we'll be left behind."

A sophomore math major from Houston, TX, Jessica Poole now has a better idea of which career path to take. "I don't want to be a teacher," the 19-year-old said. "There are other varieties of careers you can get into being a mathematics major. I want to apply my math to some type of engineering. I really feel like the summer institute has benefited me a lot."

"Getting students started into the HPC environment is what's key"

John West, Director of Scientific Computing at the ERDC MSRC

Though she also learned to build her own Web site, Kylie said her favorite part of the summer institute was learning about scientific visualization. "Now I want to go into computer animation," she said. "I just like being able to go to the computer and design something and see the finished product. It's amazing. When you get the end result, you can say, 'Hey, look what I did.'"

Kylie also said she was grateful to be accepted into the program, which exposes undergraduate students in mathematics, engineering, and science from Minority Serving Institutions (MSIs) to the



Pet Summer Institute participants watch a scientific visualization demonstration in the collaboratorium during a tour of ERDC on June 6

requirements for high-performance computing in solving problems of national significance. The objective of this effort, in the long run, is to increase the pool of HPC-trained scientists and engineers available to the DoD.

JSU, the PET MSI lead at ERDC, serves as the host institution for the HPC Summer Institute program. Since its establishment in 1997, the program has featured lectures and demonstrations given by MSRC staff and the PET university and onsite academic leads on the application of HPC to solving scientific and engineering problems.

This year, Paul Adams, Director of the ERDC MSRC's Scientific Visualization Center, demonstrated to students on how scientific visualization is used in HPC projects. "It was a great experience," said Paul. "The students really enjoyed it."

John West, Director of Scientific Computing at the ERDC MSRC, also lectured the students on scientific visualization and computational simulation. "Getting students started into the HPC environment is what's key," he said.

Students are not the only ones who benefit from the summer institute. "We have thoroughly enjoyed our work with Jackson State," said Dr. Wayne Mastin, ERDC MSRC PET Onsite Academic Lead. "This is our fifth year, and it seems like every year it gets better and better."

"I'm really pleased at how well the institute ran this year," agreed Bob Athow, MSRC PET Lead. "It has good heritage, and it gets better."

“I’m really pleased at how well the institute ran this year. It has good heritage, and it gets better.”

*Bob Athow
MSRC PET Lead*

During the first 4 years of the program, there were 72 participants from 11 MSIs from across the southeastern United States. This year, 15 students from five MSIs in Mississippi, Texas, and Florida attended the summer institute.

Joining Kylie from JSU were Jessica Gary, Christopher Grisby, Crystal Holmes, Christi Jackson, George Riley, Shayla Roundtree, Earlene Thompson, Lavaese Tillis, and Albert Williams. Other students in the program included Bojang Bakary and Ayorinde Hassan of Rust College; Andrew Francis, Jr., of Florida A&M University; Crystal Nevel of Tougaloo College; and Jessica Poole of Texas Southern University.

“This program is geared for those that might not otherwise have the opportunity to learn about supercomputing,” said Brenda Rascoe, administrative assistant to Willie Brown, JSU Vice President for Information Technology and a summer institute organizer. “It opened up a lot of opportunities for students. Now a lot of them see they have opportunities in this area.”

“I was just trying to add to my knowledge,” said Ayorinde Hassan, a 20-year-old chemistry major who hails from Nigeria, West Africa. “There are so many opportunities out there. We don’t need to limit them. We need to explore every opportunity we have right now.”

Andrew Francis, 21, a junior computer science major from Tallahassee, FL, also found the summer institute to be an educational experience. “Before I came here, I wasn’t sure what HPC was,” he said. “Now I have a better understanding. I learned a lot. I’m very thankful. Now I will look into getting my master’s in computer engineering. There are a lot of opportunities here.”

Many other summer institute students, including Crystal Holmes, are also planning to attend graduate school. “I didn’t have in my head at first about getting a master’s degree,” the senior computer science major from Jackson, MS, said. “I enjoyed all aspects of the program.”

For this summer’s institute, scientists and engineers from JSU and Mississippi State University (MSU), both ERDC MSRC PET partners, presented lectures and demonstrations. Highlights of the institute were a trip to the Engineering Research Center at MSU and the research laboratories at ERDC for a first-hand look at how researchers from various disciplines work together to solve problems of national significance to the DoD.

The ERDC tour on June 6 included a Command Briefing followed by visits to the McNary General Model in the Coastal and Hydraulics Laboratory, the Environmental Chemistry Laboratory, the Army Centrifuge Research Center, and the Information Technology Laboratory.

At the summer institute’s closing ceremony, held June 15 on the JSU campus, Ms. Rascoe praised the students for their hard work. “We are very proud of them,” she said. “I think they are very marketable individuals.”

“That is a credit to ERDC,” said Dr. Abdul Mohammed, Dean of the School of Science and Technology at JSU. “We are very fortunate to have a neighbor like ERDC very close to us,” he said. “ERDC has been very helpful to Jackson State University, very helpful to our students.”



Students in the PET Summer Institute on the JSU campus pose after closing ceremonies on June 15

Scientific Visualization Course at Jackson State University

By Ginny Miller

Dr. Kent Eschenberg, Programming Environment and Training (PET) Onsite Scientific Visualization Lead at the ERDC MSRC, taught a special topics course on scientific visualization at Jackson State University (JSU) during the summer. "It was a graduate-level class," said Dr. Eschenberg. "It consisted of lectures as well as laboratory exercises."

Dr. Eschenberg said that students worked on six different projects during the 8-week course, ranging from isosurfaces to vector fields. "Each of the projects began with a program that I would provide to get them started," he said. "We used a programming language called TCL/TK."

Seventeen computer science majors were enrolled in the course, which was arranged with the assistance of Dr. Loretta A. Moore, head of JSU's computer science department. The course was a collaboration between the ERDC MSRC and JSU, an ERDC MSRC PET partner.

"It was great to put together the course and to provide it. It was a great experience," said Dr. Eschenberg, who believes his students learned two valuable lessons. "They learned the concepts that are scientific visualization - that's useful no matter what particular program you're using," he said. "Second, they learned the applied aspects of making something work. It was more than just the concept in a vacuum."

TI-02 Benchmark Project

By Dr. William Ward



The Technology Insertion 2002 (TI-02) project requires the selection and installation of new high-performance computing (HPC) systems at the four Major Shared Resource Centers (MSRCs). As an aid in system selection, a benchmark test package is

used to measure the performance of candidate systems. The TI-02 benchmark test package was distributed to HPC system vendors on August 15, 2001. Vendors executed the tests specified in the package and returned the results in October for evaluation. New system selections will be announced in early 2002.

Test packages may contain a variety of types of tests. *Synthetic tests* consist of programs that do no useful work but to exercise various aspects of the system under test. Typically, synthetic tests are used to measure the peak performance of a particular system component, e.g., CPU or memory. *Kernel benchmarks* are computationally intensive portions of an actual code extracted for testing purposes executed by means of a driver program that often generates its own input data. *Application tests* use actual application codes and inputs. Regardless of their type, tests may be run either in *dedicated mode*, i.e., on an otherwise empty system, or in conjunction with other jobs. The former mode will measure best possible performance of a program since there is no competition for system resources. *Throughput or load tests* constitute yet another type of benchmark test; they involve imposing a mix of synthetic, kernel, and/or actual jobs on a system to mimic a production environment. The TI-02 test package includes dedicated synthetic tests, dedicated application tests, and a throughput test composed of applications. The latter two test components were constructed by the Computational Science and Engineering (CS&E) Group at the ERDC MSRC.

The dedicated application tests include seven codes: **Cobalt-60** (Computational Fluid Dynamics), **CTH** (Computational Structures and Mechanics), **GAMESS** (Computational Chemistry and Materials), **LESLIE3D** (Computational Fluid Dynamics), **NLOM** (Computational Weather and

Ocean Modeling), **Overflow** (Computational Fluid Dynamics), and **FDL3DAE** (Computational Fluid Dynamics). The first five codes were also in the Technology Insertion 2001 (TI-01) test package and accounted for more than 30 percent of the node hours consumed at the four MSRCs; they also use MPI as their primary parallel programmatic interface. Additionally, several of the codes provide other programming approaches. **GAMESS** provides a SHMEM version; **NLOM** may optionally be configured to use dual-level MPI/OpenMP parallelism; **Overflow** comes in both MPI and Multilevel Parallelism (MLP) versions; and **FDL3DAE** is a vector code that uses auto-tasking to use multiple vector processors. This diversity of both computational technology areas and parallel programming paradigms promotes the representativeness of the test package and aids in the determination of the strengths and weaknesses of the systems tested.

Typically, each of the codes in the test package is represented by two test cases (sets of input data), and each test case is to be executed on a range of processor counts in order to explore scalability on the test system. One of the test cases for each code is a so-called “large” test case that has a time target; a proposed system must meet this time target to be considered competitive. These large test cases are intended to represent the increasingly large processor count jobs observed in the DoD High Performance Computing Modernization Program (HPCMP).

It is expected that this test package will evolve over time, balancing the goals of providing a compact, easy-to-install and easy-to-use test package with that of comprehensively representing the HPCMP’s requirements. To that end, two initiatives are in progress. Increased use of kernel benchmarks in the test package is being explored to reach the first goal; for this effort to be successful, it will be necessary to show that kernels can effectively model the performance of the codes from which they are extracted. Next, long-term collection and analysis of HPCMP system utilization data are under way. These data will be used to select representative codes, test cases, and processor counts to be used in future test packages.

Adaptive Mesh Refinement for Computational Structural Mechanics

By Dr. Richard Weed

Two critical missions undertaken by Department of Defense (DoD) researchers are the modeling of large-scale explosive effects on military and civilian structures and the prediction of the response of weapons systems penetrating buried or hardened targets. Current state-of-the-art simulations utilize large parallel hydrocodes such as CTH for shock physics and explosive modeling, as well as structural dynamics codes such as DYNA3D and EPIC to model structural response to blast effects and for weapons penetration analysis. The computational requirements for a detailed analysis of an explosive event such as a terrorist attack on a military or civilian structure are enormous. A high-fidelity simulation of a high-explosive device and the resulting blast wave might require tens of millions of grid cells and several thousand central processing unit (CPU) hours on large-scale parallel systems such as the CRAY T3E to obtain an accurate solution. These large computational resource requirements limit the number of simulations that can be performed in a given time frame. A clear need was demonstrated in the early phase of the Programming Environment and Training (PET) program for improved computational methods that would reduce the time required for large-scale simulations without sacrificing the required levels of accuracy.

After an analysis of problem requirements, the ERDC PET Computational Structural Mechanics (CSM) team, led by Professors Tinsley Oden and Graham Carey of the Texas Institute for Computational and Applied Mathematics (TICAM) at the University of Texas at Austin, recognized that current and evolving technologies from the areas of adaptive mesh refinement (AMR), finite element error analysis, and mesh partitioning algorithms could be used to improve the efficiency of existing CSM analysis tools such as CTH and EPIC. In particular, AMR technologies in which computational meshes are refined locally in response to changes in predefined solution variables promised to greatly reduce the number of grid cells required for high-fidelity simulations. The research conducted by TICAM focused on the following areas: development of a posteriori error estimates and indicators, strategies for mesh adaptation, complications that arise from implementing the adaptive schemes in legacy codes such as CTH, develop-

ment of appropriate data structures and partitioning strategies for parallel analysis, and scalable parallel algorithms.

Efforts in the first year of the project focused on the development of error and feature indicators that would be used to drive the mesh adaptation and implementation of AMR into CTH and EPIC. Because of the importance of the CTH activities to both the DoD and Sandia National Laboratories, a process was put in place for collaboration between Sandia and TICAM on the CTH AMR development. TICAM focused on the development of the underlying AMR technologies, and Sandia focused on implementation issues related to its production version of CTH.

In the first year, a block refinement adaptation strategy utilizing 2:1 mesh refinement was implemented in a single-processor version of CTH. Initial results utilizing a hypervelocity impact analysis of a spherical copper projectile with target demonstrated that the AMR code produced equivalent results to a comparable uniform grid analysis with approximately one-third as many grid points.

In conjunction with the CTH work, a local simplex refinement strategy was implemented in a developmental version of the EPIC finite element code. This code was used to analyze several different error and feature indicators designed for Lagrangian impact codes such as EPIC. Other issues related to Lagrangian codes, such as mesh partitioning algorithms and mesh quality measures, were also investigated. Work was initiated on the development of a testbed for evaluating different types of error indicators.

In the next year, TICAM activities focused on the continued development of the AMR version of CTH along with work on developing error indicators for Lagrangian impact analysis. CTH development work included continued collaboration with Sandia to implement the basic AMR technology, implementation of the AMR scheme in parallel, and modifications to the existing multimaterial analysis capability in CTH to support both adaptation and parallel analysis. A significant result from this work was the demonstration of the effectiveness of AMR for blast wave simulation. This result is illustrated in Figure 1. The time required for a

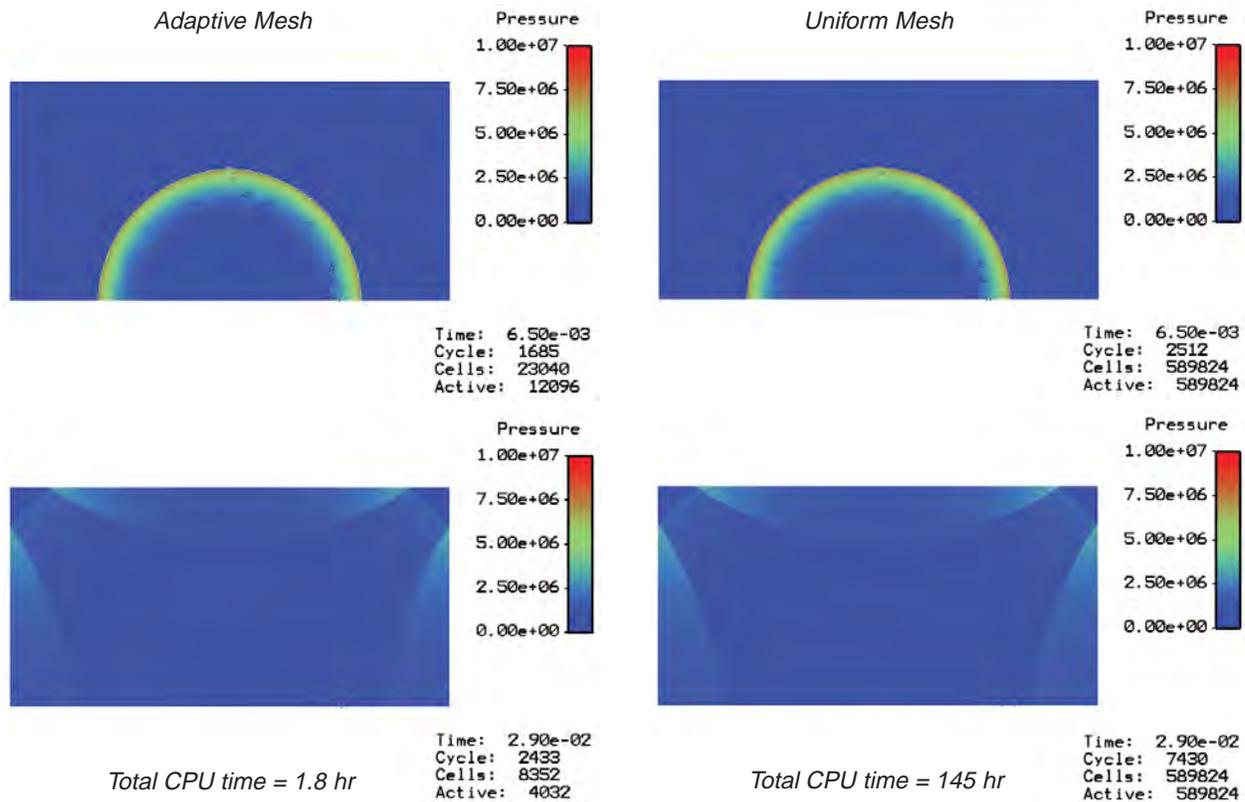


Figure 1. Comparison of adaptive mesh refinement (left) and uniform mesh solutions (right) for a 2-D blast wave computation in CTH. AMR, finite element error analysis, and mesh partitioning algorithms are being used to improve the efficiency of existing CSM analysis tools such as CTH and EPIC.

single-processor analysis of the two-dimensional blast wave propagation was reduced from 145 hours for a uniform mesh to 1.8 hours with AMR.

In the past year and a half, TICAM has worked with Sandia to deliver a production version of CTH that includes AMR. Sandia released the initial production version of CTH with AMR capability to the ERDC and other CTH users in April 2001. This version of CTH has been installed on all ERDC HPC platforms. CTH users should contact msrhelp@erdc.hpc.mil for information on how to access this version. In addition to the production version of CTH, Dr. David Littlefield of TICAM has implemented improvements to CTH requested by ERDC Geotechnical and Structures Laboratory (GSL) researchers. These improvements include implementation of a solid

obstacle analysis capability and improved methods for exporting data from CTH for use by structural response codes such as DYNA3D. This version of CTH is being implemented on the ERDC MSRC systems and will be made available to CTH users after final verification of the new features. Delivery of the production version of CTH and the enhanced version for the ERDC GSL meets the initial long-term goal of the project to utilize AMR technologies to dramatically improve the analysis capability of DoD researchers in the area of blast wave and explosive effects modeling. Continued development and implementation of the AMR technologies into DoD CSM analysis tools will be a major focus of future PET research. These technologies will greatly enhance the ability of DoD researchers to produce accurate simulations at a reduced computational cost.

A Brief History of Vector Supercomputers

By Dr. Arthur Cullati

Supercomputer is a term used for the most powerful machine available at any given time. There are two approaches to supercomputing. One is fewer, faster, specialized, more expensive processors. The other approach is many, slower, less expensive, commodity processors. This second approach is the one currently dominant in the United States.

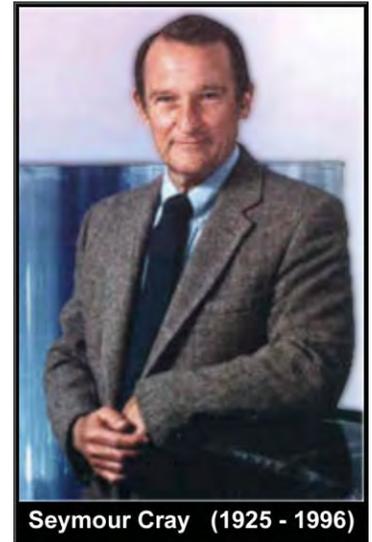
A vector supercomputer (vectors to be defined later) is one used primarily for the generation of floating-point numbers. As such, these machines have been called “number-crunchers,” and the performance has been measured in Floating Point Operations Per Second (FLOPS). These machines typically have been used in scientific simulations, computational fluid dynamics, computational chemistry, etc.

The father of supercomputing is Seymour Cray (1925-96). His genius has put him into the inventor’s hall of fame alongside such greats as Alexander Graham Bell. Cray’s design and ideas transversed three decades of supercomputing. He was one of the founders of Control Data Corporation (CDC), which was the leader in high-performance computers in the 1960s. Cray designed the first fully transistorized high-performance computer, the CDC 1604, in 1958.

Cray’s first step in the long succession of supercomputers began with his CDC 6600, which was first delivered in 1964. His idea was to take the unified CPU, which could only execute one instruction at a time, and break it up into many small processors, called functional units, each specialized to do a specific task, e.g., addition, multiplication. Once started, an add and a multiply could be going on simultaneously, as opposed to a unified processor where an add would have to wait for the multiply to finish before starting execution.

Cray’s next step was the 7600, which was delivered in 1969. It was a boxlike structure with blue glass that allowed one to see the modules and wires. A significant advance in the 7600 was the addition of segmentation into the functional units. The units were treated as if they were assembly lines, e.g., in an addition operation there are discrete steps, e.g., the exponents must be resolved, the number must be added. Each one of

these processes was separated so that not only could an addition and a multiply be going on simultaneously, but many adds could be going on simultaneously, each one at a different stage. These machines were called “pipelined” because the instructions flowed through the segmented functional units much like parts on an assembly line.



Seymour Cray (1925 - 1996)

In 1972, Seymour Cray left CDC and founded his own company, Cray Research Inc. In 1976, he delivered his first vector machine, interestingly enough called the Cray-1, to Los Alamos National Laboratory. The Cray-1 followed the evolutionary trend from the 6600 and 7600, with the addition of vector processing. A vector is a related sequence of numbers. Since much scientific work deals with matrices, the ability to manipulate horizontal, vertical, and diagonal vectors of a matrix is of great interest to the scientist. In machines before the Cray-1, arithmetic was performed in scalar mode, i.e., operation on each cell in a matrix at a time. In vector processing, once two vectors are put into two vector registers, a single math function can add, subtract, multiply, etc., the contents of both vectors together in one operation. This capacity significantly impacted the speed of scientific computing.

In 1982, the Cray-XMP was delivered, which was Cray-1 computers linked together to perform multiprocessing and multitasking among processors. Meanwhile, as the Cray-XMP was coming into fruition, Cray was busy working on his “Cray-2.” Actually, Cray never meant that the Cray-2 would ever become a real product. He intended to build a model in silicon as a proof of concept and then build his commercially available machine using gallium arsenide (GaAs).

The first four-processor Cray-2 was delivered to NASA Ames Research Center in 1985. It was very different in function from the Cray-1 and was the first supercomputer to run a commercial version of UNIX. When delivered the Cray-2 had

more central memory than all the other Cray computers ever delivered to that date. Many think that it was Cray's prettiest machine, with all the components, processors, memory, and wires totally submerged in an inert fluid and visible through glass portals.

At the same time, the Cray-XMP was undergoing a refurbishment. The XMP was a highly successful machine and the most powerful machine in the world, but the technology was getting old. Cray Research (a company once led by Seymour Cray) recast the XMP with new technology and some improvements, such as memory address space, and delivered the first Cray Y-MP to NASA Ames in 1988.

The next step in this evolutionary process was the Cray C-90. The suffix C-90 was chosen because it was to be the supercomputer of the 1990s. The C-90, first delivered in 1991, was basically a YMP with the additional feature of the capability to deliver two floating-point operations per clock period. The C-90 is considered by some to be the best vector machine ever produced in the United States.

The last player in this scenario is the T90, which was announced around 1995. This machine never became successful. It was too expensive and unreliable, and benchmarks indicated that the performance was only 1.5 to 2.0 times that of a C-90.

Cray's research into a GaAs computer was proving too costly for Cray Research Inc., so he founded a spin-off called Cray Computer. The cost became too great. GaAs is difficult to work with, and the yield of good chips from vendors was very low. Cray had to build his own foundry. Eventually Cray Computer went bankrupt in 1995.

On November 11, 1996, Cray Research Inc. announced that a 2,048-processor Cray T3E-900 broke the world's record for a general-purpose supercomputer with 1.8 teraflops performance. This computer is not a vector computer but one with commodity Alpha chips. This type of price performance made most laboratories choose to go the parallel route for future supercomputing.



Five Students Participate in ERDC MSRC HPC Summer Intern Program

By Ginny Miller

For five college students, summer internships at the ERDC MSRC were more than summer jobs. “It was hands-on experience,” said Eddie Barnes, an Alcorn State University junior majoring in industrial technology. “You get knowledge about how it actually is in the working field.”

A native of Port Gibson, MS, Eddie worked in the ERDC MSRC’s Scientific Visualization Center (SVC) during this year’s High Performance Computing (HPC) Summer Intern Program. “I did tutorials with the animation software Maya and Ensight,” the 20-year-old said. “I’d never done animation before.” He said that a favorite aspect of his internship was the mentoring he received from SVC Director Paul Adams and the entire SVC staff. “I worked with a lot of nice people who showed me a lot of different things,” he said. “That’s probably one of the best things.”

Other summer interns and their mentors included Joyce Beal, a senior at Alcorn State University who worked with Dr. Dan Duffy; Richard Anderson, a recent graduate of Clark Atlanta University who worked with Dr. Wayne Mastin; David O’Gwynn, a master’s degree candidate at Mississippi State University who worked with Dr. Nathan Prewitt; and Owen Eslinger, a Ph.D. student at the University of Texas who worked with Drs. Fred Tracy and Stacy Howington.

John West, Director of Scientific Computing and coordinator of the ERDC MSRC’s HPC Summer Intern Program, said that eighteen students applied for the internships. “All of the applicants offered unique talents, and selecting the top four was a difficult task,” he said, noting that the fifth student was not part of the competitive selection this year. “Richard Anderson returned to us this summer for his third summer internship experience at the MSRC, and we were fortunate to get back such an experienced partner.”



Summer interns at the ERDC MSRC are, front from left, Joyce Beal and Eddie Barnes, both from Alcorn State University. Back from left are Owen Eslinger, University of Texas; Richard Anderson, Clark Atlanta University; and David O’Gwynn, Mississippi State University

Richard, who received his bachelor of science degree in computer science in May 2001, learned of the HPC Summer Intern Program from a college instructor. “One of my professors told me there was a big opportunity in Vicksburg,” the 26-year-old said.

The internship gives students a unique opportunity to supplement classroom activities with real-world experience, according to John West. “While at the ERDC MSRC they are integrated into the team and become a part of real MSRC projects. This gives the students a chance to work on real projects that are typical of what happens in the Department of Defense (DoD). Thus they have the opportunity to sample the range of projects undertaken in the DoD and to be a contributing member of a large supercomputing center. In this setting they are able to see the practical application of their academic training. Getting this exposure early in their careers will also give many of the interns an opportunity to adjust their academic tracks - and perhaps pursue an advanced degree. Actually, this is how I ended up at the ERDC MSRC and as a member of the HPC community. I spent a summer at ERDC in the MSRC as an undergraduate and have been hooked ever since,” he said.

The summer intern program benefits the ERDC MSRC as well. “We have an opportunity to work with talented students while they are still receiving their academic training,” John said. “These students are the next generation of workers in the information technology, scientific computing, and math and science communities. They bring a fresh perspective to MSRC efforts, and their talents yield significant benefits to the projects on which they participate during the summer.” From a broader perspective, he said that the summer intern program affords the ERDC MSRC the opportunity to be involved with the education of the next generation of the information technology and scientific community during critical points in their academic careers. “ERDC has always had a strong commitment to the community and to education, and the MSRC and the PET program are proud to be a part of that,” he said.

David O’Gwynn, 25, from Jackson, MS, appreciates the chance to work at one of the premier supercomputing centers in the Nation. “This facility is what puts Mississippi over the top as far as information technology,” the computer engineering student said. “I do think this is going to help me a lot.”

“It’s been a productive summer,” agreed Owen Eslinger, a 24-year-old Indiana native who grew up in North Carolina. He plans a doctorate in computational and applied mathematics.

Students will again have the opportunity to advance their academic careers in 2002. To apply for the HPC Summer Intern Program, contact John West at 601-634-3629 or John.E.West@erdc.usace.army.mil.



Summer interns at the ERDC MSRC were matched with mentors. Interns are, front from left, Owen Eslinger, Joyce Beal, Richard Anderson, Eddie Barnes, and David O’Gwynn. Mentors are, back from left, Dr. Wayne Mastin, Paul Adams, Dr. Dan Duffy, John West, and Dr. Nathan Prewitt

ERDC MSRC Participation at Users Group Conference 2001

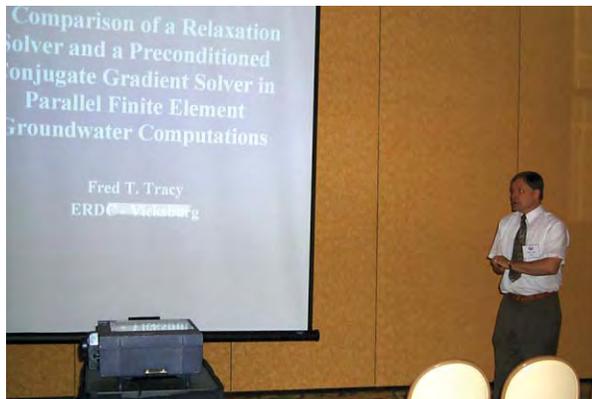
By Rose J. Dykes

The eleventh annual DoD HPCMP Users Group Conference, organized by the Shared Resource Center Advisory Panel and hosted by the Naval Oceanographic Office (NAVO) MSRC, was held at Biloxi, Mississippi, June 18-22, 2001. Eight ERDC MSRC staff members made presentations throughout the week.

Dr. Phu Luong, Environmental Quality Modeling and Simulation Onsite Lead, presented a paper comparing the grid-generation technique and domain decomposition in terms of load balance and overall performance for coastal ocean circulation modeling. He used a data set that was a 90-day simulation for the U.S. west coast under two 29-block grids, one grid generated by domain decomposition and the other by the multiblock technique.



Dr. Phu Luong



The implementation of a parallel conjugate gradient solver using an incomplete lower-upper factorization preconditioner was described by Dr. Fred Tracy for solving practical groundwater problems in his paper presentation entitled "A Comparison of a Relaxation Solver and a Preconditioned Conjugate Gradient Solver in Parallel Finite Element Groundwater Computations." He showed how a relaxation solver could be completely inadequate where point sources/sinks are needed to model an application.

Dr. Fred Tracy

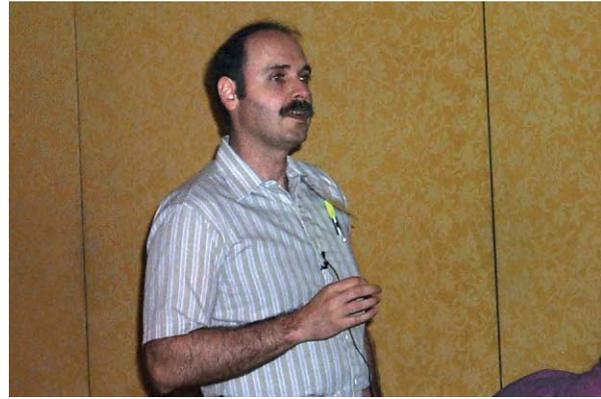
In his paper presentation entitled "Mixed-Model Parallel Implementation of Chimera Grid Assembly," Dr. Nathan Prewitt, Computational Fluid Dynamics Onsite Lead, talked about improvements in the parallel implementation of the grid assembly function within the Beggar code based on a mixed programming model, combining message passing with shared-memory programming constructs using standard POSIX calls.



Dr. Nathan Prewitt

Dr. Brian Circelli and Mr. Paul Adams, of the Scientific Visualization Center, presented a paper examining some of the issues associated with the storage, transport, and investigation of large data sets; methods for effectively processing the data; and how the difficulties associated with investigating large data sets can be overcome. They explained that by using techniques such as data mining, byte-scaling, and internal hardware texture mapping, they were able to visualize and interpret more than 6 terabytes of data.

Dr. Brian Circelli



Dr. Daniel Duffy Speaks at Plenary Session

During a plenary session, Dr. Daniel Duffy on behalf of Dr. William Ward presented the Technology Insertion 2001 benchmark project. The objective of the project was to provide the High Performance Computing Modernization Program with performance data on available high-performance computing systems as an aid to acquisition decision making. Dr. Duffy stated that this test package should be viewed as a starting point for a more comprehensive, modular, and portable test package and that the package must be revised on a yearly basis to remain up to date.

Dr. Stephen Wornom, Climate/Weather/Ocean Modeling and Simulation Onsite Lead, participated in the Poster Session of the conference illustrating timing charts for a domain decomposition version of the Jacobi iterative method applied to solving a particular Helmholtz differential equation. Pthreads were used to emulate the algorithm specifications that were originally given in Co-Array Fortran. The timings showed that thread programming could provide a portable and efficient alternative to Co-Array Fortran.



Dr. Stephen Wornom discusses his poster with Dr. Virginia Ross, High Performance Computing Center Manager, Air Force Research Laboratory/Information Directorate



Dr. Kent Eschenberg

A new file protocol under development that adds a layer upon Hierarchical Data Format to provide support for several higher level concepts including time-varying data, unstructured cells, metadata about each component, and notes created during evaluation was discussed by Dr. Kent Eschenberg in a paper presentation. One of the benefits of the new protocol will be that postprocessing tools, such as visualization packages, can more completely “understand” files without knowledge of the source of the file.

COL John W. Morris III

New Commander of ERDC

Colonel John W. Morris III assumed command of the U.S. Army Engineer Research and Development Center (ERDC) on July 11, 2001. The ERDC is the U.S. Army Corps of Engineers' distributed research and development command and consists of seven unique technical laboratories.

Prior to assuming command of ERDC, Colonel Morris was Director of Management, Office of the Chief of Staff of the Army. He is a graduate of the U.S. Military Academy at West Point and holds a master's degree in engineering from the University of Florida at Gainesville. His service schools include Airborne School, Ranger School, Infantry Officer Basic Course, Engineer Officer Advanced Course, Command General Staff College and the Army War College.

Colonel Morris visited the ERDC MSRC on September 5 while touring the ERDC Information Technology Laboratory.



COL Morris (center), new ERDC Commander, shown with Brad Comes (left), ERDC MSRC Director, and Tim Ables (right), ERDC Information Technology Laboratory Acting Director

Impact Expands for Users with New PET Program

By Dr. Wayne Mastin

The DoD HPC Modernization Program on May 25, 2001, announced a new 8-year effort in Programming Environment and Training (PET) at the Shared Resource Centers. The PET program is composed of four components, one facilitated by each of the MSRCs. A contract for the first 3 years of the PET component at the ERDC MSRC was awarded to the MOS University Consortium, comprised of Mississippi State University, the Ohio Supercomputer Center, and the San Diego Supercomputer Center. Dr. Joe Thompson, Distinguished Professor of Aerospace Engineering at Mississippi State, will serve as Program Director for PET at the three components operated by the MOS Consortium.

ERDC MSRC Lead for CSM and CFD

The new PET, which began on June 1, 2001, includes a major restructuring of the program. Each PET component is assigned functional areas and provides support in these areas for the entire DoD HPC community. The functional areas assigned to the ERDC MSRC include the DoD Computational Technology Areas (CTAs) of Computational Structural Mechanics (CSM) with Dr. Richard Weed as its lead and Computational Fluid Dynamics (CFD) with Dr. Nathan Prewitt as its lead. In addition, the ERDC MSRC PET component is responsible for the PET Online Knowledge Center (OKC) and Education, Outreach and Training Coordination (EOTC). The lead for the CFD functional area is Dr. Bharat Soni, Director of the Center for Computational Systems at Mississippi State University. Dr. Tinsley Oden, Director of the Texas Institute for Computational and Applied Mathematics at the University of Texas at Austin, is the lead for the CSM functional area. The PET OKC will provide a repository of programmatic and technical information in all functional areas. This knowledge center will have ready access to software tools and training, as well as current information on PET projects. It will allow the DoD HPC users to enter a single Web portal to navigate and search through vast amounts of information and expertise collected from the PET program's activities. Dr. Geoffrey Fox of Indiana University is leading the development and implementation of the OKC. The EOTC functional area addresses the efficient and productive delivery of training to DoD HPC users, opportunities for minority serving institutions, and visiting student and faculty programs. A major activity in this area

Consortium Partners

Central State University
Clark Atlanta University
Computer Sciences Corporation
Florida International University
Indiana University
Jackson State University
Mississippi State University
National Center for Supercomputing Applications
Ohio Supercomputer Center
Ohio State University
San Diego Supercomputer Center
Science Applications International Corporation
University of Tennessee - Knoxville
University of Texas - Austin
University of Hawaii
Wright Technology Network

is to promote careers in computational science and engineering, and thereby ensure a sufficient supply of well-trained HPC users to meet the future needs of the DoD. The lead for EOTC is Don Frederick of San Diego Supercomputer Center. Mr. Frederick will be coordinating training at the DoD sites, as well as promoting PET participation in outreach forums such as conferences, workshops, and symposia.

Expertise Remains in the Program

From the outside observer's perspective, the PET program will appear much the same. The ERDC MSRC PET staff will continue to include onsite leads in the CTAs of CFD, CSM, Climate/Weather/Ocean Modeling and Simulation, and Environmental Quality Modeling and Simulation. Mr. Bob Athow will continue to serve as the ERDC MSRC's Government PET Lead. The past Onsite Academic Lead, Dr. Wayne Mastin, will also continue in PET as the Component Point of Contact. The PET training program will continue to offer instruction on the latest HPC software, tools, and machines, while taking full advantage of emerging collaborative and distance learning technologies.

Mission Remains the Same

Although the organizational structure of the PET program has changed, its mission has not. PET is responsible for gathering and deploying the best ideas, algorithms, and software tools emerging from the national high-performance computing infrastructure into the DoD user community. However, the mission field has expanded to include not only the MSRCs, but all of the HPCMP Shared Resource Centers (MSRCs and Distributed Centers) and other DoD locations.

Converting from PBS to LSF

By Drs. Daniel Duffy, Mark Fahey, and Jeff Hensley

Over the past 2 years, users at the ERDC MSRC have become comfortable using the Portable Batch System (PBS) to submit their jobs to the batch queue. Even though PBS has provided a uniform queuing system across the MSRC, the Load Sharing Facility (LSF) provides an attractive alternative to PBS for the new Compaq systems. Furthermore, with just a few changes, users' existing PBS scripts can be converted to LSF in a short time.

Just as with PBS, users can submit batch jobs in one of two ways, either by entering information on the command line or by creating a batch script. This article focuses on the batch script approach to submitting jobs.

First, there are just a few job commands that one needs in order to control the submission, monitoring, and killing of a batch job. Table 1 shows four common PBS job control commands and their equivalent LSF commands. Therefore, instead of issuing a PBS `qsub` command to submit a batch file, the user uses the command `bsub` when using LSF. Furthermore, while monitoring batch jobs, the PBS `qstat` command and its many options are often used. The equivalent command in LSF is the `bjobs` command.

Table 1 Common PBS job control commands and their equivalent command in LSF		
PBS Command	LSF Command	Function
<code>qsub batch file</code>	<code>bsub batch file</code>	Submit the batch file to be run.
<code>qstat</code>	<code>bjobs</code>	Show the current status of jobs in the queue.
<code>qdel jobid</code>	<code>bdel jobid</code>	Delete the job with the specified jobid.
<code>qhold jobid</code>	<code>bstop jobid</code>	Suspend the job with the specified jobid.

Next, the batch script itself must be modified before it can be submitted via the `bsub` command. Table 2 shows some common PBS batch options and their equivalent LSF options. A short description of the functionality of these options is also given in the table. Perhaps the most important of these alternatives is the specification of both the number of CPUs and the wall-clock time. While in PBS the `-l` option is used for declaring both the needed number of processors and the desired wall-clock time, LSF breaks these into two separate commands. Hence, both the `-n` option followed by the requested number of processors and the `-W` option followed by a time limit in the `hh:mm:ss` format must be included. Other options are available for specifying the number of processors and their arrangement across the machine. More information may be found in the man pages. Finally, a complete PBS batch script converted to the corresponding LSF batch file is included in Table 3.

The conversion between PBS and LSF batch files for most users will be quite simple. Along with the ease of use of LSF, users will find additional commands in the LSF documentation and man pages that will aid users to better utilize the ERDC Compaq platforms.

Table 2
Equivalent PBS and LSF batch commands

PBS Option	LSF Option	Function
-N	-J	Declare the name of the job.
-l nodes=1:ppn=4	-n 4	Specify the number of CPUs for the job.
-l walltime=00:25:00	-W 00:25:00	Specify the time limit of the job.
-A project	-P project	Charge to the specified project.
-q queue	-q queue	Run the job in the specified queue.
-o outfile	-o outfile	Rename standard output to outfile.
-e errfile	-e errfile	Rename standard error to errfile.
-m b	-B	Send mail to the user after the job has begun.
-m e	-N	Send mail to the user upon completion of the job.

Table 3
A comparison between a typical PBS and LSF batch script and job control commands. Note that equivalent commands occur on the same line of the text.

PBS	LSF
#!/bin/ksh	#!/bin/ksh
# This is a typical PBS Script	# This is the equivalent LSF Script
#	#
#PBS -N jobname	#BSUB -J jobname
#PBS -l nodes=1:ppn=4	#BSUB -n 4
#PBS -l walltime=00:25:00	#BSUB -W 00:25:00
#PBS -A project	#BSUB -P project
#PBS -q primary	#BSUB -q primary
#PBS -o outfile	#BSUB -o outfile
#PBS -e errfile	#BSUB -e errfile
#PBS -m b	#BSUB -B
#PBS -m e	#BSUB -N
#	#
# Change directories to where	# Change directories to where
# the job is to be run.	# the job is to be run.
#	#
cd \$WORKDIR/mydirectory/	cd \$WORKDIR/mydirectory
#	#
# Run the MPI executable	# Run the MPI executable
#	#
mpirun -np 4 executable < input	prun -n 4 executable < input

Dr. Fred T. Tracy—An Impressive Past and an Enthusiastic Future

By Rose J. Dykes

Dr. Fred T. Tracy, a member of the ERDC Information Technology Laboratory, recently joined the MSRC Team. Dr. Tracy has written numerous papers and reports and given several presentations on a wide variety of engineering and scientific computer applications. Examples include interactive computer graphics, numerical grid generation, finite element seepage and groundwater modeling, computer-aided structural stability, and Fast Fourier Transform (FFT) algorithms for structures applications. His paper “Graphical Pre- and Post-Processor for 2-Dimensional Finite Element Method Programs,” presented at the ACM SIGGRAPH Conference, was the founding paper for the whole class of marching front grid generation algorithms. The program Seep2D that he wrote is being widely used in Government, industry, and colleges and universities, and it is in the Groundwater Modeling System (GMS).

Dr. Tracy’s enthusiasm for high-performance computing (HPC) was kindled by his successful parallelization of the 3-D finite element groundwater model FEMWATER under the Common High Performance Computing Software Support Initiative project. Some of his current interests include improved linear and nonlinear solvers, more

efficient parallelization algorithms and paradigms, and using the in-the-future development of computational grid technology and satellite data to obtain timely answers to critical situations from remote locations using HPC. He also works with various university personnel under the PET program to help facilitate the determination and dissemination of state-of-the-art methodology to the users. For instance, he has recently worked with Professor Mary Wheeler of TICAM, University of Texas.

In 1991, Dr. Tracy received the first Ph.D. in computational engineering awarded by Mississippi State University. He was awarded the U.S. Army Engineer Waterways Experiment Station (WES) Herbert D. Vogel Award for Scientist in 1989, the WES Director’s Research and Development Achievement Award in 1992, the Department of the Army Research and Development Achievement Award in 1992, the WES Employee of the Year with Disability Award in 1992, and the Department of the Army Meritorious Civilian Service Award in 1999. Fred’s goal at the MSRC is to help users solve problems of such magnitude that would otherwise not be possible without the resources of HPC. He looks forward with enthusiasm to addressing the challenges that will be presented.



*Dr. Fred Tracy
Computational Engineer at ERDC*

Scientific Visualization Workshop—Bring Your Own Data

By Paul Adams

John West, Director of Scientific Computing, and Paul Adams, Director of Scientific Visualization, initiated a Bring Your Own Data (BYOD) workshop for MSRC users. The first workshop was held June 25-26 in the ERDC MSRC Scientific Visualization Center (SVC). Drs. Joe Werne and Michael Gourlay, Colorado Research Associates, and Chris Stone, Georgia Tech, participated.

The purpose of the workshop is to highlight the capabilities of the SVC in meeting the needs of the users. In addition, by bringing together a diverse group of users, it creates a synergy, or an exchange of ideas, among the users that could not be otherwise achieved. Finally, users can meet other groups within the ERDC MSRC, such as Computational Science and Engineering (CS&E) and Programming Environment and Training (PET), and leverage these assets in their future work.

The first BYOD workshop was definitely a benefit to the users. Chris Stone, in particular was able to see phenomena in his calculations that he would not otherwise have seen. "I had no idea that there was a ring within a ring reaction," he said as he viewed isosurfaces on the Panoram screen. As he viewed the wake created by the injection of particles, he commented, "I never knew that structure existed."

Dr. Werne was also impressed with the workshop. "This was a very productive 2 days," he said, referring to the porting of the ABL code to the Origin 3000 with the help of the CS&E group.



Glen Browning, Chris Stone, and Tom Biddlecome discuss the Gas Turbine Engine Challenge Project and possible future endeavors

Members of SVC, PET, and the Gas Turbine Engine Challenge Project attended Dr. Gourlay's tutorial on the use of OGLE, a visualization tool that he developed. While OGLE has been used for the AirBorne Laser and Wake Turbulence Challenge Projects, Mr. Stone of the Gas Turbine Engines Project saw how it could benefit his project as well.

The 2-day workshop was a great example of the two-way technology transfer that can take place. Additional workshops will be conducted about every 6 months.



Attendees of the OGLE tutorial



Upgrades to the ERDC MSRC's Computational Environment

By Rebecca Fahey

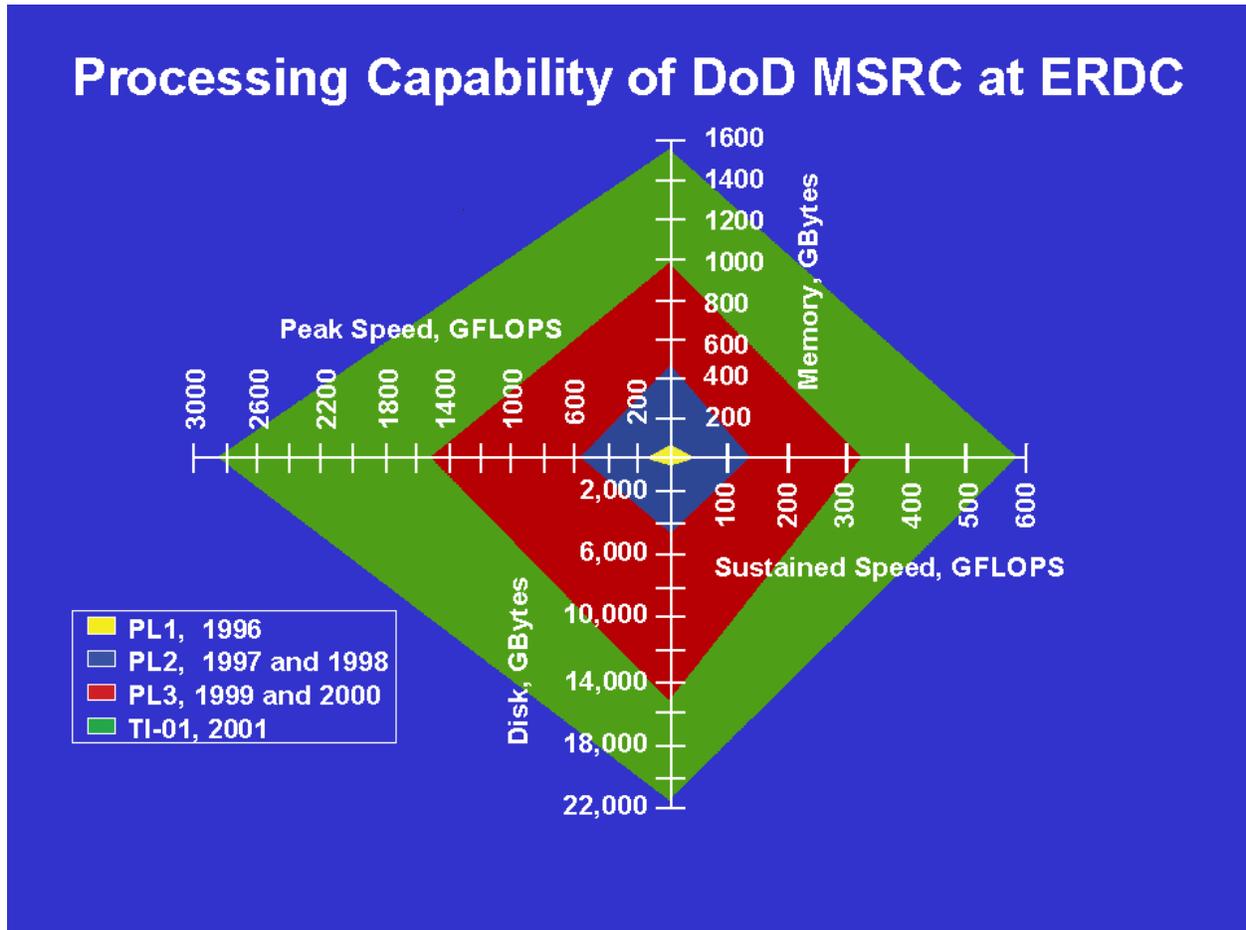
Three major upgrades to computing hardware have occurred in the recent months at the ERDC MSRC, with one additional upgrade scheduled to occur in the spring of 2002. The first upgrade was a modernization of the Mass Storage Facility (MSF) that occurred early this summer. Following that event was the addition of more processing elements and disk space to the Cray T3E that occurred in July. In October, the ERDC MSRC welcomed the first of two Compaq supercomputers to the site, with the second to follow in the spring.

The future need for long-term data storage exceeds the capability of the old Cray J90 system that has served the ERDC MSRC users for several years. That system has been modernized with a dual, highly available Sun-based system. The new mass storage archival system, dubbed "Data Management System," is comprised of two Sun E6500 servers and a 3.2 terabyte disk cache, which is a

whopping 16-fold increase over the 200 gigabyte disk cache present on the outgoing J90-based system. The upgrade also includes replacement of aging "Timberline" and "Redwood" tape technology with newer "Eagle" tape drives and cartridges. The new DMS system has been in operation since July, and the data that were stored on the old MSF are currently being migrated. The data migration is scheduled to be complete in December. At that time, the MSF system will be decommissioned.

"The benefits to DoD research and development will be enormous, enabling considerable advances in the science areas that are critical to the Nation's defense"

*Bradley Comes,
Director of the ERDC MSRC*



Processing Capability of DoD MSRC at ERDC

In July, the Cray T3E was upgraded with additional processors and disk. Since the system stays very busy, the additional processors were added to decrease the queue wait time of submitted jobs. In the computational pool, 256 Alpha EV56, 675-MHz processors with 512 MB of memory were added. The Cray T3E now has 768 processors in the computational pool, which are available for batch processing of parallel jobs. The queue structure and limits have not been changed. Thus, the maximum job size is still 512 processors. To support the additional processors, a large amount of disk was added to the system to bring the available storage to 2.7 TB. Two terabytes of the disk is used for the \$WORKDIR file system that is used as scratch space for executing jobs. This is an addition of 1.3 TB of \$WORKDIR disk space. The remaining 0.7 TB of disk is used for user home directories and system space.

In October, the ERDC MSRC welcomed the first of two Compaq systems. The new Compaq systems, which together cost more than \$11 million, will provide 1.4 teraFLOPS peak computational capacity (a teraFLOPS is 1,012 floating-point operations per second). The first Compaq system is a Compaq AlphaServer SC40, which contains 256 Alpha 833 MHz processors configured as a single system with 64, 4-way shared memory nodes that contain 4 GB of memory each.

The system also contains 4 TB of disk space. The system was available to users in October and will be followed by a Compaq AlphaServer SC45 next spring. The Compaq SC45 will contain 508 Alpha 1 GHz processors configured in a single system with 127, 4-way nodes that contain 4 GB of memory each.

As part of the transition to Compaq supercomputers, the IBM systems are being decommissioned. In the first phase of the decommissioning, which occurred in July, the IBM SMP was downsized to a 16, 8-way, node system. Thus, the original 512-processor system now consists of only 128 processors. At the end of September, the two IBM SPs, pandion and osprey, were decommissioned. In the final phase, the remaining 128-processor IBM SMP system will be decommissioned in late December.

Together, these system upgrades will boost the ERDC MSRC's aggregate peak computational capability to 2.8 TFLOPS, providing world-class capabilities to support the DoD. The accompanying chart shows the current capabilities of the ERDC MSRC and the history of the growth of ERDC's capabilities over time. With the completion of the Compaq installations, the ERDC MSRC will be capable of providing approximately 18 million hours of computational time per year to the Nation's DoD researchers.



A Technical Overview of the Compaq Supercomputer at the ERDC MSRC

By Dr. Mark Fahey

The ERDC MSRC recently installed a Compaq SC40 system. The Compaq SC40 provides a distributed/shared-memory environment where each node has its own memory and address space, and the address space is not shared between nodes. These systems are programmed using a message-passing model that involves exchanging blocks of data between cooperating processes. The shared-memory model can be exploited at the node level.

The operating system for the Compaq SC40 is Tru64 UNIX, a 64-bit advanced kernel architecture. Tru64 UNIX provides symmetric multiprocessing (SMP), real-time support, and numerous features to assist programmers in developing applications that use shared libraries, multithread support, and memory-mapped files. The Alpha

architecture is based on a 64-bit microprocessor and the Tru64 UNIX 64-bit operating system. These facts introduce a number of extended capabilities beyond 32-bit architectures. For example, 64-bit addressing allows Tru64 UNIX to naturally support file system sizes greater than 2 gigabytes.

64 bit versus 32 bit

When porting a 32-bit application to the 64-bit environment of Tru64 UNIX, one faces the same issues that would be faced in porting that application to 64-bit IRIX, the SGI Origin 3800's operating system. Most porting concerns for applications written for a 32-bit environment are caused by three facts about the 64-bit environment.





- * Pointers are 64 bit, not 32 bit.
- * Long values (long) are 64 bit, not 32 bit.
- * Integer values (int) are 32 bits and not the same size as longs and pointers, as is the case in a 32-bit environment.

In the Tru64 UNIX 64-bit environment, long and pointer data types are 64 bits. Table 1 presents the Tru64 UNIX data types.

Table 1 C Language Data Types	
Data Type	Tru64 UNIX Bits
char	8
short	16
int	32
float	32
pointer	64
long	64
long long	64
double	64
long double	128

Big Endian versus Little Endian

Computer memory is referenced by addresses that are positive integers. It is “natural” to store numbers with the least significant byte coming before the most significant byte in the computer memory; however, computer designers sometimes prefer to use a reversed order version of the representation.

The “natural” order, in which less significant binary digits come before more significant digits in memory, is called little-endian. Many vendors, such as IBM, Cray, and SGI, prefer the reverse order that is called big-endian. Big-endian and little-endian derive from Jonathan Swift’s “Gulliver’s Travels,” in which the Big Endians were a political faction that broke their eggs at the large end (“the primitive way”) and rebelled against the Lilliputian King who required his subjects (the Little Endians) to break their eggs at the small end.

On the Compaq SC40, one can compile codes with the “-convert big_endian” option to ensure that the executable can read big endian files. Table 2 lists machine type and endian status.

Table 2 Endian Type of Computer Types	
Machine	Endian
Compaq	Little
Cray	Big
IBM	Big
MAC	Big
PC	Little
SGI	Big
Sun	Big
Vax	Little

Some other useful compiler options on the Compaq SC40 are presented in Table 3.

Math Libraries

The Compaq SC40 also provides the Compaq eXtended Math Library (CXML), a set of mathematical subroutines optimized for high performance on Alpha systems. These subroutines perform numerically intensive subtasks that occur frequently in scientific computing. They can therefore be used as building blocks for the optimization of various science and engineering applications. This was formerly known as the Digital eXtended Math Library (DXML).

The CXML contains the following:

- * Basic Linear Algebra library (BLAS)
- * Linear System and Eigenproblem Solvers (LAPACK)
- * Sparse Linear System Solvers
- * Signal Processing

To link in this library at compilation, use “-lcxml” at the linking stage.

In addition, the Compaq SC40 provides Cray SciLib portability support. It does this by providing SCIPOINT, which is implementation of the Cray Research scientific numerical library, SciLib. SCIPOINT provides 64-bit single precision floating point and 64-bit integer interfaces to underlying CXML routines for Cray users porting programs to Alpha systems. SCIPOINT also provides almost all Cray Math Library and CF77 (Cray Fortran 77) Math intrinsic routines.

Table 3 Common Compiler Options	
Option	Purpose
O0, -O1, -O2, -O3, -O4, -O5	Different levels of optimization. -O4 is the default. -O5 is the most aggressive.
-qarch host or -qarch ev6	Produces object code specifically for 21264 (Ev6) processors.
-qtune host or -qtune ev6	Tunes the object code specifically for 21264 (Ev6) processors.
-qfast	Sets the following options: -align dcommons -arch host -assume no accuracy_sensitive -math_library fast -O4 -tune host -align sequenc -assume bigarrays -assume nozsize. Use with care.
-omp or -mp	Enables parallel processing using directed decomposition following the OpenMP application program interface.
-assume byterecl	Specifies for unformatted data files that the units for the OPEN statement RECL specifier (record length) value are in bytes, not long words (four-byte units).

Parallel Programming Paradigms Supported on the Compaq SC40 at the ERDC MSRC

MPI, SHMEM, and HPF programs are all run in the same manner on the Compaq SC40. That is, they are all invoked with the parallel run command prun. For example, to run the executable prog.exe with four processors, use the following command:

```
prun -n 4 prog.exe
```

MPI - Highly Optimized

AlphaServer SC systems support the distributed memory model through a highly optimized implementation of MPI that uses Elan hardware, software, and switches to provide low latency and high bandwidth communications. This hardware-accelerated communication incurs very little system overhead; the Alpha system simply writes its data to a location in memory. The Elan interconnect handles the entire communications task and uses hardware-accelerated Direct Memory Access (DMA) to interact with the host Alpha processor.

The MPI library is a standard message-passing library for parallel applications. Using MPI, parallel processes cooperate to perform their task by passing messages to each other. MPI includes point-to-point message passing and collective operations between a user-defined group of processes.

The AlphaServer SC MPI library is an optimized implementation of the MPI-1 specification and is based on MPICH Version 1.1.1 from Argonne National Laboratory (ANL) and Mississippi State University (MSU). Fortran and C interfaces are provided. The AlphaServer SC MPI library is layered on top of tagged message-passing routines that are especially designed for the AlphaServer SC Elan cards, which results in highly efficient communication between nodes connected by the AlphaServer SC Interconnect. In addition, highly efficient communication within a node occurs through direct memory transfers done with minimal data movement. A number of environment variables can be set to help optimize the performance of MPI programs.



Support is also included for a large subset of the MPI-2 I/O interface via the ROMIO package from ANL. Performance of the MPI-2 I/O operations is related to the attributes of the Parallel File System.

To compile a Fortran MPI program, it must be linked to the `fmmpi`, `mpi`, and `elan` libraries. For example,

```
f90 -O3 prog.f -o prog.exe
      -lfmpi -lmpi -lelan
```

SHMEM

The SHMEM library provides direct access, via put and get calls, to the memory of remote processes. A message-passing library, such as MPI, requires that the remote process issue a receive to complete the transmission of each message; the SHMEM library, by contrast, provides the initiating process with direct access to the target memory. The one-sided communication used by SHMEM maps well onto the DMA hardware in the Compaq AlphaServer SC network adapter. A consequence of this is that SHMEM latencies on the Compaq are even lower than MPI latencies.

The SHMEM library also includes a number of initialization and management routines. SHMEM routines provide high performance by minimizing the overhead associated with data-passing requests, maximizing bandwidth, and minimizing data latency (the time from when a process requests data to when it can use the data). By performing a direct memory-to-memory copy, SHMEM typically takes fewer steps to perform an operation than a message-passing system. It requires only one step, to either send the data or get the data. However, additional synchronization steps are almost always required when using SHMEM. For example, the programmer must ensure that the receiving process does not try to use the data before it arrives.

To compile a Fortran SHMEM program, it must be linked to the SHMEM library.

For example,

```
f90 -O3 prog.f -o prog.exe
      -lshmem
```

HPF

High Performance Fortran (HPF) 3,4 extends Fortran 90 with data distribution directives to

facilitate computations done in parallel. To compile an HPF program, it must be compiled as follows:

```
f90 -hpf -hpf_target
      [smpi,cmapi,gmpi,pse]
      prog.f -o prog.exe
```

See the `f90` man page for more information on the `hpf_target` options.

OpenMP

OpenMP is a set of shared-memory parallel directives for Fortran, C, and C++ programs. Its purpose is to standardize the parallel compiler directives developed by different vendors. In addition to loop-level parallelism, OpenMP supports parallel sections (i.e., task-level parallelism) and orphaned directives. Additional information can be found at the OpenMP Web site at <http://www.openmp.org>. An ERDC MSRC report entitled "Using OpenMP and Threaded Libraries to Parallelize Scientific Applications" can be found at http://www.wes.hpc.mil/pet/tech_reports/reports/pdf/tr_00_35.pdf. Compilation of a Fortran OpenMP code would proceed as follows:

```
f90 -omp prog.f -o prog.exe
```

To run an OpenMP code, one must first set the `OMP_NUM_THREADS` variable to the desired number of threads before running the code. Then the code can be run directly without using the `prun` command.

In summary, the Compaq SC40 is a distributed/shared memory high-performance computer. The system offers scalability, high-speed internal communications, and a mature set of compilers and libraries. If you need assistance getting started or optimizing an existing code on the Compaq SC40 at the ERDC MSRC, please do not hesitate to contact our Customer Assistance Center at 1-800-500-4722 or msrhelp@erdc.hpc.mil.

(Note: Much of the information for this article was obtained from Compaq documentation available at www.compaq.com/hpc.)



Access Grid Multimedia Conferencing System

By Mark Green and John West

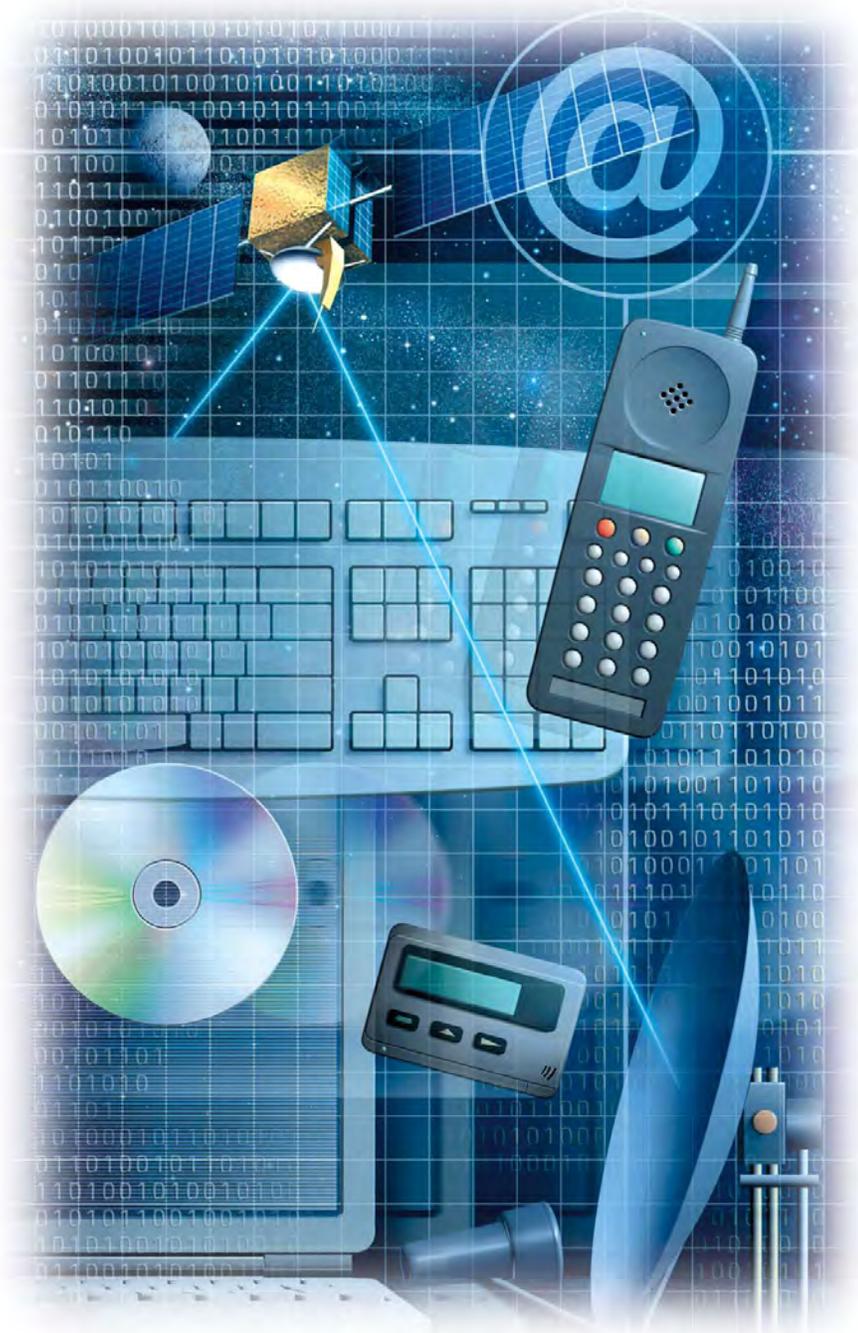
The Access Grid (AG) is a collection of hardware and software resources that enable groups located throughout the world to collaborate and interact with each other over the Internet. Teams formed by the National Computational Science Alliance (a nationwide partnership of academic, Government, and business organizations) with funding from the National Science Foundation started work on the AG in the late 1990s. Today nearly 60 AG sites are operational and participating in large-scale

distributed meetings, group work sessions, distance training classes, seminars, and lectures. AG nodes are “planned spaces” that have large-format multimedia displays and are designed for group-to-group communication rather than the individual communication achieved by the many commercially available desktop-to-desktop tools.

The ERDC MSRC is installing an AG system in the ERDC Information Technology Laboratory.

This system uses standard commercial multimedia equipment, i.e., projectors, microphones, speakers, video cameras, etc., interfaced with multiple personal computers connected to a high-speed, multicast-enabled network.

Open-source, interactive software provided by Argonne National Laboratory is used to control the system and manage the multiple video and audio streams transmitted to and received from each participating site. The ERDC system features a 20-foot-wide multimedia display with life-size projections to provide the capability for collaborating with HPC users located across the DoD. The system is scheduled to be fully operational this fall and is complemented by a similar system being installed by ERDC MSRC PET partner Jackson State University (JSU) in its new e-Center facility. In addition to providing JSU with a state-of-the-art communications infrastructure on a par with the Nation’s leading technology institutions, JSU and ERDC will team to use the Access Grid infrastructure to support leading efforts in distance training and education and to provide a rich environment for research collaborations.





Robert Athow (third from left), ERDC MSRC, and Dr. Will McMahon (fourth from left), ERDC Geotechnical and Structures Laboratory, with DoD-Army visitors, September 11



John Wolfin (left), Chesapeake Bay Field Office, U.S. Fish and Wildlife Service, and Tony Caligiuri (center), Chief of Staff for Representative Wayne T. Gilchrest, R-Maryland, August 20



Dominic Izzo (left front), Principal Deputy Assistant Secretary of the Army (Civil Works), August 16



Dr. Ingrid Padilla (left front) and Professor Ismael Pagan (right front), University of Puerto Rico, in the MSRC's Programming Environment and Training Facility, July 31



Dr. Joseph W. Westphal (center), Acting Secretary of the Army, April 18



BG Peter Madsen (far right) and Officer Professional Development Tour, South Pacific Division, U.S. Army Corps of Engineers, May 17

ERDC MSRC Publications

- 01-05 "Performance Comparison of SGI Origin 2800 and SGI 3800 on Application Codes," Jeff Hensley, Daniel Duffy, Mark Fahey, Tom Oppe, William Ward, and Robert Alter
- 01-06 "Parallel Finite Element Simulation of Wave Interacting with Ships in Motion," Shahrouze Aliabadi, Andrew Johnson, Bruce Zellars, Charlie Berger, and Jane Smith
- 01-07 "Parallel I/O for EQM Applications," David Cronk, Graham Fagg, Shirley Moore, and Victor Parr
- 01-08 "Visualization for EQM," M. Polly Baker and Alan M. Shih
- 01-09 "Developing Multi-Threaded Fortran Applications Using the PARSA Software Development Environment," Jeff Marquis and Geoffrey Wossum
- 01-10 "Development of Parallel 3-D Locally Conservative Projection Codes for Reduction of Local Mass Errors in Hydrodynamic Velocity Field Data," Mary F. Wheeler, Clint Dawson, Victor J. Parr, Eleanor W. Jenkins, and Jichun Li
- 01-11 "A Discontinuous Galerkin Discretization for the Mass Conservation Equation in CE-Qual-ICM Code," Krzysztof Bana and Mary F. Wheeler
- 01-12 "Library of Grid Interpolation Modules (INLib)," S. Gopalsamy and Bharat K. Soni
- 01-13 "A Full 2-D Parallel Implementation of CH3D-Z," Clint Dawson and Victor Parr
- 01-14 "Improvements in Parallel Chimera Grid Assembly," Nathan C. Prewitt, Davy M. Belk, and Wei Shyy
- 01-15 "CFDTool: A Web-Based Training Tool for CFD," Roy P. Koomullil and Bharat K. Soni
- 01-16 "Enhancement, Evaluation, and Application of a Coupled Wave-Current-Sediment Model for Nearshore and Tributary Plume Predictions," D.J.S. Welsh, K.W. Bedford, Y. Guo, and P. Sadayappan
- 01-17 "Metacomputing Support for the SARA3D Structural Acoustics Application," Shirley Moore, Dorian Arnold, and David Cronk
- 01-18 "Computational Science and Information Technology: Distance Education and Training," Geoffrey Fox
- 01-19 "Audio Video Conferencing," Geoffrey Fox, Gurhan Gunduz, and Ahmet Uyar
- 01-20 "Architecture and Implementation of a Collaborative Computing and Education Portal," Geoffrey Fox
- 01-21 "Ubiquitous Access for Computational Science and Education," Geoffrey Fox
- 01-22 "Improved Parallel Performance for Environmental Quality Models," Victor J. Parr and Mary F. Wheeler
- 01-23 "SPEEDES Installation and Training at ERDC MSRC," Wojtek Furmanski
- 01-24 "Adaptive Mesh Refinement in CTH: Implementation of Block-Adaptive Multi-Material Refinement and Advection Algorithms," David L. Littlefield, J. Tinsley Oden, and Graham F. Carey
- 01-25 "Review of A Priori Error Estimates for Discontinuous Galerkin Methods," S. Prudhomme, F. Pascal, J.T. Oden, and A. Romkes
- 01-26 "Comparison of Multiblock Grid and Domain Decomposition in Coastal Ocean Circulation Modeling," Phu Luong, Clay P. Breshears, and Le N. Ly.
- 01-27 "A Prototype File Protocol for Application Data Sets Based on HDF," Kent E. Eschenberg and Mike Folk
- 01-28 "STWAVE: A Case Study in Dual-Level Parallelism," Rebecca Fahey and Jane Smith

These and other publications can be accessed at

www.wes.hpc.mil

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

AG	Access Grid
AMR	Adaptive Mesh Refinement
BYOD	Bring Your Own Data
CDC	Control Data Corporation
CFD	Computational Fluid Dynamics
CHSSI	Common High Performance Scalable Software Initiative
CPU	Central Processing Unit
CS&E	Computational Science and Engineering
CSM	Computational Structural Mechanics
CTA	Computational Technology Area
DMS	Data Management System
DoD	Department of Defense
EOTC	Education, Outreach and Training Coordination
ERDC	Engineer Research and Development Center
GMS	Groundwater Modeling System
GSL	Geotechnical and Structures Laboratory
HPC	High Performance Computing
HPCMP	HPC Modernization Program
JSU	Jackson State University
LSF	Load Sharing Facility
MPI	Message-Passing Interface
MSF	Mass Storage Facility
MSI	Minority Serving Institution
MSRC	Major Shared Resource Center
MSU	Mississippi State University
NAVO	Naval Oceanographic Office
OKC	Online Knowledge Center
PBS	Portable Batch System
PET	Programming Environment and Training
SMP	Symmetric Multiprocessing
SVC	Scientific Visualization Center
TICAM	Texas Institute for Computational and Applied Mathematics
TI-02	Technology Insertion 2002
WES	Waterways Experiment Station

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 or the use of such commercial products.

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

info-hpc@wes.hpc.mil

ERDC MSRC Newsletter Editorial Staff

Interdisciplinary Coordinator

Mary L. Hampton

Outreach Coordinator

Rose J. Dykes

Chief Editors

Rose J. Dykes

Virginia "Ginny" Miller

Visual Information Specialist

Betty Watson

ERDC MSRC Customer Assistance Center

Web site: www.wes.hpc.mil

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

Telephone: 1-800-500-4722



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.



ERDC

Major Shared Resource Center

ERDC MSRC

In Touch
with Tomorrow

