

Resource


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

NEWSLETTER

In this issue:

- Modeling High-Energy Compounds
- HPCMP Users Advocacy Group
- Cpusets on the Origin 3800
- User Disk Striping on the T3E
- HPCMP Users Group Conference 2002
- Metacomputing Grid Workshop

Just prior to September 11, the section of the Pentagon hit by the plane had been renovated and fitted with blast-resistant windows designed in part by blast simulations performed on DoD HPCMP supercomputers.



ERDC MISRC

Major Shared Resource Center



From the Director's chair



Exciting material can be found in this edition of *The Resource*. More new leaders at the U.S. Army Engineer Research and Development Center (ERDC) are introduced: Dr. Jeffery Holland, new Director of the Information Technology Laboratory (ITL) and Dr. Rick Morrison, new Deputy Director, ERDC. In my early years, Jeff and I worked together in the ERDC Hydraulics Laboratory back when we were REAL users of high-performance computers. Both of us have stayed engrained in high-performance computing (HPC) over the years, so we have not drifted too far apart. I'm looking forward to working with Jeff again. Dr. Holland is a strong advocate of HPC, and I am certain he will continue to support HPC initiatives in his new position as the ITL Director. Some of you may recognize Dr. Morrison, as he served at the Army Research Laboratory (ARL) for many years before assuming a position as Director for Research and Laboratory Management, Office of the Assistant Secretary of the Army (Acquisition, Logistics and Technology). Dr. Morrison also has a strong track record of supporting HPC. With more new leadership in place, the ERDC Major Shared Resource Center (MSRC) has a bright future.

A number of very informative technical articles also can be found in this issue. A cross section of HPC applications; HPC technology; and HPC tools, tips, and techniques are included. Everyone will find something interesting and educational. Most significant for the ERDC MSRC is its recent involvement in metacomputing. It has been seriously working with this technology for a little over a year now and recently tested the first cross-MSRC computing grid. The ERDC MSRC also held the first Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) Metacomputing Grid Workshop this past February. This new computing architecture, most commonly known as "Grid Computing" or "the Grid," is the next generation computing environment. What "the Web" did for the general public, the Grid will do for the large-scale computing community. This technology is nothing less than a distributed computing environment's operating system that provides a full suite of logically connected HPC capability. Geographically distributed institutions, individuals, and resources support this capability. The resources are most commonly, but not limited to, compute and storage resources, instruments, and visual systems. Grid Computing will be used to solve large-scale scientific applications. The technology is new and the learning curve is steep. However, the trek has begun, and there is no turning back now. Look for this new computing environment to be introduced at select HPCMP Centers within the next 12 months as a testbed environment with pioneer users. This environment will continue to grow and mature with time until it becomes the status quo operating environment. This is an exciting step toward maturing the existing computing environment into a seamless "system of systems." Grid Computing technology is the means toward this end.

This edition of *The Resource* will be introduced just prior to the 2002 HPCMP Users Group Conference—the one time each year of gathering to meet, share, and learn. The ERDC MSRC is the lead MSRC to support the HPCMP's Users Advocacy Group in making sure the conference is a success. To support this mission, the MSRC will have a few extra people at the conference. The conference will be an opportunity to meet some of the staff and to get to know them. They are a great bunch of individuals who are truly committed to providing the best HPC experience possible for the DoD HPCMP user.

A handwritten signature in blue ink that reads "Bradley M. Comes".

Bradley M. Comes
Director, ERDC MSRC

About the Cover:


Technology in Action. The picture of the Pentagon is prior to September 11. Blast simulations performed on DoD HPCMP supercomputers helped to save lives. ERDC MSRC supercomputer was featured on CBS "60 Minutes II" (see article, page 4).

Group	MPer	CPUs	MBPerSec
0	500	100	100
1	1000	200	200
2	1500	300	300
3	2000	400	400

p.22




p.24



p.10



p.11



p.14



p.20

Features

ERDC MSRC Supercomputer Featured on "60 Minutes II"	4
Users Advocacy Group – Ombudsman for the User	8
Programming Environment and Training On-line Knowledge Center	9
High-Energy Compounds Have Potential as Effective Rocket Fuels	10
Welcome to the 2002 Users Group Conference	11
University of Texas at Austin Has Major Role in Pet Program	12
Active Control of Combustion Dynamics in Gas Turbine Engines	14
Supercomputing 2001 Conference – “Beyond Boundaries”	20
Metacomputing Grid Workshop	28
Exploring the Future of High-Performance Computing via the Access Grid Node	32

Departments

announcements	2
accomplishments	6
word search	7
interview with Dr. Michael Stephens	18
technology update	22
off-campus	31
visitors	33
tool tips	35
community outreach	35
acronyms	36
ERDC MSRC publications	36
training schedule	36



p.2



p.4



p.12

Dr. Jeffery P. Holland, New ERDC Information Technology Laboratory Director

By Rose J. Dykes

Dr. Jeffery P. Holland became Director of the ERDC Information Technology Laboratory (ITL) in December 2001. During his 22 years at ERDC, he has served as a first and second line supervisor, program manager, and technical director.

Dr. Holland received a bachelor's degree (with honors) in environmental engineering from Western Kentucky University; a masters degree in environmental and water resources engineering from Vanderbilt University; and a doctorate in civil engineering from Colorado State University. Among the numerous awards that he has received during his career are the Army R&D Award and the Federal Laboratory Consortium Award for Excellence in Technology Transfer. He has authored or coauthored over 100 publications.

The ERDC MSRC is fortunate to have Dr. Holland as the new ITL Director. He has been a long-time advocate of HPC technologies. He serves as the DoD Computational Technology Area Leader for Environmental Quality Modeling and Simulation and the ERDC High-Performance Computing Coordinator and views HPC as a key component necessary to cost-effectively meet the future challenges in DoD science and technology and test and evaluation.



Dr. Jeffery P. Holland



Dr. Holland (center) and his wife, Janet (left), shown with Timothy Ables (right), new Assistant to the ERDC Director, at the ITL reception held for Dr. Holland

Dr. Holland (center), Bob Athow (left), ERDC MSRC PET Technical Advisor, and Dr. Leslie Perkins, HPCMP, at the DoD HPC Users Group 2001 Conference, Biloxi, MS, June 2001



ERDC New Deputy Director Views High-Performance Computing as a Cornerstone for Future Capabilities

By Rose J. Dykes

Dr. Walter F. “Rick” Morrison assumed the position of Deputy Director of ERDC in October 2001. Dr. Morrison says “I consider high-performance computing to be a cornerstone of our future capabilities in engineering and science.”

As the ERDC Deputy Director, Dr. Morrison is a member of the leadership team for one of the most diverse research organizations in the world. ERDC comprises seven laboratories located at four geographical sites. With over 2,000 employees, 1,000 of whom are engineers and scientists, ERDC boasts \$1.2 billion in facilities and an annual program exceeding \$550 million.

Prior to accepting the ERDC position, Dr. Morrison served as Director for Research and Laboratory Management, Office of the Assistant Secretary of the Army (Acquisition, Logistics and Technology). His responsibilities included the Army Basic Research, Applied Research programs for the Army Research Laboratory, Army Research Institute, Corps of Engineers, and Simulation, Training and Instrumentation Command, as well as several Army-wide programs including Environment Quality Technology, Manufacturing Technology, and Army High Performance Computing. Overseeing laboratory management policy for all Army laboratories and research centers was another one of his assignments.

Dr. Morrison received his bachelor’s, master’s, and doctorate degrees in physics from the Georgia Institute of Technology. He is an R.O.T.C. Distinguished Military Graduate and served 4 years on active duty with the U.S. Army. His commitment to HPC is a great asset for the ERDC MSRC.



Dr. Rick Morrison



ERDC MSRC Supercomputer Featured on “60 Minutes II”

By Ginny Miller

Work performed on supercomputers at the ERDC MSRC was featured last fall on a broadcast of the CBS weekly news magazine “60 Minutes II.” The program, a spinoff of “60 Minutes,” aired November 28, 2001.

For the episode, entitled “Miracle of the Pentagon,” CBS correspondent Scott Pelley reported on life-saving building modifications at the Pentagon, where American Airlines Flight 77, en route to Los Angeles when it was hijacked by terrorists, hit the building on September 11.

“Even though 125 people were killed in the Pentagon on September 11, there was something miraculous about that day,” Mr. Pelley said in his report. “The plane obliterated the first and part of the second floor, but the third, fourth, and fifth floors remained suspended in midair for 35 minutes. Hundreds of people escaped. How is that possible? The answer lies in a partially classified Government study of the bombings that have come before.”

Hundreds of people escaped. How is that possible?

The answer lies in a partially classified Government study of the bombings that have come before.

To design the Pentagon’s new protection, the U.S. Army Corps of Engineers studied the 1983 Marine barracks bombing in Lebanon and the 1995 bombing of the Alfred P. Murrah Building in Oklahoma City, as well as terrorist attacks on the Khobar Towers barracks in Saudi Arabia and the U.S. embassies in east Africa.

“At Khobar Towers, for example, most of the damage and casualties were caused by flying debris from the structure and the glass, etc.,” Lt. Gen. Robert Flowers, Commander of the Corps of Engineers, said in the “60 Minutes II” report. “And so based on that, we designed things to prevent flying debris and flying glass. At Oklahoma City, the bulk of the casualties was caused by the collapsing structure. So one of the things we studied was how to put redundant capability in a structure to prevent it from collapsing if it was attacked. So by applying the lessons that you learn from doing those studies, you can better protect structures in the future.”




The “60 Minutes II” report identified the ERDC MSRC as “one of the most powerful (high-performance computing centers) in the Nation”

One way the Corps of Engineers has attempted to make safer, more bomb-resistant buildings is by subjecting its new designs and retrofits to the detonation of its own bombs. Leading the work is Dr. Reed Mosher of the ERDC Geotechnical and Structures Laboratory, who used the high-performance computers provided by the High Performance Computing Modernization Program (HPCMP) to simulate blast effects from the bombs used in their experiments. The “60 Minutes II” report identified the ERDC MSRC as “one of the most powerful (high-performance computing centers) in the Nation.”

“The computer can test various kinds of bombs against different buildings without breaking any glass,” said Mr. Pelley, whose piece featured footage of supercomputers located at the ERDC MSRC. A portion of the program was also taped in the Information Technology Laboratory’s collaboratorium, a special three-dimensional imaging room where Dr. Mosher showed Mr. Pelley how the supercomputer re-created the blast wave that hit Khobar Towers, predicting the path of shards of glass from a breaking window.

That kind of information was critical to the renovation of the 60-year-old Pentagon. “We made several modifications to the building as part of that renovation that we think helped save people’s lives (on September 11),”



Mr. Lee Evey, in charge of the billion-dollar Pentagon renovation project, said in his interview with Mr. Pelley.

Just prior to September 11, the section of the Pentagon hit by the plane had been renovated and fitted with blast-resistant windows designed in part by blast simulations performed on DoD HPCMP supercomputers.

Just prior to September 11, the section of the Pentagon hit by the plane had been renovated and fitted with blast-resistant windows designed in part by blast simulations performed on DoD HPCMP supercomputers. Neither the heat nor the force of the exploding jet fuel was able to shatter the new windows, reducing the number of deadly shards of glass.

The renovation project also strengthened the building's limestone exterior with a web of steel beams and columns. When the Boeing 757 hit the Pentagon at 350 miles an hour, the limestone layer shattered. A shield of geofabric that lined the entire section of the building caught the shards of stone.

Together, the building modifications held up the third, fourth, and fifth floors of the Pentagon for 35 minutes after the 9:45 a.m. attack. Only after the evacuation did the heat melt the new steel away, and those who survived the attack are convinced that without the reconstruction, the floors might have collapsed immediately.

“It’s a testament to the work that the people in the renovation did and the engineers,” said John Yates, a civilian security manager in the Pentagon who was badly burned in the attack. “If it hadn’t been done, if there’d been no structural hardening, I can’t imagine what the death and destruction would be. It would have been more catastrophic than what it is. Ten times, 100 times worse.”

“It’s a testament to the work that the people in the renovation did and the engineers. If it hadn’t been done, if there’d been no structural hardening, I can’t imagine what the death and destruction would be.”

John Mauldin Named Distinguished Fellow at Mississippi State University

By Ginny Miller

The key to John Mauldin's success is the understanding that challenging problems are not always solved on the first day. "Early in my career I looked at other engineers and realized they did not have the answer when they started, but they worked at it every day," John said. "I learned that there's no problem that can't be solved given enough time."

That commitment has paid off for John, who serves as the ERDC MSRC's Deputy Program Manager and business office manager. In February, he was inducted as a Distinguished Fellow of the College of Engineering at Mississippi State University (MSU).

"This came as a complete surprise," said John, 45, one of 10 College of Engineering alumni selected for the honor. "I am really honored."

John, a 1978 MSU graduate, holds a bachelor of science in electrical engineering. In 1984, he received a master of science in computer science from MSU's College of Engineering.

John's 23-year career is varied. "I started out as a line engineer at RCA," he said, working with the Missile and Surface Radar Division on the U.S. Navy's AEGIS weapons system. He has also worked as an electronic engineer at the National Security Agency and began working with Nichols Research Corporation (now Computer Sciences Corporation (CSC)) as a systems engineer in 1987. After a 2½-year absence, during which time he worked as an independent consultant, Mauldin returned to CSC. He moved to the Vicksburg office, which handles the ERDC MSRC's high-performance computing integration effort, in September 2000.

The praise humbles Mauldin, who is still trying to get used to his new title of Distinguished Fellow. "I don't have patents and I don't have great engineering breakthroughs," he said. "Really and truly, the big problem that I solved wasn't even at work."

While living in Huntsville, AL, Mauldin was a member of the Madison County chapter of Habitat for Humanity, a nonprofit, nondenominational Christian housing organization. By applying good engineering principles and logical thinking, "We literally took an organization that was dead – we had less than \$1,000 in the bank and two unfinished houses – and on that foundation, we re-engineered an organization that has now completed more than 75 houses," Mauldin said.

Now a resident of Clinton, MS, Mauldin continues to serve his community. He is a member of the Clinton Parks and Recreation Commission and the Clinton Soccer Board, and also serves as an industry advisor to MSU's Department of Computer Science.



John Mauldin

ERDC MSRC Team Member Is Graduate of Leadership Development Program

By Rose J. Dykes

Greg Rottman, Government lead for HPC systems operation at the ERDC MSRC, has recently completed the Army's 3-year premier leadership development program for engineers and scientists—one of only two people from ERDC with this accomplishment. Greg is also the project manager for the High Performance Computing Modernization Program's (HPCMP) metacomputing initiative and serves as a member of the HPCMP metacomputing working group.

In 2000, Greg received the Commander's Award for Civilian Service and the Commander's Award for Excellence in Community Services in 1999. In addition to his busy position at the ERDC MSRC, he is also a member of the Diocese of Jackson School Board and serves on the Development Committee at Vicksburg Catholic Schools. Greg also serves as a trustee for the Vicksburg Knights of Columbus in which his leadership as the Grand Knight resulted in their receiving the highest award in the State of Mississippi for the 2 years he served.

The ERDC MSRC feels fortunate to have Greg as a member of its team. With the leadership ability that he possesses, Greg will help to ensure that the mission of the MSRC, Army, and DoD is accomplished.




Greg Rottman

Courtesy of Pickett Photography

Word Search – Metacomputing Terminology

- | | |
|------------------|---------------|
| ACCESS | RESOURCES |
| ALGORITHMS | SCIVIS |
| ARL | SERVICE |
| ASC | SIMULATIONS |
| CENTER | STORAGE |
| COLLABORATION | SUPERCOMPUTER |
| COMPUTE | TRAINING |
| COORDINATION | USERS |
| DATA | |
| DATA SETS | |
| DOD | |
| EDUCATION | |
| ERDC | |
| GIGABYTE | |
| GLOBUS | |
| GRIDS | |
| HIGH PERFORMANCE | |
| HPCMP | |
| KERBEROS | |
| MEGABYTE | |
| MSRC | |
| NAVO | |
| PARALLEL | |
| PORTALS | |
| PROCESSOR | |
| PROGRAMS | |



answers on page 27

Users Advocacy Group – Ombudsman for the User

By David Stinson

When high-performance computing users need their requirements represented to the HPCMP, where do they turn for help?

One possible source is the Users Advocacy Group (UAG). Although this group may not sound familiar, it has been around since the early days of the Program under various other names, the most recent being the Shared Resource Center Advisory Panel (SRCAP).

In the beginning, the focus of this group was to represent users in deciding what software should be purchased for the Army's Cray YMP, located at the Waterways Experiment Station in Vicksburg, Mississippi. This was one of the first supercomputers available to Department of Defense (DoD) users through the HPC Modernization Program. In 1996, four Major Shared Resource Centers (MSRCs) were established within the DoD. These four centers were located at the U.S. Army Engineer Waterways Experiment Station in Vicksburg, Mississippi (now called the Engineer Research and Development Center or ERDC), the Aeronautical Systems Center (ASC) at Wright-Patterson Air Force Base in Dayton, Ohio, the Naval Oceanographic Office (NAVO) located at the Stennis Space Center near Bay St. Louis, Mississippi, and the Army Research Laboratory (ARL) located at Aberdeen Proving Ground in Maryland. Other HPC resources have been added to the HPCMP through various Distributed Centers (DCs) around the country. The MSRCs and the DCs are collectively referred to as the Shared Resource Centers (SRCs). As the HPC Modernization Program has grown and computer architectures have changed, the needs of the user have changed as well.

In December 2001, the Shared Resource Center Advisory Panel charter and name were changed to reflect the changing mission of the group.

In December 2001, the SRCAP charter and name were changed to reflect the changing mission of the group. The mission of the new Users Advocacy Group is to provide a forum for users of the DoD HPCMP's resources to influence policies and practices of the Program, to facilitate the exchange of information between the user community and the HPCMP, to serve as an advocacy group for all HPCMP users, and to

advise the DoD High Performance Computing Modernization Office (HPCMO) on policy and operational matters related to the HPCMP. The UAG provides a forum for the following:

- ✦ Articulating user needs and issues to the Shared Resource Centers.
- ✦ Sharing of programs and techniques to exploit HPC technology.
- ✦ Advising on system hardware and software configurations.
- ✦ Advising Program participants on workload forecasting as well as the process for, composition of, and timing of system software and hardware upgrades.
- ✦ Assessing the quality and value of service provided by the SRCs.
- ✦ Recommending standard tools, techniques, or software used to develop, enhance, or maintain HPC applications.
- ✦ Providing recommendations on policy and operational issues involving other components of the program, such as networking, software applications and support, requirements, and allocations.

In addition, the UAG has a specific responsibility for the following:

- ✦ Providing a forum for discussion and resolution of issues arising between the user community and program participants.
- ✦ Soliciting information from the user community on specific topics of interest to the user community, Program participants, or the HPCMO; compiling this information for use by the HPC Advisory Panel or HPCMO, as appropriate.
- ✦ Providing recommendations on policies and procedures, hardware, software, and placement of equipment at the Shared Resource Centers and the HPCMO.
- ✦ Providing user community perspective on HPC strategies and policies to the HPCMO, including serving on other HPCMP advisory panels.
- ✦ Working with the HPCMO and other Program participants in planning the annual DoD HPCMP Users Group Conference. The UAG will have primary responsibility for planning and organizing the technical program.
- ✦ Responding to actions and issues as requested by the HPCMO.

↳ Reporting results of each of its meetings to the HPCMO within 2 weeks after that meeting. This report will provide recommendations to the HPCMO on issues or opportunities identified during the meeting.

As the users' needs have changed, the HPCMP has grown and the computer architectures deployed have changed.

Users are represented by their service members of the UAG. Each service (Army, Navy, Air Force) appoints four members with one additional member selected to


represent other DoD agencies. These members come from Science and Technology as well as Testing and Evaluation communities and serve 2-year renewable terms. The services are encouraged to appoint members who are active users in the HPCMP to best represent the user community. The UAG chairperson is selected by the UAG membership after approval by the Director of the HPCMP for a 2-year, nonconsecutively renewable term. Meetings are scheduled at least twice a year. Representatives from the SRCs, although not members, have a standing invitation to attend. To find out who their service representatives are, users may contact the HPCMO.

Programming Environment and Training On-line Knowledge Center

By Bob Athow

Finding information on current and past activities within the Programming Environment and Training (PET) program will become much easier and faster with the advent of the On-line Knowledge Center (OKC). The summer of 2002 will witness the unveiling of the OKC so that all DoD HPC users will have at their fingertips current, accurate, and robust tools and information. More than a Web portal, the OKC will have training scheduling and registration, and a complete archive of the past 6 years of the articles, reports, and HPC tools developed by the PET team during that term.

The creation of the OKC is the responsibility of Indiana University's PET OKC point of contact, Dr. Geoffrey Fox. Dr. Fox's university team has developed the framework for the OKC. Indiana University then handed that framework to the ERDC MSRC operational team for inclusion into the operational PET OKC. The fully operational OKC will include customer assistance, interactive Web pages, literature searches, DoD computational scientist POCs' training offerings, highlights of ongoing projects, and much more. In the future, if you are looking for training in the areas that PET has provided in the past, or for a tool to enable your code to be more productive, check into the PET OKC.



High-Energy Compounds Have Potential as Effective Rocket Fuels

By Drs. Mark S. Gordon, Graham D. Fletcher, and Thomas C. Oppe

A major effort funded by the Air Force Office of Scientific Research (AFOSR) is the development of new high-energy compounds that have the potential to serve as effective fuels for rockets. These compounds are referred to as high-energy density materials, or HEDM. Rob Schmitt and Jeff Botaro of Stanford Research Institute (SRI) have proposed an intriguing series of potential high-energy compounds that are being explored computationally. The procedure is to use a modified version of the G2/G3 computational models to predict the heats of formation of these compounds and then to use these predictions to assess the specific impulse, I_{sp} . The specific impulse is directly related to the energy released when a fuel is burned, and a good measure of this energy is a high-positive heat of formation.

The procedure is to use a modified version of the G2/G3 computational models to predict the heats of formation of these compounds and then to use these predictions to assess the specific impulse, I_{sp} .

The G2/G3 methods require geometry optimizations using second order perturbation theory with a reliable basis set, followed by higher level calculations at these geometries, and then followed by higher order terms that are designed to correct for any remaining errors caused by the level of theory used. Because of the large memory and operation count requirements of these calculations, the computer simulations have been performed on the ERDC MSRC Cray T3E using GAMESS. A typical modeling run requires 256 processors and 64 GB memory and may run for 48 hours. This work is supported by a grant of computer time under the auspices of the HPCMP Challenge Project program.

The first of these species to be examined is shown in Figure 1 and denoted as Compound 1.

Using a combination of isodesmic reactions and the G2 model, the heat of formation for Compound 1 is predicted to be 456.8 kcal/mol. This translates to an I_{sp} of 240 seconds, compared with 230 seconds for hydrazine. Since hydrazine is a popular monopropellant, Compound 1 appears to be a very promising fuel. However, it is always important to consider the stability of such high-energy species to various reactions before asserting their viability as fuels. One possible

reaction is the loss of molecular nitrogen from the center of Compound 1 to form the smaller species Compound 2, shown in Figure 2.


Using the highly scalable second order perturbation theory code with a large basis set on the T3E, this decomposition process was found to be exothermic by ~60 kcal/mol. It may be that the barrier for decomposition is large, in which case Compound 1 may still be a viable HEDM candidate. The search for the transition state for this reaction will be initiated shortly.

Additional studies will be needed in which the stability of potential HEDM compounds must be assessed with respect not only to various unimolecular decompositions but also to bimolecular interactions and attack by environmental species such as water.

Compound 1 could also break an N-N single bond at either or both ends to form a mono- or di-azide, Compound 3.



Calculations performed on the T3E suggest that Compound 3 is considerably higher in energy than Compound 1, an unexpected result indicating that it also is a potential HEDM. Jeff Botaro has suggested that Compound 4 (shown in Figure 3), an isomer of Compound 1, may be more viable because it cannot undergo the analogous N-N bond cleavage to form azides. Calculations on the T3E suggest that it will have a heat of formation and I_{sp} comparable with those for Compound 1. Calculations are proceeding on these and related high nitrogen-content species.




Welcome to the 2002 Users Group Conference

By David Stinson

This year the Users Advocacy Group (UAG), the HPCMO, the ERDC MSRC, and the University of Texas at Austin have teamed to bring the 12th Annual DoD Users Group Conference to Austin, Texas, during the week of June 10-13. This year's conference is different from previous conferences in that it is only open to U.S. Government agencies and their contractors. Steve Scherr, the Master of Ceremonies and outgoing Chairman of the UAG (formerly SRCAP) will bid us farewell and introduce the new UAG Chairman, Steve Finn. Looks like a precedent is being set here, as this is the third Steve to head this group (there is no running out of "Steve's" yet).

Monday arrivals will have the opportunity of attending one or two of the seven half-day/full day tutorials being offered. However, the main part of the conference will start with the Tuesday morning plenary session and will include the usual assortment of keynote speakers, invited speakers, a Director's report, and a panel discussion. The afternoon sessions will include Common High Performance Scalable Software Initiative (CHSSI) and Challenge papers, and Birds-of-a-Feather sessions. Tuesday evening has been set aside for the poster session. The traditional social event will take place on Wednesday night at the Texas State History Museum. Guests will have the opportunity to view the newly released IMAX Theater three-dimensional presentation of *Space Shuttle* and will have access to all three floors of exhibits and to multiple showings of the special effects *Texas Spirit Theater*.

While in Austin, attendees might want to go downtown and visit the Congress Street Bridge and see the largest urban colony of Mexican Free-Tail Bats in North America that migrate from Mexico each spring. Austin is the state capital of Texas and is also home to the Lyndon Baines Johnson Presidential Library. About 60 museums and galleries are in and around Austin. Riverboat tours, quarter horse racing, the zoo, the University of Texas, and the Lady Bird Johnson Wildflower Center are also part of the Austin scene. On Thursday night, the DoubleTree Hotel will provide bus transportation to the Warehouse District on 6th street. Conference attendees will have an opportunity to experience the "Live Music Capital of the World" as well as some fine restaurants in this historic district. Here is hoping all will enjoy the conference and take advantage of the many things that Austin has to offer.



University of Texas at Austin has Major Role in PET Program

By Dr. Wayne Mastin

The University of Texas at Austin has been a Member of the PET university team since the PET program began in March 1996. From that date until the present, Texas has led the Environmental Quality Modeling and Simulation (EQM) effort under the direction of Dr. Mary F. Wheeler, Director of the Center for Subsurface Modeling. Dr. Wheeler's team also includes Drs. Tinsley Oden and Jay Boisseau.

Dr. Wheeler's team has been engaged in a broad range of activities that have advanced the state of modeling and simulation of the environment.

Dr. Wheeler and her associates at Texas have been instrumental in the parallelization and improvement of codes for circulation and contaminant transport in rivers, bays, and estuaries. Much of this work has been done in collaboration with users in the ERDC Environmental and Coastal and Hydraulics laboratories. With the addition of Dr. Phu Luong as Onsite EQM Lead in 1999, Texas EQM activities expanded to include ocean modeling.

These activities have included basic parallelization techniques for porting codes to HPC machines, new

algorithms to improve the accuracy and efficiency of codes, and projection techniques for coupling circulation and transport models. Recent efforts have also been directed to solving multiphysics problems, where the simulation more accurately models the real-world conditions by simulating several physical processes in a tightly coupled computational model.

The scope of Texas activities in PET broadened to include Computational Structural Mechanics (CSM) in 1998, when Dr. Oden, Director of the Texas Institute for Computational and Applied Mathematics (TICAM), assumed the role as CSM Academic Lead.

Under Dr. Oden's leadership, the Texas CSM team has enhanced the CTH shock physics code to meet the needs of DoD researchers. Team members have added an adaptive grid capability to improve the efficiency of the code in solving very large-scale problems, and have also added features such as a solid-wall boundary condition that has been available in other CSM codes.



HPCMP users can expect a continuing close relationship with the University of Texas at Austin.

In conjunction with the adaptive grid efforts, major advances have been made in the development of reliable error estimators and techniques for grid refinement and coarsening. While efforts have been targeted at CTH, a major applications code in the ERDC Geotechnical and Structures Laboratory and the Army Research Laboratory, the results have been applicable to other codes such as EPIC, a penetration mechanics code. The effectiveness of the Texas CSM team is attributed not only to a skilled and experienced at-university team but also to the presence of an Onsite CSM Lead, Dr. Rick Weed, who aids in porting the new codes to the ERDC MSRC HPC machines and assists users with the new features.

The synergy between the EQM and CSM teams has contributed to significant advances in the development and application of grid partitioning and discontinuous Galerkin algorithms. Both of these algorithms are used to solve problems arising in EQM and CSM, as well as other areas such as computation fluid dynamics (CFD).

The PET program was restructured in 2001, and with that came additional support from the University of Texas at Austin. Dr. Boisseau moved to Texas as Director of the Texas Advanced Computing Center and is now leading the Climate/Weather/Ocean Modeling and Simulation (CWO) effort for PET. Dr. Boisseau

also serves as the Chair of the Technology Council for the Mississippi State University (MSU)/Ohio State University (MOS) consortium of which MSU is the lead. MOS serves as the lead for the ERDC, ASC, and NAVO MSRC components of the PET program. The Technology Council is a group of leading academics on the PET team that provide advice and direction for the program. Along with all of his other duties, Dr. Boisseau leads the Tiger Teams for PET. The Tiger Teams, a unique part of the program, consist of a select group that is formed for attacking a specific problem of high importance to the DoD researcher.

Dr. Boisseau leads the Tiger Teams for PET. The Tiger Teams, a unique part of the program, consist of a select group that is formed for attacking a specific problem of high importance to the DoD researcher.

Many users at the ERDC MSRC know Drs. Oden, Wheeler, and Boisseau. They have been frequent visitors, giving seminars, workshops, and PET training courses to DoD scientists and engineers. HPCMP users can expect a continuing close relationship with the University of Texas at Austin.

The ERDC MSRC would like to thank Drs. Oden, Wheeler, Boisseau, and others at the University of Texas at Austin for planning, organizing, and co-hosting the 12th Annual DoD HPCMP Users Group Conference.



Drs. Tinsley Oden (far left), Graham Carey (second from left), and David Littlefield (right foreground), participate in discussions

Active Control of Combustion Dynamics in Gas Turbine Engines

By Dr. Suresh Menon, Chris Stone, and Vaidyanathan Sankaran

For next-generation combustion systems, stringent emission constraints on gas (e.g., CO, NO, and unburned hydrocarbons) and solid particulates (e.g., soot) pollutants are likely to become major design criteria. For land-based power generation gas turbines, it appears that lean premixed combustion may be able to meet most, if not all, of these emission constraints. However, design and deployment of such systems will have to deal with the consequences of lean combustion, including flame instability.

The application for the U.S. Army is for the gas turbines in battle tanks and helicopters... [and] for gas turbines used in long-range bombers and freight transporters used by all Department of Defense components.

The application for the U.S. Army is for the gas turbines in battle tanks and helicopters. This work is also valuable for gas turbines used in long-range bombers and freight transporters used by all Department of Defense (DoD) components.

Flame instability manifests itself when the low-frequency pressure (acoustic) oscillation in the combustor suddenly grows into large-amplitude nonlinear waves that can cause flame extinction and, under certain conditions, even structural damage. Combustion instability can also occur in spray combustion gas turbines that are used in nearly all propulsion systems (aircraft, helicopters, tanks, submarines, etc.). Lean combustion systems are particularly sensitive to small perturbation in heat release that can result in flame extinction. Therefore, in order to avoid this phenomenon in production gas turbines (both power and propulsion versions), combustion always occurs far away from the lean limit by design. On the other hand, if stable combustion can be achieved in the lean limit, then not only can emission be reduced drastically, but it may also result in significant reduction in fuel consumption. This goal is now of particular relevance because of the increasing environmental concerns of fossil fuel combustion.

The physics of combustion instability has been well-known for a long time. Combustion instability is the unsteady release of heat (caused by combustion) in-phase with the acoustic oscillation (often called the

Rayleigh Criterion since it was first observed by Lord Rayleigh in the 19th century). This in-phase heat release adds energy to the oscillation, resulting in the instability. Because of its dynamic nature, combustion instability is nearly impossible to control by passive means. Therefore, active control strategies are being explored to suppress this instability. Current experimental and numerical efforts are focusing primarily on active control of fuel flow in the combustor to stabilize combustion in the lean limit. This approach controls combustion dynamics and instability by controlling the phase between the unsteady heat release and the acoustic mode in the combustor. This is accomplished by explicit manipulation of where, how, and when the fuel is introduced into the combustor.

Active control of swirl, another approach to achieving active control of combustion dynamics, has received little attention in the past. All gas turbine engines employ swirl to increase fuel-air mixing and to stabilize the flame in the combustor. Thus, active control of swirl could be another avenue to suppress combustion instability. Swirl manipulation may be practically feasible, and perhaps safer, than fuel injection control since control actuation occurs in the nonreacting flow in the inlet. Swirl manipulation could be achieved by the controlled injection of air into the inlet swirl vane to either directly change the velocity direction or to actively change the shape of the swirl vanes by creating a modified boundary layer.


Numerical prediction of highly swirling turbulent reacting flows is very difficult since conventional time-averaged methods (used in numerous commercial codes) cannot accurately predict these flows. Fortunately, a technique based on large-eddy simulations (LES) has the ability to capture the physics of these flows. In LES, all scales of motion larger than the grid resolution are resolved by the numerical scheme (that is, of high temporal and spatial accuracy), and the scales of motion smaller than the grid size are represented using subgrid models. The accuracy of the LES depends not only on the resolution of the resolved scales but also on the models used to represent the unresolved small scales. At Georgia Tech, advanced subgrid models of LES were developed that allow these complex flows using relatively coarse grids to be simulated.

LES has been used to study premixed combustion in a General Electric (GE) LM 6000 combustor and spray


combustion in a GE DACRS combustor. All conditions simulated represent realistic values for full-scale devices. For premixed combustion, when swirl is below a critical value, the flow entering the combustor behaves more like a free jet and forms circular vortex rings, similar to smoke rings, which undergo a twisting shear motion that leads to their eventual breakdown. The vortex shedding also forces the flame to pulsate in-phase. Figures 1a and 2a show these features. When swirl is increased beyond a critical value, the high swirling motion of the fuel-air mixture and the rapid expansion at the entrance of the combustor causes an adverse axial pressure gradient along the centerline (in contrast, a favorable pressure gradient exists in the low-swirl case). This pressure gradient causes rapid breakdown of the vortex rings, slows down the axial motion, and forms a recirculating bubble near the centerline (Figure 2a). The coupling between the vortex motion and flame is broken, and the flame is pushed radially outward and also upstream, finally stabilizing very close to the inlet (Figure 2b). For a fixed incoming fuel-air mixture (fixed equivalence ratio), increase in

swirl, therefore, stabilizes the combustion process (Figure 3a). This pressure amplitude reduction is similar to what has been experimentally observed. Decreasing the equivalence ratio (i.e., in lean mixture) for a fixed-swirl number increases the pressure oscillation amplitude (Figure 3b). This phenomenon is a precursor to combustion instability. The current effort is now on active control of the instability, whereby the pressure signal is recorded and analyzed on-line and then used to change the incoming swirl to stabilize combustion when the fuel-air mixture is made leaner.



For spray combustion, the gas-phase LES model is coupled to a Lagrangian droplet-tracking model in which droplet groups are tracked explicitly in the flow field. The droplet field is fully coupled to the gas field and both gas and liquid motions affect each other. As droplets vaporize, the gas (fuel) species appear as a source term in the fuel species equation. The spray LES code is being used to study swirl control in a GE DACRS combustor. The liquid (n-heptane) is injected in the inlet and the droplets are entrained into the



Figures 1a and b. Isocontours of axial velocity in the combustor for two different swirl numbers (left $S = 0.56$, right $S = 1.12$). The colors red and green indicate left-to-right flow, whereas blue and black indicate reverse flow. As shown, increase in swirl creates a recirculating (reverse flow) bubble in the combustor that pushes the flame outward and also stabilizes the combustion.



Figures 2a and b. Increase in swirl (left to right) in a premixed system stabilizes pressure oscillation in a gas turbine and also directly impacts the flame (yellow)-vortex (gray) interaction process. With high swirl, the shear layer at the dump plane breaks down, and a recirculation bubble forms in the combustor that stabilizes the flame; hence the combustion process.



Figures 3a and b. The effect of increasing swirl (for a fixed equivalence ratio) is to decrease the pressure oscillation amplitude and the effect of decreasing equivalence ratio (for a fixed swirl) is to increase the pressure oscillation amplitude. Data such as this can be used to devise an active swirl control strategy.

swirling air as it enters the combustor. Results show that increasing the swirl changes the spray dispersion and the combustion process, which in turn affects the pressure fluctuations. All these observations are in good agreement with experimental data. For a nonreacting case, increase in swirl increases the spray dispersion and vortex breakdown process. Large droplets tend to accumulate in regions of low vorticity (i.e., they tend to “surround” the vortices and are not present inside them). Thus, as the vortices break down into smaller features and get dispersed because of the increased swirl effect, the droplets also get dispersed (Figures 4a and b). This phenomenon impacts combustion as well, since as the droplets vaporize and become very small, they and the vaporized fuel get entrained into the vortices where they are under mixing. Thus, in high-swirl flows, flame surfaces are highly wrinkled and localized, whereas in low-swirl flows, the mixing process is delayed and the flame surfaces are less wrinkled (Figures 5a and b). Since fuel-air mixing is significantly enhanced in high-swirl flows, as confirmed here and observed experimentally, nearly all gas turbine combustors employ swirl to increase combustion efficiency. However, as noted earlier, the swirl condition that is optimal for cruise (or load) conditions need not be optimal for idle or takeoff-landing conditions. The LES model developed here can be used to understand the physics of swirl, spray dispersion, fuel-air mixing, and combustion heat release within a single formulation. Such a capability does not exist anywhere else at this time. Future studies will focus on control of the spray combustion zone when inflow swirl is changed.

The LES model developed here can be used to understand the physics of swirl, spray dispersion, fuel-air mixing, and combustion heat release within a single formulation.


The simulation code developed here is highly optimized for parallel simulations and executes without any modification on all currently available high-performance computing (HPC) systems. In fact, the gas-phase LES solver, called LESLIE3D, is now a benchmark code for DoD HPC centers. Computations such as these are only feasible on massively parallel systems such as the Cray T3E, IBM SP3, SGI O3K, and Compaqs available at HPC centers supported by the DoD. Computational support was provided by the HPC Modernization Program at the U.S. Army Engineer Research and Development Center (ERDC) Major Shared Resource Center (MSRC), the Naval Oceanographic Office (NAVO) MSRC, and the Aeronautical Systems Center (ASC) MSRC under Challenge Projects. Most of these simulations were carried out primarily at the ERDC MSRC, but some of the premixed simulations were conducted at the NAVO MSRC and the ASC MSRC. This work was also supported in part by the Army Research Office and General Electric Power Systems.

These simulations are also quite expensive. A typical premixed combustion LES using 1,000,000 grid points requires 10 GB of runtime memory and 40,000 single-processor hours on the Cray T3E to obtain sufficient


data for statistical analysis (typically, 10 flow-through times). A nonreacting spray LES using 1,000,000 grid points and 100,000 droplet groups requires 20 GB of runtime memory and 200,000 single-processor hours. On the other hand, spray combustion for the same conditions requires 100,000 single-processor hours (as the droplets vaporize, the Lagrangian tracking cost decreases rapidly). Note that although these time estimates are considerable, the rapid increase in processor speed (e.g., the SP3 is around five times faster than the T3E-900, and the new Compaq appears to be around five times faster for LESLIE3D) implies that such simulations may become practical in the near future. However, computational efficiency of codes by itself is not enough. Physically consistent and accurate models are needed in order to achieve reliable

predictions. The simulation methodology discussed here has already demonstrated its ability to capture fundamental physics of turbulent combustion. Implementation of this approach on advanced parallel systems will result in a new, advanced predictive capability.

A typical premixed combustion LES using 1,000,000 grid points requires 10 GB of runtime memory and 40,000 single-processor hours on the Cray T3E to obtain sufficient data for statistical analysis (typically, 10 flow-through times).




Figures 4a and b. Increase in swirl (left to right) also can be used to control the spray dispersion into the combustor. This in turn directly controls the combustion process in the combustor. Spray dispersion is increased in high swirl because of the breakdown of the vortex structures, which enhances local entrainment, and the formation of the recirculation bubble, which provides an effective bluff body to push the spray droplets around it and hence increase dispersion. The combustion efficiency (not shown) is also increased with an increase in swirl.



Figures 5a and b. Effect of swirl (left to right) on spray combustion is to increase fuel-air mixing because of rapid breakdown of the vortices. This in turn results in a more wrinkled flame surface and produces a more dispersed flame zone. The color red indicates a temperature isosurface at $T = 2000$ K. Yellow indicates an isosurface of the azimuthal vorticity component.

Dr. Michael Stephens

By Ginny Miller



Dr. Michael Stephens
ERDC MSRC
Scientific
Visualization Center

An onsite visualization lead in the ERDC MSRC's Scientific Visualization Center (SVC), Dr. Michael Stephens realizes the importance of communication. "Basically, my work is trying to allow for better communication between the data that's generated on our supercomputers and the scientist whose data it is," Stephens said. "What I do, in a nutshell, is draw pictures of the output data from these simulations. It serves several purposes, but underlying it all is the communication. It's 'the picture is worth a thousand words' type of approach."

Stephens holds a bachelor of science degree (1979) and a Ph.D. (1987) in chemical engineering from the University of Kentucky. He honed his scientific visualization (SV) skills at the University of Wyoming in Laramie, where he worked as an assistant professor of chemical engineering from 1984 to 1990. Stephens left academia in 1990 to join the Army High Performance Computing Research Center (AHPARC) as an onsite visualization lead at ERDC.

"The opportunity here presented itself that I could work with state-of-the-art equipment," Stephens said. "That's important, but I also enjoy the people I work with. It's a good environment that's afforded me a lot of opportunities."

Since joining AHPARC, Stephens has worked with the nationally recognized Mississippi State University/Engineering Research Center, the National Center for Supercomputing Applications, and NASA Ames. "It's a pretty small community," he said, "but being here has helped me open the doors to other groups doing visualization and graphics."

"The Resource" asked Dr. Stephens a few more questions.

How exactly does scientific visualization benefit the DoD scientist? To be blunt, most of the projects being supported in the HPCMP could not succeed without using SV in some manner.

SV is used in code development, runtime monitoring, and in better understanding of the meaning of hundreds of gigabytes of data that took thousands of CPU hours to produce. SV is the best way of distilling an enormous amount of data into an understandable form. SV also allows researchers to quickly explore much more of their data, which often leads to new questions. The visualization process brings to bear the fastest and most advanced parallel computer ever – the human visual system – to the task of understanding the meaning of the data.

In what other ways is scientific visualization important? SV is very much an organic process and as much art as it is science, perhaps even more art. SV does not stand on its own. It is not a discipline like mathematics or physics. It must be coupled with a scientist's data so that something new is created, some new understanding. A volume-rendering algorithm wouldn't be worth much if it never rendered a volume of data that was important for some purpose.

The goal is to convey the maximum amount of understanding in the simplest way using the least amount of space. An enormous number of factors go into producing a successful visualization. Many of these deal with human psychology and visual perception that to be honest we don't know a lot about.

How closely must visualization specialists work with researchers and scientists?

The researcher must be involved in the production of the visualizations in the beginning of any project. They know what they are trying to communicate or find. Because it is an iterative process, it is beneficial for the visualization specialist to know some of the language of the researcher. It makes the discussions of visualization less prone to misunderstanding and allows things to proceed more quickly. In my case, even though I'm not the scientist, having an engineering background allows me to understand what they are working on so we can generate a better visualization product for them.

How has scientific visualization changed over the years?

It has paralleled the rest of HPC in that it hasn't changed all that much fundamentally, but rather has grown in magnitude and complexity.

This is really being driven by the nature of the computational modeling (large and complex) being done today. Also, SV has matured over the past decade into an integral part of HPC. It is still the best and most compelling way for a scientist or engineer to show his or her work to sponsors, research peers, and the lay community.

The one thing that has changed is the ubiquitous nature of SV. Everybody's doing it. Commodity SV is no longer in the domain of the large computing facilities. The great advances in graphics hardware and the enormous amount of effort put in the development of visualization toolkits and algorithms have allowed SV to be incorporated into most computational researchers' projects. This advancement in technology has had the effect of pushing greater and greater graphical capabilities

into the lower end platforms – PCs running Windows or Linux. This growth in graphical power parallels the growth seen in CPU power and network speeds.

Even though there is a great deal being done on low-end platforms, there is still a great need for the so-called high-end visualization centers, such as the ones found at the national labs and the MSRCs. For it is in these facilities the new challenges of SV will be met. The complexity and sheer magnitude of the data that are to be processed into image form make it absolutely necessary that there be significant HPC resources available for visualization activity. Indeed, one of the challenges facing visualization specialists now is an efficient way of delivering visualizations to remote and graphically dispersed users. There have been several attempts at solving this problem, but this much is clear: The solution will involve the big three components of HPC – fast, high-capacity compute servers, high-speed networks, and powerful rendering hardware (visualization server).

Also, there is now more emphasis on melding data visualization with computer animation and video to better tell the story of the projects.

What kinds of visualizations are there? There are three broad categories of visualizations: personal, peer, and public. The personal visualization is what the researchers use on a daily basis in order to gain a greater understanding of the problem they are working on. These visualizations are very minimal as the researcher is using them to guide his research. The researcher is so familiar with what he or she is working on that there isn't need for decorations.

The peer type visualization is geared for an audience that is very familiar with the science of the problem and is slightly more polished than the personal visualizations. These are to be presented in journals and/or subject matter conferences where the audience has a great deal of understanding of the problem and what the issues are.

The last type (public) is the “Gone With the Wind” productions that are geared for the broadest audience. The computer animations help put the research in a context that is more clearly understood by lay persons. Examples of such visualization can be seen on NOVA or The Discovery Channel. Additionally (and more importantly to the researchers), they are used to report progress to sponsors. To this end, such visualizations have been used to great effect.

What are some emerging visualization techniques? Some of the emerging visualization methodologies are large data visualization, collaborative visualization, information visualization (InfoVis), visual simulation, virtual environments, parallel visualization algorithms, visual data mining, feature detection, distributed visualization, and remote visualization.

InfoVis is one of the two pillars of visualization, the other being Data Visualization, of which we are most familiar. InfoVis is concerned with representing information in a form so that relationships are more understandable and deals with the representation of things for which there is no physical analog. Included under the umbrella of InfoVis are the visualization techniques that deal with highly multidimensional data. An example of this pertinent to the DoD is battlespace visualization. Such a visualization must represent such items as terrain, weather, force strengths and positions, types of equipment, weapons and vehicles, types structures, soil types in different

areas, sea conditions, etc. This must be represented in ways that are easily understood by the commanders.

Another use of visualization is in code development and execution monitoring. SV is used by researchers to guide where “the interesting stuff” is and thereby allows them to refine the study on specific areas of interest, which makes better use of the resources; e.g., they don't refine the whole mesh (thereby wasting HPC resources), but rather refine the mesh in a limited area - where the action is taking place. In execution monitoring, images can be produced to show how the computer simulation is proceeding. If some error or instability begins to manifest the job can be halted, freeing HPC resources for others, and thereby preventing the churning of results that are erroneous and worthless.

What tools are popular, and why? The visualization toolkit (VTK) is wildly popular at this point. A product of Kitware, Inc., VTK uses the open source model and as such has an enormous number of people using and extending it. It is a class library API that can be used on the most widely used platforms (Windows and Unix). While VTK is natively a C++ API, it offers wrappers to other widely used languages, namely TCL/TK, Java, and Python. VTK is very powerful and covers polygonal models and image processing classes and it is FREE. I have used it to good effect in several projects.

Other widely used packages are AVS, OpenDX, Ensignt, and FAST. The first three (AVS, OpenDX, and Ensignt) can be used on data from a wide variety of disciplines (structures, CFD, etc.), while FAST is geared for the large area of CFD visualization.

What satisfaction do you get from your work? The visualization helps the scientist, who in turn helps the decision maker. It's part of a bigger picture. The proof of concept is you take somebody's data, and you try to do something with it. When you show it to the scientist you can tell the excitement factor. It sounds childlike, but it's almost a sense of awe whenever you can convert a bunch of numbers from some equations... For a lot of mathematicians and engineers, that's where a lot of people stop. In their minds they have painted the picture as much as they can. But when you take something that physically exists...it's as if you put on special glasses and you see things in a special way. And that is so cool!




Dr. Michael Stephens illustrates a visualization of Chesapeake Bay on the ImmersaDesk in the Scientific Visualization Center.




Supercomputing 2001 Conference – “Beyond Boundaries”

By Rose J. Dykes

The Institute of Electrical and Electronics Engineers/Association for Computing Machinery’s high-performance networking and computing conference for 2001 “SC2001: Beyond Boundaries” convened in Denver, Colorado, on November 10-16, with over 5,000 attendees gathering at the Denver Convention Center. Linking the Denver Convention Complex with constellation sites across the United States and worldwide with advanced Access Grid collaboration technology, SC2001 was a multinational and multicultural meeting place for communication and discussion of HPC and its impact on science and society.



(Left to right) Bob Athow and Roy Campbell help to support the DoD HPCMP booth activities



(Left to right) Paul Adams and Tom Biddlecome, ERDC MSRC scientific visualization staff, demonstrate scientific visualization at the DoD HPCMP booth


The ERDC MSRC participated in SC2001 by presenting tutorials and papers, giving technical talks and scientific visualization demonstrations at the DoD HPCMP booth, attending presentations by others, viewing and discussing the state of the art in HPC, and helping support the HPCMP booth.

A one-half day tutorial entitled “Mixed-Mode Programming Introduction” presented by Dr. Dan Duffy discussed the benefits and pitfalls of multilevel parallelism using the Message Passing Interface combined with threads. Dr. Duffy also discussed the pros and cons of various tools that can be used across platforms to help the application developer to optimize




Greg Rottman discusses the DoD HPCMP Metacomputing Initiative





and debug an MP program and showed sample codes illustrating a comparison of different MP methods and presented their resulting speedups.

In his paper presentation entitled “Coastal Ocean Modeling of the U.S. West Coast with Multiblock Grid and Dual-Level Parallelism,” Dr. Phu Luong discussed how the performance results from the Multiblock Grid Princeton Ocean Model demonstrated the effectiveness of both the MPI-Only and MPI-Pthreads code versions by lessening memory requirements and long processing times and improving load balance.




**Technology Enhancements
in the ERDC MSRC
Computational Environment**

By John E. West

As the FY01 hardware installations near completion, plans have been finalized for follow-on configuration enhancement activities at the ERDC MSRC. Planned technology enhancements to the HPC target configuration include expansion of the existing Compaq SC40 system, expansion of the Cray T3E, and the addition of storage capacity in the mass storage and archival system.

Planned technology enhancements to the HPC target configuration include expansion of the existing Compaq SC40 system, expansion of the Cray T3E, and the addition of storage capacity in the mass storage and archival system.

The current Compaq AlphaServer SC40, delivered during the FY01 acquisition cycle, consists of 64 nodes, with each node containing four 833 MHz Alpha EV68 CPUs and four gigabytes of RAM. An additional 64 nodes will be added to this system as part of the FY02 enhancements, bringing the total number of processors to 512. Each node in the expanded system will have the same configuration as the original 64 nodes, resulting in a homogeneous system. Two of the expansion nodes will be configured as file server nodes, adding to the four file server nodes currently in service. These nodes will not be available for computational jobs. In addition to the computational nodes, the proposed expansion will add approximately 2.4 terabytes (raw) of Fibre Channel RAID5 storage to the SC40 configuration. The expansion will add approximately 420 gigaFLOPS of computational capacity to the SC40, bringing the final net computational capacity to 853 gigaFLOPS peak computational capacity (a gigaFLOPS is 10^9 floating-point operations per second).



The next major computational enhancement at the ERDC MSRC involves the combination of the NAVO MSRC and ERDC MSRC T3Es. This will boost the capability of ERDC's T3E from 960 gigaFLOPS to 1.9 teraFLOPS via the addition of 1,088 Alpha EV56 processors operating at 450 MHz (a teraFLOPS is 10^{12} floating-point operations per second). The T3E combination increases the available disk storage by 1.5 TB to approximately 4.5 TB (raw). The proposed T3E final configuration is as shown. The table illustrates the logical arrangement of contiguous processor elements. As this system is integrated into the ERDC environment, ERDC will also assume the data and remaining FY02 user allocations from the NAVO MSRC.

These system upgrades will boost the ERDC MSRC's aggregate peak computational rating to 4.2 teraFLOPS, providing world-class capabilities to support the DoD.

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

Strategy for Performance Benchmarking in the HPCMP

By Dr. William A. Ward, Jr.

The DoD HPCMP, like any other large corporate user of computer systems, often requires a series of benchmark tests to evaluate the performance of installed systems and to compare that of proposed, competing systems. The quality of any benchmark test package depends on the presence of several important characteristics. These characteristics serve as goals for the test package constructor and guide the constructor in the selection of test package components (e.g., synthetic tests and application codes). First, and most importantly, the benchmark must be representative of the current and projected workload on DoD HPC systems. Also, in the context of large system acquisition activities, portability and ease of use are also highly desirable test package features.

The benchmark must be representative of the current and projected workload on DoD HPC systems.

Historically, benchmark test packages have contained codes chosen from one or more of the following types:

- ↳ Relatively short synthetic programs, such as Whetstone and Streams.
- ↳ “Toy” benchmarks, such as Quicksort and Prime Sieve.
- ↳ Widely used off-the-shelf codes or “package” kernels, such as Linpack and ScaLAPACK.
- ↳ “Application” kernels, that is, sections of code extracted from actual application programs that perform a significant fraction of work.
- ↳ Complete applications.

The two most recent test packages, one each for the Technology Insertion-Fiscal Year 2001 (TI-01) and TI-02 HPCMP procurements, have used a combination of synthetic benchmarks to measure peak system performance and actual applications to more closely model expected operational performance. These latter tests, in addition to being run in dedicated mode on candidate systems, were also used to construct a “load” test involving the submission of a stream of application jobs.

A problem with this approach is that significant work is required of the vendors in order to complete the required benchmarks. Because of the size and complexity of the test package and the limited time allowed for the vendor to respond, a vendor may be able to only get the benchmarks to run in an unoptimized fashion or

may submit incomplete or erroneous results. The TI-01 test package contained 13 application codes, each of which had one or two sets of input data. Each code/input combination was run at several different numbers of CPUs. A sizeable suite of synthetic tests were also included to limit the complexity of the test package and reduce the effort necessary to run the tests. The TI-02 test package required only five application codes, Cobalt-60, CTH, GAMESS, LESlie3D, and Navy Layered Ocean Model (NLOM). With the exception of NLOM, however, the TI-02 test package still required two sets of input data for each application code and, of course, varying numbers of CPUs. This most recent test package still proved to be a stiff challenge for participating vendors, and at least one vendor chose not to run the tests.

To construct these [synthetic applications], actual codes first are instrumented and profiled to determine crucial resource utilization characteristics, e.g., number of floating-point operations, number of loads, number of stores, amount of I/O, and number and size of messages passed between processors.

To remedy this situation, the HPCMP is exploring the use of simpler, yet still representative, test cases. However, construction of such tests will not be easy. An example of this new type of test package may include (a) synthetic tests as before; (b) perhaps two applications run for a total of eight different numbers of CPUs and using one or two input cases; and (c) synthetic applications that model the behavior of actual applications. Both *a* and *c* would be easy to compile on a variety of platforms because they would be small and self-contained. Tests in *b* should be easy for most vendors to implement because they would either be chosen from a previous year’s test package or would be released earlier than the rest of the tests.

Component *c* would be the core of the test package. To construct these, actual codes first are instrumented and profiled to determine crucial resource utilization characteristics, e.g., number of floating-point operations, number of loads, number of stores, amount of I/O, and number and size of messages passed between processors. This first step will probably be the most difficult



Dr. William A. Ward, Jr.

part of the project. Next, small synthetic applications or kernels with known, parameterized input would be identified for use as building blocks to construct the synthetic applications; these building blocks would be called from a yet-to-be-constructed main program. The third step would require use of linear regression or linear programming to determine the number of subroutine calls for each building block so as to match the resource utilization in the actual program. These building blocks would be called by this main program in some varying order to mimic overall code behavior (e.g., time-stepping).

The final stage in this process is educational in nature. Graphical performance profiles will be produced illustrating that the resource utilization of the synthetic application tracks that of the actual application with respect to elapsed wall time and other significant metrics. Hopefully, this will convince scientists that such synthetic applications really do model actual code behavior.

The HPCMP is exploring the use of simpler, yet still representative, test cases.

Cpusets on the Origin 3800

By Dr. Jeff Hensley

The Origin 3800 (O3K) series machine (Ruby) at the ERDC MSRC recently underwent a configuration change designed to enhance the overall efficiency of the system. In February 2002, Ruby was configured to use cpusets, a method of logically partitioning the processors and memory on the machine.

The implementation of cpusets was undertaken in response to users who observed large inconsistencies in the time required to run identical or similar jobs.

The implementation of cpusets was undertaken in response to users who observed large inconsistencies in the time required to run identical or similar jobs. ERDC MSRC staff performed tests on a 256-processor machine and on Ruby (512 processors) to evaluate the magnitude of the performance degradation. The results of these tests indicate that using cpusets on Ruby greatly improves the overall throughput efficiency of the machine.

The results of these tests indicate that using cpusets on Ruby greatly improves the overall throughput efficiency of the machine.



SGI Origin 3800 "Ruby"

A throughput test was created to run on a 256-processor O3K made available by SGI. For the throughput test, a collection of jobs was submitted to the queue to simulate the runtime environment. The jobs consisted of actual user codes and requested different processor counts. The throughput test was run three times with the machine configured without cpusets and three times with cpusets implemented. The results of these tests are summarized in Table 1.

When cpusets were used, the total time for the test was remarkably consistent. However, when cpusets were not used, the time was quite variable and significantly greater than when cpusets were used. A similar but simpler and less demanding experiment was executed

**Table 1
Total CPU Time for Throughput
Tests on 256-Processor SGI O3K**

Test Number	CPU Time (hours) for Entire Test	
	With cpusets	Without cpusets
1	747	1,118
2	747	1,200
3	747	1,208

on Ruby (the 512-processor machine) and also showed a significant performance improvement when using cpusets.

The performance degradation that occurs without cpusets does not affect all jobs the same. Some jobs showed very little variation in runtime performance, while others showed a slowdown of as much as a factor of four when running without cpusets.

What is the cause of this performance loss? On the SGI O3K, tasks are not “stuck” on specific processors. The operating system searches for an available processor for any task that needs to be performed. So, if a task is momentarily in a sleep state on a processor, the system may decide to use that processor for another task. When the sleeping task starts back up, it may then be moved to another processor. Although the tasks may migrate from processor to processor, the data stored in memory do not migrate. On Ruby, there are four processors and four Gbytes of memory on an “ssinode.” A processor’s memory access is much faster for memory on the ssinode. If tasks migrate across ssinodes, it is possible that a given task winds up on a processor that is far removed from the data that it needs, resulting in slower performance.

While the use of cpusets is for the most part transparent to the user, there are some issues of which users need to be aware.

Ruby is now configured with a “boot cpuset” consisting of eight processors. The boot cpuset is used to

handle system processes. All of the other processors are available for dynamic cpusets. When PBS launches a user’s job, a (dynamic) cpuset is created for that job; other jobs are not allowed access to the processors or memory.

While the use of cpusets is for the most part transparent to the user, there are some issues of which users need to be aware.

1. The syntax for requesting processors in PBS has changed. To request <ncpus> processors, use the “-l ncpus=<ncpus>” rather than the syntax “-l nodes=1:ppn=<ncpus>” that was used previously.
2. PBS builds cpusets in terms of “ssinodes,” which are groups of four processors together with four Gbytes of memory. The number of processors allocated by the system will be the smallest multiple of four processors greater than or equal to the number of processors requested. In particular, if a job asks for one processor, the system allocates four processors. User allocations will be charged according to what the system allocates. Users who run serial jobs should take special note of this. One strategy is to bundle several serial jobs together in order to use all four processors on an ssinode. For additional information about bundling jobs on the O3K, see http://www.ercd.hpc.mil/faq/tips/comp_tips/bundling.htm.
3. Jobs that require large amounts of memory may require special treatment. Each ssinode has four Gbytes of memory, and a job can use only memory contained within its cpuset. If a job needs more than the memory available in the cpuset, the task will use virtual memory and swapping will occur (with a likely degradation of performance). The alternative is to request more memory (and more processors). For example, the PBS option “-l mem=12gb” asks for 12 Gbytes of memory (three ssinodes or 12 processors).

The experiments that have been run show that using cpusets should provide a significant overall improvement in machine performance that far outweighs the inconveniences mentioned above. The use of cpusets removes the large variation in performance that some users have observed and provides for more efficient use of the machine, which is beneficial for everyone.




David Sanders

User Disk Striping on the T3E

By David Sanders and Dr. William A. Ward, Jr.

Many computer systems implement some form of hardware disk striping (spreading a file across multiple cylinders or spindles) to improve I/O performance. As part of the effort to meet the TI-01 NLOM target time on the Cray T3E at the ERDC MSRC, it was discovered that the T3E also implements software disk striping. At Cray Inc.'s suggestion, software disk striping was applied to seven NLOM input files ranging in size from 44 Mb to 1.8 Gb. As a result, the execution time for the NLOM NA825 test case



Dr. William A. Ward, Jr.

was reduced from 1,308 to 1,162 seconds. Unlike most other forms of disk striping, users without special system privileges may perform this technique. Cray Inc.'s guidelines indicate that the approach is more easily applicable when the file size is known in advance (e.g., an input file) and is most effective for files larger than 10 MB. The steps involved in this approach can be used on either input files or output files and are somewhat different for each case.

For input files, the steps are as follows:

```
/etc/fck -b <file>
```

List the number of blocks in the input file by disk partition. Typically, files reside on only one partition (unstriped), although the file may be broken into one or more "chunks."

```
df -p <filesystem>
```

List the available blocks, by partition, in the filesystem in which the striped file will reside, e.g., /tmp. Select the partitions across which the file will be striped; heavily loaded partitions should generally be avoided.

```
mv <file> <save>
```

Move the file to a temporary location.

```
setf -n <#blocks>b -p <partitions> <file>
```

Preallocate the file so that it takes up the specified number of blocks and is split across the desired partitions. Partitions are specified as a range (e.g., 7-9), as a set (e.g., 2:3:11), or a combination of both (e.g., 2:3:7-9:11).

```
fdcp <save> <file>
```

Copy the file to its former location preserving the partition structure created using **setf** (use of **cp** will destroy this structure). This effectively stripes the input file across the previously specified partitions.

```
assign -p <partitions> -n <#blocks> -F bufa:<bufsize>:<numbufs> f:<file>
```

Allocate memory buffers for reading the file. The partitions and #blocks are the same parameters as specified for the **setf** command. Declare the size ("bufsize") and number ("numbufs") of buffers to use for this file. Buffer sizes are in units of 4,096-byte blocks. Buffer sizes less than approximately 128 blocks are generally not effective. If possible, the number of buffers should be at least twice the number of partitions. Be aware that these buffers reduce the amount of memory available for the application itself. For output files, the process is much simpler. It is helpful, but not required, if the output file size is known prior to execution.

```
assign -p <partitions> [-n <#blocks>] -q <chunk-size> -F bufa:<bufsize>:<numbufs> f:file
```

As before, use **-p** to declare the partitions across which the file is to be striped, and use **-F** to declare the number of memory buffers and their size. The guidelines for number of buffers and sizes are the same as for

input files. If the total file size is known in advance, then use `-n` to specify that value in units of 4,096-byte blocks. Also, specify the size of each contiguous “chunk” of data on each partition. If `-n` is being used, then the chunk-size value should be no larger than “#blocks” divided by the number of partitions. If `-n` is not used, then the chunk-size value can vary as desired.

The above process will vary from application to application as file sizes and numbers of files change. In general, the larger the files the more effective the technique. Performance will also vary depending on the total amount of I/O in the application. Some experimentation will most likely be required to determine the best numbers to use for the various `setf` and `assign` parameters. A more detailed I/O tutorial can be found at <http://hpcf.nersc.gov/training/tutorials/T3E/intro/>.



Cray T3E “Jim”

Word Search – Metacomputing Terminology

Answers for word search (from page 7).



Metacomputing Grid Workshop

By Rose J. Dykes

The 1st DoD High Performance Computing Modernization Program Metacomputing Grid Workshop was held at the ERDC Information Technology Laboratory, Vicksburg, Mississippi, on February 25-28, 2002, hosted by the ERDC MSRC. Greg Rottman, ERDC MSRC and Project Manager of the HPCMP Metacomputing Grid Initiative, planned and coordinated the Workshop with assistance from Rose Dykes and many others. Drs. Dan Duffy and Jeff Hensley, Tim Dunaway, and Jerry Morris, all of the ERDC MSRC, conducted DoD Specific Grid Training (hands-on) on the last day of the workshop.

Attendees from across the Nation expressed satisfaction with the 4-day Workshop by stating that they found it informative, interesting, and professionally valuable. After spending long days in the Workshop sessions, everyone welcomed the nighttime events planned for their enjoyment and relaxation. Several expressed their appreciation for the hospitality shown.



Dr. Jeffery Holland, ITL Director, welcomes attendees to ERDC ITL



Bradley Comes, ERDC MSRC Director, gives welcome and presents introduction



Greg Rottman, Workshop Coordinator, discusses the agenda and metacomputing initiatives



Dr. Dan Duffy, ERDC MSRC, presents the DoD Grid



Dr. Michael Gourlay, Colorado Research Associates, presents Uniform Command Line Interface




Charlotte Coleman, ASC MSRC, discusses Information Environment



Dr. Marlon Pierce, Indiana University, presents Gateway Project



Steve Thompson, ARL MSRC, talks about Job Migration Tool



Charles Bacon, Argonne National Laboratory, presents Globus training



Mike D'Arcy, Argonne National Laboratory, presents additional Globus training



Dr. Aram Kevorkian, Space and Naval Warfare Systems Command, presides over the DoD HPCMP Metacomputing Working Group Meeting



A full house for the workshop



A few attendees eating lunch on the patio at ITL



Valerie Thomas, HPCMP, David Cronk, University of Tennessee (both on left), and others enjoy southern culinary delights




Christine Cucchi, NAVO MSRC, watches tugboats coming down the Mississippi River



Tom Cortesi, NAVO MSRC, entertains on the piano



Attendees enjoy dinner at Jacques in the Park



Lance Terada, Maui HPC Center, performs at Karaoke Night

ERDC MSRC Represented at IEEE Visualization Conference

By Dr. Michael Stephens and Tom Biddlecome

The 12th annual Institute of Electrical and Electronics Engineers (IEEE) Visualization Conference was held at Paradise Point in San Diego, CA, October 21-26, 2001. Dr. Michael Stephens and Tom Biddlecome from the ERDC MSRC Scientific Visualization Center were among the attendees. This highly technical conference, attended by more than 600 people, offered participants the three following major tracks:

↳ **Visualization Techniques.** This track offered papers and case studies in the following areas:

- Information Visualization
- Human Perception
- Data Simplification
- Multivariate Visualization
- Large Data Visualization
- Virtual Environments

↳ **Visualization Algorithms.** This track had papers in the following areas:

- Volume Rendering
- Flow Visualization
- Data Compression
- Vector and Tensor Visualization

↳ **Visualization Applications.** Papers in this track offered the widest range of practical use of visualization in science and engineering. Talks covered the use of visualization in the following areas:

- Archaeology
- Astrophysics
- Biomedical Research
- Chemistry
- Education
- Mathematics
- AVS
- Data Explorer
- VTK
- Khoros

- ↳ **Visualization Techniques**
- ↳ **Visualization Algorithms**
- ↳ **Visualization Applications**

Prior to the conference, which began on October 23, several tutorials and two, 2-day symposia took place. The tutorials were as follows:

- Large Data Visualization and Rendering
- Information Visualization for Beginners
- Rendering and Visualization in Parallel Environments
- Multiresolution Techniques for Surfaces and Volumes
- From Transfer Functions to Level Sets: Advanced Topics in Volume Image Processing

A symposium was conducted on Information Visualization, which deals with visualizing data that are categorical in nature and whose relationships are not typically described by mathematical relationships. Examples of this are tree structure graphs, which describe the time and space relationships of the evolution of a species or the hierarchical structure of the World Wide Web. The second symposium was on Parallel and Large-Data Visualization and Graphics, an area of particular interest to the ERDC MSRC visualization team.



From data simplification to vector visualization, all aspects of the conference are valuable to DoD applications. The next IEEE Visualization Conference will be 27 October - 01 November 2002 at the Park Plaza Hotel in Boston, MA. Conference information is available at <http://vis.computer.org/vis2002>

Exploring the Future of High-Performance Computing via the Access Grid Node

By Rose J. Dykes

The Future Technologies domain of Pacific Northwest National Laboratories Computational Science and Engineering's Technical Network hosted a session entitled "High Performance Computer Requirements and Projections to 2010" via the Access Grid Node. The opportunity was available for worldwide participation in a discussion of what will be required when high-performance computers increase in computational power by four orders of magnitude over the next 10 years, which has been projected.

Interesting points were presented and explored in this particular session via the Access Grid Node regarding computational power increase. The computational power increase projected above requires new and creative facilities; this, of course, will impact how computers in the future will look. Ubiquitous computing, the ability to store petabytes of data in a single location, and other new technologies will be enabled with the rate of change for networking and storage devices.



From the Access Grid Node located in the ERDC Information Technology Laboratory, John West (left) and Mike Donovan (right) explore the future of HPC



(Left to right) William Bennett, Spence Cobb, Mike Donovan, Roy Campbell, Nathan Bill, ARSC, and Greg Rottman discuss mass storage with the Arctic Region Supercomputing Center (ARSC) via the Access Grid Node



MG Robert Griffin (third from right), Director of Civil Works, U.S. Army Corps of Engineers, Washington, DC, shown outside the Joint Computing Facility Control Room, April 18, 2002



Lawrence E. Clark, Deputy Director for Science and Technology, and Ron L. Marlow, Director, Conservation Engineering Division, both of the Natural Resources Conservation Service, visit the ERDC MSRC Scientific Visualization Center, April 15, 2002. (Left to right, Mr. Clark; Bob Welch, ITL Technical Director; Mr. Marlow; Bob Athow, ERDC MSRC PET Technical Advisor; and Milton Myers, Special Assistant to the Director, ERDC Geotechnical and Structures Laboratory)



U.S. Army Engineer Division, South Atlantic, Atlanta, Emerging Leaders are briefed by Dr. Jeffery Holland, ITL Director, before touring the ERDC MSRC Scientific Visualization Center, April 9, 2002



Environmental Quality in Progress Review attendees view a scientific visualization of the Chesapeake Bay, April 9, 2002



Dr. Ed Theriot (far left), ERDC Environmental Director, escorts visitors from the Japanese Defense Facilities Administration Agency; U.S. Army Corps of Engineers, Japan District; and LTC William Buechter (far right), ERDC Construction Engineering Research Laboratory, February 21, 2002



COL Neal Gafney (second from left) and other Liaison Officers/Planners, U.S. Army Corps of Engineers, shown in the Joint Computing Facility, March 4, 2002

visitors



(Left to right) Jerry Satterlee, Chief, Engineering Division, U.S. Army Engineer District, New Orleans; Drs. Ellis (Buddy) Clairain and Russell Theriot, ERDC Environmental Laboratory; John Saia, New Orleans District Deputy District Engineer for Project Management; and David Stinson, ERDC MSRC, discussing mass storage for HPC, February 14, 2002



AMEC (formerly AGRA Earth & Environmental) visitors including Roger Jinks (far left), President of the Earth and Environmental Operations, February 5, 2002



(Left to right) Bob Athow, ERDC MSRC, LTC Ray Dunton, Dr. Michael Grodowitz, and other ERDC Environmental Laboratory visitors in front of the SGI Origin 3800 supercomputer, February 11, 2002



(Left to right) CPT Eric Betts, ERDC ITL; Dr. Peggy Wright, Program Manager for the Corps of Engineers Enterprise Infrastructure Services; LTC Bill Flynt, Director, Threats to Critical Infrastructures, Foreign Military Studies Office, TRADOC DCSINT, Fort Leavenworth, Kansas; and Brad Comes, ERDC MSRC Director, shown in the Joint Computing Facility, January 8, 2002



Dennis Gilman (left), ERDC MSRC, and Dr. Vellore S. Gopalaratnam, University of Missouri-Columbia, discuss HPC outside the control rooms of the Joint Computing Facility, January 28, 2002



Dr. Richard E. Price (far left), ERDC Environmental Laboratory, escorts attendees of the Hazardous Toxic Waste Program Management Meeting, U.S. Army Engineer Division, South Atlantic, Dallas, Texas, as they tour the Joint Computing Facility with Bob Athow (far right), ERDC MSRC, January 15, 2002

Stereo Viewing for Scientific Visualizations – Not With Liquid Crystal Displays

One of the benefits of humans having two eyes is that each eye has a slightly different view of the world. The brain takes these two images and combines them to create a three-dimensional (3-D) image. As visualization in the computer world becomes increasingly more realistic, the desire to see in 3-D also increases. One way in which this can be accomplished is by wearing shutter glasses.

Suppose a display device, such as a cathode ray tube (CRT) monitor, displays 90 frames per second of an image. Half of those frames (45) could be drawn to be seen by the right eye; half could be drawn for the left eye. The image for the right eye would be slightly offset from the image meant for the left eye. The shutter glasses would allow the correct image to be shown to the right and left eye by blocking the image to the opposite eye. This allows the brain to combine the two slightly different images into a 3-D image.

One problem with this approach to 3-D imaging is that the image for one eye must substantially decay before

the image for the next eye is displayed. If this does not happen, then the user will see a double image, or “ghost,” as the image meant for the right eye bleeds into the image for the left eye. CRT displays have a decay rate fast enough to prevent this ghosting. Liquid Crystal Displays (LCDs) currently do not have a fast enough decay rate, however, which leads to ghosting. This makes stereo viewing using shutter glasses on LCDs, for now, impractical.

A second problem with LCDs and stereo viewing is that the display device needs to be able to support a high vertical refresh rate in order to prevent flicking of the image. Many CRTs, though not all, are capable of providing a vertical refresh rate of 100 Hz or more. Many LCDs are limited to a vertical refresh rate of 75 Hz. This also contributes to making stereo viewing using shutter glasses on LCDs impractical.

Information in this article is based on an article from the Internet entitled “StereoGraphics Products and Modern Displays” by Brick Bradford.

Buffered Input/Output on Compaq SC Platforms – Significantly Increases Performance

On the Compaq SC40 and SC45, unbuffered input/output (I/O) is the default for the Fortran 90 compiler. This forces the computer to write the data to the disk immediately at each write command. If a code has many writes, especially if they each write a small piece of data, this “no buffered I/O” state can cause a drastic increase in execution time. When a program is compiled with the `-assume buffered_io` option,

during execution the data is not written to the disk until the buffer is filled. With a particular program containing a vast amount of small writes, a speed-up in execution time of a factor of ten was achieved by compiling with this option. A disadvantage of buffered I/O is that data in the buffer will be lost if a program fails or is killed during execution.

(Left to right) Brad Comes, Jane Giffin, Christie Baker, Rose Dykes, and Anne Page, all of the ERDC MSRC, enjoy a “Working Lunch” with Mississippi Senator Mike Chaney



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.

AFOSR	Air Force Office of Scientific Research	HPCMP	High Performance Computing Modernization Program
AHPCRC	Army High Performance Computing Research Center	InfoVis	Information Visualization
ARL	Army Research Laboratory	ITL	Information Technology Laboratory
ARSC	Arctic Region Supercomputing Center	LCDs	Liquid Crystal Displays
ASC	Aeronautical Systems Center	LES	Large-Eddy Simulations
CFD	Computation Fluid Dynamics	MOS	Mississippi State University/Ohio State University
CHSSI	Common High Performance Scalable Software Initiative	MSRC	Major Shared Resource Center
CPU	Central Processing Unit	MSU	Mississippi State University
CRT	Cathode Ray Tube	NAVO	Naval Oceanographic Office
CSC	Computer Sciences Corporation	NLOM	Navy Layered Ocean Model
CSM	Computational Structural Mechanics	OKC	On-Line Knowledge Center
CWO	Climate/Weather/Ocean Modeling and Simulation	PET	Programming Environment and Training
DCs	Distributed Centers	SRC	Shared Resource Center
DoD	Department of Defense	SRI	Stanford Research Institute
EQM	Environmental Quality Modeling and Simulation	SV	Scientific Visualization
ERDC	Engineer Research and Development Center	SVC	Scientific Visualization Center
HEDM	High-Energy Density Materials	TI	Technology Insertion
HPC	High-Performance Computing	TICAM	Texas Institute for Computational and Applied Mathematics
HPCMO	High Performance Computing Modernization Office	UAG	Users Advocacy Group

ERDC MSRC publications

- 01-29 "An MPI Quasi Time-Accurate Approach for Nearshore Wave Prediction Using the SWAN Code," Stephen F. Wornom
- 01-30 "High Performance Computing Summer Intern Program - Summer 2001 Final Report," John E. West
- 01-31 "A Comparison of Wave Hindcasts for a Gulf of Mexico Storm Using Quasi Time-Accurate and Time Accurate Methods," Stephen F. Wornom, Richard Allard, Y. Larry Hsu
- 01-32 "Minimizing Runtime Performance Variation with Cpusets on the SGI Origin 3800," Jeff Hensley, Robert Alter, Daniel Duffy, Mark Fahey, Lee Higbie, Tom Oppe, William Ward, Marty Bullock, Jeff Becklehimer

training schedule

- July 30-31:** Workshop on Computational Fluid-Structure Interactions
- August:** Interlanguage Programming Practices
- September:** Dense and Sparse Linear Algebra Libraries

Courses will be taught at the ERDC MSRC. Agendas and on-line registration are available at <http://www.erd.c.hpc.mil/training/schedule/schedule.htm>.

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

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.erdchelp@erdc.hpc.mil

E-mail: msrchelp@erdc.hpc.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.



ERDC

Major Shared Resource Center

ERDC MSRC

In Touch

with Tomorrow

