

The Resource

U.S. ARMY CORPS OF ENGINEERS, WATERWAYS EXPERIMENT STATION

NEWSLETTER



ERDC!!!



Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 1999		2. REPORT TYPE		3. DATES COVERED 00-00-1999 to 00-00-1999	
4. TITLE AND SUBTITLE The Resource. Spring 1999				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Engineer Research and Development Center,ATTN: ERDC MSRC HPC Resource Center,3909 Halls Ferry Road,Vicksburg,MS,39180-6199				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

From the Director's chair...

I was prepared to write an article about the HPCMP, the CEWES MSRC, and recent related events. I got started on it this evening and then received a phone call from a friend. I have known this gentleman for more than a decade. He was involved with supercomputers in the early days. He can tell you stories about GE, Honeywell, CDC, Harris, Prime, VAXes, Ncube, TI machines - the gambit. In the 80s and early 90s, he ventured into Cray vector systems and in the mid-90s, into large scalable systems. He has an intuitive feel for what will play in Peoria and a wisdom about high performance computing. He sets his sights high and strives for excellence. He once told me that he played competitive Ping-Pong and that his father played to his weak spots to help him improve his game. He applies this to his daily life, always taking risks to strengthen his game. He is an aggressive player, but he also plays by the rules. He is a fair man.

We have attended technical conferences and we have attended funerals. We have worked all night. We have missed airplanes. We have gotten traffic tickets at one o'clock in the morning. We have shouted at each other, we have laughed, and . . . we have cried. He informed me tonight that he is leaving the Information Technology Laboratory, the place he transformed from "the computer center" with less than two dozen employees into a nationally recognized multi-disciplinary laboratory with several hundred employees. I feel sad not only because



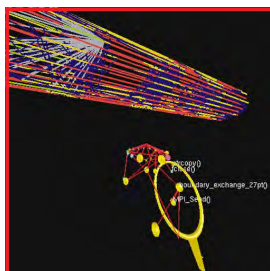
*Brad Comes, CEWES MSRC Director (left),
and Dr. N. Radhakrishnan, ITL Director*

we are going to miss him but also because I know how hard it is for him to move on. However, I am also happy for him, knowing that the challenges he faces are what he lives for. Dr. Radha, we wish you the very best. Thank you for all that you have done. And remember - your legacy will remain at ITL forever.



The Jamie L. Whitten building, headquarters of the Information Technology Laboratory founded by Dr. N. Radhakrishnan

Bradley M. Comes
Director, CEWES MSRC



page 6

Feature Articles

Parallelized Solvers for Sparse Linear Systems on the SGI Origin2000	6
--	---

Virtue Software	6
---------------------------	---

The End of an Era	7
-----------------------------	---

NASA Liaisons Collaborate on Key Projects	11
---	----

It Takes Two to TANGO	12
---------------------------------	----

ITL Dedicates New Addition	15
--------------------------------------	----


Challenge Projects Up Close: Integrated Topside Design	16
--	----

Interview with Jamie Grantham	17
---	----

Scientific Illustration in HPC	18
--	----



page 7



page 12

MSRC Regulars

Announcements	2
-------------------------	---

Off-Campus	3
----------------------	---


Visitors	5
--------------------	---

UNIX Tool Tips	20
--------------------------	----

1999 Training Schedule	20
----------------------------------	----



page 16



page 18

announcements

CEWES Changing Name to ERDC

Mary Gabb


On 1 October 1999, the CEWES MSRC will undergo a name change along with eight other research laboratories in the Corps of Engineers. The CEWES MSRC will then be known as the ERDC MSRC, or the Engineer Research and Development Center MSRC.

The CEWES MSRC is part of the U.S. Army Corps of Engineers, Waterways Experiment Station (WES) Information Technology Laboratory, one of the five laboratories in WES. The others are the Coastal and Hydraulics Laboratory (CHL), the Environmental Laboratory (EL), the Geotechnical Laboratory (GL), and the Structures Laboratory (SL). Those four, along with the Topographic Engineering Center (CETEC), the Cold Regions Research and Engineering Laboratory (CECRL), and the Construction Engineering Research Laboratory (CECER), all located at sites other than WES, will comprise the research laboratories under ERDC.

The name change reflects a reorganization that will place a single director of the laboratories in Washington, D.C.

Corps of Engineer research laboratory support elements have already undergone the name and organizational change in government fiscal year 1999 (FY99). FY00 will be the green light for the research laboratories and the MSRC will officially change its name 1 October 1999. The physical site where WES is located is now called The ERDC at the Waterways Experiment Station.

So look for our name change in the coming months. How do you pronounce it? Listen to the frog.



Pictured left to right are Charles Ray, Brad Comes, and Patrick Heraghty

CEWES Passes DISA Inspection - Grade A

Patrick Heraghty

From 1-5 February 1999, the Defense Information Systems Agency Security Test and Evaluation (DISA ST&E) team conducted an evaluation of the security environment for both the unclassified and classified processing areas of the CEWES MSRC. The ST&E Director Ms. Marie Green of DISA INFOSEC Program Management Office acknowledged the contributions of site personnel involved with the test, "The site is commended on the pre-inspection assistance and the concern and awareness of security throughout the certification process."

Hundreds of systems and devices were checked. A very small number of vulnerabilities were found by the ST&E team, of which fifty percent were corrected while the test team was still on-site. The rest were satisfactorily addressed by 31 March 1999. Mr. Charles Ray, CEWES MSRC Information Systems Security Officer and Mr. Patrick Heraghty, Nichols Research Corporation ST&E Lead accepted, on behalf of the CEWES MSRC team, the DOD Certification of Compliance Plaque.

T3E Upgrade, New IBM

Michelle Morgan Brown

The CEWES MSRC was recently awarded a contract that secures the center's position as one of the top ten most powerful high performance computing sites in the world.

As part of the contracted improvements, the existing SGI/Cray T3E was upgraded in February 1999 from a 336-processor system to a 544-processor system. The new system contains 136,000 MBytes of central memory and 750 GBytes of disk space. This large system allows users to utilize up to 520 processors on a single computer program, significantly increasing the scientist's ability to address large research and development challenges. During the first day of use, the new T3E executed a



SGI/Cray T3E

500-processor computer program, the largest single computer program ever executed at the CEWES MSRC.

The recent contract enhancements also provide for the installation of the newest of the IBM supercomputers, an IBM high-node SP (8-way SMP). The new system, which is scheduled for delivery in early fall of 1999, will include 512 processors, 256,000 megabytes of memory, and 4,600 gigabytes of disk space. The new IBM high-node SP will complement the existing CEWES IBM system to create an integrated compute complex of 894 IBM processors.

The latest MSRC additions round out the CEWES MSRC capability, which includes a 16-processor Cray C90, a 128-processor SGI Origin2000, the 894-processor IBM complex, and the 544-processor T3E. With the recently contracted upgrades, the CEWES MSRC will have a peak performance rating of 1.4 trillion operations per second.



New IBM high-node SP to be installed in the fall of 1999

CEWES MSRC

off-campus

Mary Gabb

First Southern Conference on Computing

The University of Southern Mississippi was host to the First Southern Conference on Computing on 4-5 December 1998. The aim of the conference was to bring together researchers in all areas of computation and, in an informal atmosphere, attempt to develop relationships between its various threads. More than 70 presentations were given over the two-day conference in a wide range of topics including Ocean Modeling, Parallel/Distributed/Concurrent Methods, Computational Predictions of Molecular Structures, and Artificial Intelligence. Dr. Clay Breshears, CEWES SPP Tools Lead, served on the Programming Committee and presented the paper "A Computation Allocation Model for Distributed Computing Under MPI_Connect," co-authored with Dr. Graham Fagg (University of Tennessee, a PET partner).

The conference was sponsored by the School of Mathematical Sciences of the University of Southern Mississippi in cooperation with the Department of Mathematics and Statistics and the Department of Computer Science of Mississippi State University, and the Department of Computer Science of Louisiana State University. There are plans to make this a bi-annual event.

MPIDC'99

MPIDC'99 (the Message Passing Interface Developer's and User's Conference 1999, Atlanta, 10-12 March) centered on design, implementation, and realization of cluster parallel and dedicated parallel message passing systems, applications, and related software technology.

Trey White, CEWES Computational Migration Group, and Dr. Steve Bova, CEWES computational fluid dynamics lead for PET, co-authored a paper entitled "Where's the Overlap? An Analysis of Popular MPI Implementations," which Mr. White presented at the meeting. In the paper, White and Bova question the efficacy of the common technique of hiding communication latency by coding the communication in such a way that it can be performed simultaneously with floating-point calculations. They designed experiments and examined the behavior of Message-Passing Interface (MPI) implementations on the Cray T3E, IBM SP, and SGI/Cray Origin2000 located at CEWES. The authors concluded that because there is no mechanism in these machines to provide third-party handling of the communication, it does not behoove the programmer to go to extra effort to program the communication in this way.

For more information on this work, please see the CEWES MSRC PET Technical Report 99-09.

off-campus

Concurrent Programming with Pthreads

Drs. Clay Breshears and Henry Gabb visited the Center for Research on Parallel Computation at Rice University on 29-31 April to teach a two-day course entitled "Concurrent Programming with Pthreads." Pthreads is a standard UNIX library for thread programming. On symmetric multiprocessors, Pthreads programs may run in parallel. Pthreads may be combined with Message-Passing Interface (MPI) to take advantage of distributed, shared-memory cluster architectures. The course focused on a core of Pthreads functions and the application of Pthreads to scientific programming.

This course is available on request. Interested parties should contact the Customer Assistance Center (800-500-4722, info-hpc@wes.hpc.mil).

Specialty Workshop on Adaptive Grids


A Specialty Workshop on Adaptive Grids was held 21 March 1999 at The University of Texas at Austin (the CEWES PET partner for computational structural mechanics) and was hosted by Professors Graham Carey and Tinsley Oden. The purpose of the meeting was to discuss current issues related to error analysis, adaptive grid methodology, and representative applications such as those arising in nonlinear finite element modeling of problems in computational mechanics.

Dr. Alan Stagg of the CEWES Computational Science and Engineering Group (CS&E) gave a presentation entitled "Developing a Parallel, Adaptive Grid Capability for a New DoD Multidisciplinary Flow Solver." Dr. Stagg described the current software development effort to enable parallel grid adaption for applications utilizing unstructured tetrahedral and/or triangular grids. The presentation focused on the methodology used to parallelize the grid refinement and coarsening algorithms and the data structures developed to simplify implementation. Dr. Stagg stressed the importance of considering all primary components such as refinement, coarsening, and load balancing when developing the parallel algorithms and data structures. He also presented preliminary results indicating high efficiency of the parallel adaption scheme.

The grid adaption software described in the presentation is being developed in collaboration with Jackie Hallberg of the CEWES Coastal and Hydraulics Laboratory and is being tested and demonstrated in the environmental quality management CHSSI code ADH. For more information on this work, please contact Alan Stagg (stagg@rusty.wes.hpc.mil).

CEWES MSRC

visitors



Sen. Cochran welcomes WES employees and visitors to the dedication of new the ITL expansion

The Honorable Thad Cochran, the senior United States Senator from Mississippi, was on hand for the dedication of the 32,500 square-foot expansion to the ITL building in December 1998. Senator Cochran welcomed the more than 200 WES employees and visitors to the ceremony before the high-tech ribbon-cutting ceremony. After the dedication, visitors were able to tour the ITL building and see some of the technologies within (see related story on page 15).

In April 1999, Dr. Robert E. Foster and Mr. Robert Boyd of the Office of the Secretary of Defense visited ITL and the CEWES MSRC. Dr. Foster is the Director of Bio Systems for the Director of Defense Research and Engineering and reports to the Deputy Under Secretary of Defense (Science and Technology). He is responsible for the coordination and oversight of the DoD biomedical, environmental, chemistry, and civil engineering programs.



Dr. Foster (left) gets a tour of the CEWES MSRC computer facility by Brad Comes, Director



Challenge user Dr. Manry with his workhorse, the CEWES IBM SP

Dr. Charles W. Manry, Jr., of the Space and Naval Warfare Systems Center in San Diego, CA, visited the CEWES MSRC in March of 1999. Dr. Manry is an engineer in the Electromagnetics and Advanced Technology Division and is working with the CEWES MSRC on a DoD Challenge Project studying Integrated Topside Design (ITD). The ITD project is working on new ways to meet requirements for advanced communications capability with greater imagery and data transfer capacity, while also working to meet aggressive signature reduction goals for U.S. Naval surface combatants (see related story on page 16).

Parallelized Solvers for Sparse Linear Systems on the SGI Origin2000


Tom Oppe, Ph.D.

The selection of a solver for sparse sets of linear equations is often of critical importance in migrating codes to one of the CEWES MSRC parallel platforms. For the SGI Origin2000 (O2K) platform, there are at least three solvers that have already been parallelized. They are the SGI proprietary solvers PSLDLT and PSLDU and the public domain solver SuperLU from the University of California at Berkeley.

These solvers are all direct solvers, meaning that the solution is obtained in a predictable number of steps through variants of the Gaussian elimination algorithm, as opposed to iterative solvers, which obtain the solution through a succession of approximations. The choice of the “best” solver is often difficult to determine in advance, and there are situations where either direct or iterative solvers are indicated. Generally, direct solvers are a good choice if the bandwidth of the matrix corresponding to the linear system is small relative to the system size or if the matrix is ill-conditioned. All three solvers considered here have reordering strategies to reduce the bandwidth.

The PSLDLT solver solves linear systems with symmetric real-valued matrices, while PSLDU can handle matrices that are nonsymmetric in coefficients but symmetric in their nonzero structure. Of course, a matrix can always be forced to have a symmetric nonzero structure by storing some additional matrix coefficients as zeros. The user must store the nonzero coefficients of the matrix in a compressed column sparse storage

scheme called the Harwell-Boeing format. There are separate steps for preprocessing (re-ordering and symbolic factorization of the matrix), numerical factorization, and solution. This is helpful if the user is solving a sequence of problems whose matrices are the same or have the same nonzero structure, in which case the numerical factorization or preprocessing steps may only need to be done once. This feature is especially important in parallel environments since only the numerical factorization step has been parallelized. These solvers do no pivoting during the factorization, and only one right-hand side can be treated at a time with each call. Please see the PSLDLT and PSLDU man pages on the O2K for more details.

The SuperLU (single-threaded) and SuperLU_MT (multi-threaded) solvers are also available on the CEWES O2K and have more flexibility. The matrices can be unsymmetric in nonzero structure as well as in coefficients. There are routines for all four precisions (real, double precision, complex, and double complex), and several right-hand sides can be treated with a single call. There are routines to equilibrate (i.e., scale) the matrix, estimate its condition number, reorder the matrix, compute a factorization, and solve the linear system. The factorization uses partial pivoting for improved numerical stability, and equilibration can improve the matrix’s condition number beforehand. Finally, there are routines to refine the obtained solution and compute error bounds. 

More information:

www.wes.hpc.mil/news/sw_tools/

Virtue Software

Clay P. Breshears, Ph.D.


To address the problem of analyzing large trace files resulting from scalable parallel codes, the CEWES MSRC has begun experimenting with 3-D virtual reality immersion using the Virtue system developed at the University of Illinois. By using Virtue on an Immersa-Desk or a graphics workstation, the user can view state changes and processor-to-processor communication behavior for the tasks making up a parallel execution in the form of a 3-D time tunnel display that can be interactively manipulated and explored. The dynamic call

tree for a process can be viewed as a 3-D tree with sizes and colors of nodes representing attributes such as duration and number of calls. The amount of data that can be visualized at one time using 3-D virtual immersion is an order or two of magnitude more than with a 2-D display.

Virtue also has multimedia tools that allow remote collaborators to view and manipulate the same Virtue display (although scaled down and with fewer

navigation capabilities) from their desktop workstations. Users can also attach voice annotations to specific objects within the Virtue display for asynchronous collaboration.

Virtue has been successfully installed on the CEWES MSRC Scientific Visualization Center systems. A converter program that translates VAMPIRtrace files (the standard trace tool installed on all CEWES scalable parallel systems) into the data format understood by Virtue has been written, and experimentation with Virtue to view large VAMPIRtrace files produced by CEWES MSRC applications has begun. In collaboration with the Virtue team, CEWES is developing new Virtue features, such as source code annotations and other labeling to increase the effectiveness of the displays and better relate the visual phenomena to the application. Look for additional developments and capabilities in future CEWES MSRC publications. [\[link\]](#)



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

The End of an Era

Mary Gabb

June 1997 was a pivotal month for the CEWES MSRC. We seized the opportunity to more aggressively pursue parallel computing, an initiative that supports the goals

of the High Performance Computing Modernization Program (HPCMP).

The Cray C90 is scheduled for decommissioning at the end of September 1999, but the CEWES MSRC staff has been preparing for this since 1997. Almost 700 users were working on the C90 in 1998. The CEWES MSRC supports a special team known as the Computational Migration Group (CMG) to assist C90 users who wish to migrate to the latest and greatest high performance computing (HPC) technology: scalable, parallel programming.

The change from vector to parallel computing has been a slow process in the HPC community, which has experienced paradigm shifts with surprising frequency during its short lifetime. Remember when 3.0 GFLOPS was considered state of the art? That was 1988, the year of the first Supercomputing conference.

The Cray C90 was first introduced to the supercomputing community in 1992, the same year that the HPCMP was established. "For its day, it had a lot of memory (8 GBytes)," says Dr. Henry Gabb, Director of the CMG. "At the time, the largest problems could only be processed on a machine that had that much memory. But by today's standards, 8 GBytes is not a lot."

Dr. Gabb added, "But its main advantage was its processor. The processor in a C90 has special hardware for vector processing, which allows it to pull large pieces of data from memory and process it simultaneously."



In a move toward parallel computing, the C90 is scheduled to be decommissioned 30 September 1999.

CEWES Computing: You've Come a Long Way

IBM 650	August 1957 - November 1962
GE-225	August 1962 - October 1968
GE-437	August 1968 - June 1972
GE-635	March 1972 - July 1981
T1-ASC	January 1979 - November 1980
DPS 1	July 1981 - December 1984
Honeywell DPS 8/70	December 1984 - October 1991
Cray Y-MP	November 1989 - March 1997
CDC 962	April 1990 - present
nCUBE	July 1992 - October 1995
CD 4000	January 1993 - March 1996
Cray C90	July 1993 - September 1999
Cray J916 Mass Storage Facility	March 1994 - present
Sun 2000	March 1995 - present
Sun 6000	June 1996 - present
SGI POWERCHALLENGE	February 1997 - June 1998
IBM SP	March 1997 - present
Cray T3E	April 1997 - present
SGI/Cray Origin2000	May 1997 - present



The CMG team members' specialties include mathematics, physics, chemistry, and computer science, all with a strong background in computational modeling.



“Memory allows you to tackle much larger problems. So if you have a memory-intensive problem and it’s going to become more complex, you’re going to need a parallel machine for the memory to do the calculations.”

**- Dr. Henry Gabb,
Director, Computational
Migration Group (CMG)**



Tipping the scales toward scalable computing

The advent of user-friendly parallel computing sparked the shift to what is now considered “supercomputing.” Dr. Gabb explained, “There are desktop computers that can get the same performance as a C90, if coded correctly. The hardware technology has come a long way; compiler technology has come a long way, such that a good workstation can do almost as much as a Cray C90 can for a fraction of the price.”

So who would want a C90? “It’s still a good machine,” Dr. Gabb emphasized. “It is still a powerful computer. And there are some users who have highly vectorized codes - codes that are highly optimized for the C90. It’s tough to beat vectorized code running on a C90, even for a parallel computer.”

The advantage of the C90 was that you could get good performance without a lot of effort. C90 users knew that they would get the best performance through vector processing, so they wrote their programs accordingly. The Cray compilers for the C90 do a lot of the optimization for you, including vectorization. Dr. Gabb warns, “There aren’t many compilers that will take your program and make it run on a large, scalable, parallel system. We’re still years off from that in terms of compiler technology.”

With parallel computers you get much better performance—orders of magnitude decreases in run times, much larger memory, but at the expense of complexity. The programmer is required to do a lot more to his program to get that performance.

Dr. Gabb has watched parallel computing mature into an elegant, feasible system. “Parallel computing has come a long way. Parallel computers have been around for at least 30 years, but they were too difficult to program. There were no standard parallel programming models. You wrote code, you parallelized your code for one computer, the company went out of business and your code was useless on any other machine.

A Cray C90 processor, with its special cooling systems and hardware, power consumption, and price, is a significant investment. Commodity, off-the-shelf processors like those in the SGI O2K, Cray T3E, or the IBM SP are about four to six times slower than a single CRAY C90 vector chip. But they are very cheap – cheap enough to place several hundred in a computer. The price: performance ratio is clearly in favor of parallel machines.

So just getting into that level of processing was prohibitive. People didn’t have the skills and they didn’t have the time to write a parallel code that wouldn’t be portable. So getting the kind of performance that the vendors promised was nearly impossible.”

Today, parallel programming models have become standardized. Several are recognized as industry standards: Message Passing Interface (MPI), High Performance Fortran (HPF), OpenMP, Pthreads.

The Bottom Line

What concrete advantage does parallel programming offer? Dr. Gabb experienced the golden eggs laid by this goose first-hand. He was part of a team that won the “Most Effective Engineering Methodology” at the SC98 challenge competition. The CEWES MSRC team used OpenMP and MPI to optimize a program (CGWAVE) used by the Navy in modeling harbors. When the Navy wants to evaluate a specific harbor for dangerous spots, CGWAVE does the modeling calculations for them. Prior to parallelization, those calculations would have taken between six months and one year. At SC98, the team had calculation times down to less than 72 hours using a combination of OpenMP and MPI. With further optimization, those calculations are now being performed in less than a day.

Dr. Gabb offered other examples where parallel programming offers significant scientific advantages: “A materials science code that might run on the C90 for a few thousand atoms can now simulate millions of

atoms. It's also possible to tackle linear systems on the order of 40,000 simultaneous equations."

Parallel Computing Expertise

So parallel programming can open up whole new worlds of possible problems that can be addressed computationally, with ever-increasing complexity. The methods are becoming standardized. What does the user need to parallelize his code?


At the CEWES MSRC, the CMG is in place to migrate the current C90 users to the parallel machines. The CMG not only assists users in performing the actual parallelization but also offers support and consulting so users can parallelize their own codes in the future.

In addition to parallel coding assistance, the CMG provides services that assist users in migrating code to parallel machines themselves: FAQs on the Web site for all of the machines (see "More Information" at the end of this article), teaching a course on a particular parallel programming tool or method, answering ques-

tions from the Customer Assistance Center on parallel programming.

Users often choose to migrate on their own for one of two reasons: they are not aware that CMG exists, or the user knows more about his code than the CMG does. According to Dr. Gabb, "We're here to provide assistance, mainly because the user may not know everything that we know as a group. We've been working on these machines for a long time. We're competent on all the parallel systems and move between them, depending on the project. We're experienced with all of the compilers, which is also important. It takes time for a user to get accustomed to a particular compiler."

The CEWES MSRC training courses in parallel programming can be specialized to focus on any of the models with real user code. They are typically offered here at CEWES because we have the facilities for "bring-your-own-code" workshops. In some cases, a CMG member, along with a representative from the Programming Environment and Training (PET) Program will teach the class at an offsite location, provided that enough interest is generated. PET leads are also called upon when the user's code is in their area of expertise (e.g., computational fluid dynamics, scalable parallel programming tools).

The CMG is a mix of scientific and computational expertise, synergized to offer cutting-edge performance to all CEWES MSRC users. As a group, the CMG has expertise in all of the standard parallel methods (Pthreads, MPI, HPF, OpenMP). 

"It takes that diversity of background to do this kind of work. No one person can know as much as everybody in the group. Working separately, we would fail."

- Dr. Henry Gabb, CMG

More information:

The History of Supercomputing Display from SC98: www.supercomp.org/history/1992/index.html

Hardware Descriptions for On-site Supercomputers: www.wes.hpc.mil/hardware

Information on CMG: www.wes.hpc.mil/news/parallel

User Guides and FAQs for On-Site Parallel Computers: www.wes.hpc.mil/documentation/index.html

NASA Liaisons Collaborate on Key Projects

Judith Utley and Michelle Morgan Brown

In 1998, NASA and the DoD High Performance Computing Modernization Program's Major Shared Resource Centers (MSRCs) began a collaborative effort. The goal of this collaboration is to transfer the technology and the experience gained during the Phase I NASA MetaCenter to establish a similar MetaCenter between the Aeronautical Systems Center (ASC) MSRC, in Dayton, OH, and the U.S. Army Corps of Engineers Waterways Experiment Station (CEWES) MSRC in Vicksburg, MS, with possible expansion to other DoD HPC resources.

The earlier effort for the NASA MetaCenter, which was between the NASA Ames and NASA Langley Research Centers, resulted in a production MetaCenter joining two cooperating IBM SPs, one at each center. This allowed approximately 500 users to reap the benefits of a successful MetaCenter. The technology developed as well as the challenging nontechnical issues encountered with the NASA MetaCenter as it moved from an idea to full production in 1996, running successfully until 1998, is now being shared with the two MSRCs. The software user for this venture is the Portable Batch System (PBS). This metacentring technology can balance the job mix, decrease time to solution, provide researchers with a wider range of resources, enable large jobs to run more frequently, and automatically migrate jobs between available resources.

A technical team from each MSRC is working with NASA to implement this technology. In addition, a team from each center representing the accounting and utilization groups is also working to make this happen. Leading this effort for the CEWES MSRC is Judith Utley, who led the original effort for NASA Langley Research Center. Utley, a senior systems analyst for MRJ Technology Solutions, has nearly 16 years of experience in computer systems work.

In addition to the MetaCenter work, the CEWES MSRC supports an initiative to provide a common user environment across all HPC systems. As part of this initiative, a common queueing system for all CEWES HPC production platforms is required. The portable batch system (PBS) has been selected to provide this capability. NASA liaison James Jones is on-site performing the port of PBS to the CEWES Cray T3E.

When implemented, users will have the ability to submit a batch job to any HPC system and to run on any HPC system using the same interface. This interface provides a transparent environment that allows users to concentrate on work rather than learning multiple batch systems. Porting PBS to the T3E brings this closer to reality.

Jones is also collaborating with CEWES personnel to write and maintain PBS schedulers for all CEWES HPC systems. Jones is a PBS systems analyst with MRJ and a member of the parallel systems group at NASA Ames. Jones has been working with PBS schedulers for several years.

The NASA liaisons at the CEWES MSRC have facilitated numerous opportunities for technology transfer with the NASA team working on the NASA Information Power Grid (IPG), the next generation interconnection technology required to support the exchange of information across geographically distributed computing resources.



NASA Liaisons Judith Utley and James Jones

It Takes Two to TANGO

Mary Gabb

In our last issue of *The Resource* we introduced TANGO Interactive (or TANGO), a “web-based collaboratory system” developed by the Northeast Parallel Architectures Center (NPAC) at Syracuse University, a CEWES Programming Environment and Training (PET) partner. Using TANGO, the CEWES MSRC has participated in two for-credit, college-level courses (undergraduate and graduate), taught at Syracuse University and offered on-line through Jackson State University (Mississippi). But college classes will not be the only use for this cutting-edge technology at the CEWES MSRC.

“TANGO is one of the most exciting new tools that’s come along in a while,” says John Eberle, CEWES PET Training Coordinator. Mr. Eberle has been instrumental in not only getting TANGO up and running for

our staff, but also looking at future uses. He is the on-site expert.

The most significant advantage with TANGO is that it allows synchronous (two-way, real-time) communication. “For us, the big interest is to lessen the dependence on classroom seats, and accessing those 500+ users who are not physically here,” Mr. Eberle added. “Since 1 January 1998, we had 338 people sign up for classes here. Remote students had to apply for travel fare and per diems for lodging and food.” Distance education can decrease travel budgets and can increase the likelihood of class attendance. Users can go to a class in their local environment, saving the time, trouble, and cost of travel.

The benefits of TANGO extend beyond the classroom. TANGO can be used for any two-way collaboration. Although the two parties are referred to as Instructor and Student, the names indicate who initiates the session (and therefore controls the content to be presented) and who “plugs in.” The Instructor status, however, can be relinquished to a Student at anytime during a session.

John Eberle’s role is to facilitate the use of TANGO by the CEWES MSRC staff who want to collaborate with outside partners. Several of the staff have taken advantage of this new capability.

TANGO has many tools to facilitate these partnerships. The Web Wisdom Object Manager allows the Instructor to accumulate data from different presentations. It places data from different presentations into an object-oriented database, where all data (e.g., PowerPoint slides, Excel graphs, JPEG files) are seen as “objects” and not as different file types. The computer can then interact with all types of data in the same way.

The second version of TANGO, due to be released in early 1999, offers new features that reach out to other potential users. A common debugging tool will allow the Instructor and Student to look at source (or computer) code simultaneously, find the problem, and perform test runs. For the CEWES MSRC Computational Migration Group (CMG), tasked with assisting users in migrating their codes to the parallel systems,

“Continuing life education — embrace it or be left behind.”

- Dr. Geoffrey Fox



In conjunction with the training component of the CEWES MSRC PET program, TANGO was developed at the Northeast Parallel Architectures Center, a center founded by Dr. Geoffrey Fox of Syracuse University. Dr. Fox has a long history in parallel computing, dating back to 1981. At that time, he was part of the original “Cosmic Cube” team that developed the first parallel computer at Caltech University.

John Eberle, the CEWES MSRC PET Training Coordinator, acts as a catalyst among potential TANGO converts, hoping to allay fears through educating people on the TANGO tools. He answers their questions: “How do I use it? What can it do for me? How do I make my data available to others?” With a 20-year career in training, John is keenly aware of what works in the classroom, and how that can translate into effective collaborations.

synchronous debugging is an attractive feature. Dr. Henry Gabb, Director of CMG, is looking forward to the second version: “This will be an extremely useful collaboration tool for us. With remote users, we can look at specific areas of code that are causing problems - together, in real-time.”

Great idea but how does it work?

“It’s a well-thought-out tool, but it takes getting used to,” says Mr. Eberle. “There’s a time management issue and a screen real estate issue.”

The technology of TANGO can be intimidating for those used to traditional teaching tools, such as slides, overheads, videos, PowerPoint presentations, and live classrooms. The Instructor is required to keep track of both the classroom and the hardware. And learning not to be overwhelmed by that is the most important aspect of learning TANGO, according to Mr. Eberle. A camera rests on top of the computer to view the instructor as he or she is speaking. The instructor, as in a live classroom, controls what information goes into a shared browser that everyone can see. For example, Instructors can use PowerPoint presentations, a drawing program, and a white board, as well as answer questions from students using live chat windows (Figure 1). In this way, the flow of thought is not interrupted, yet the instructor can tailor the lecture to a question on the chat window as they go.

Dr. Geoffrey Fox of Syracuse University led the movement in developing TANGO. Dr. Fox, who gave a talk at the CEWES MSRC on 16 September 1998, was visiting CEWES and Jackson State University (JSU) as part of his regular visits in support of the CEWES MSRC distance learning program. He interacts with the JSU curriculum mentor and technical expert on a regular basis to ensure smooth operations.

During his talk, Dr. Fox reviewed the process of creating a course on-line. He emphasized that the instructor has to do work ahead of time: design the curricula/content; define the architecture; assign an author for the material; organize student management (i.e., the forms of quizzes, tests, and grades); and determine the delivery method (asynchronous, synchronous, or collaborative). However, with the exception of architecture, these same planning decisions are made in the preparation of a traditional course.

For the Student, the requirements for using TANGO are less complicated. They can simply jump right in. The student needs to know how TANGO works (functioning more and more like Microsoft environments with each new release), and the location of the Instructor’s server.



Some institutions offer mirror sites so the students can access files locally, instead of waiting for downloads across the Internet.

The student session is controlled by what the Instructor does. Students also must adapt to frequently changing screen real estate. “If you aren’t expecting it and don’t understand why it’s happening, it can be disorienting,” explains Mr. Eberle.

A “Wish List” of features

Syracuse University is very interested in getting feedback on TANGO. With the imminent release of the second version, John Eberle plans to develop a “like/don’t like list” with users to submit to Dr. Fox’s team at Syracuse. Syracuse has been saying: Let’s make it useful. They want their product to meet some real need. Using the CEWES MSRC user base as an audience to improve the tool is a mutually beneficial situation.

“It’s an evolving piece of software,” Mr. Eberle emphasized. “The academic team at Syracuse (under Dr. Fox) is pushing the envelope with assistance from the DoD.”

And Dr. Fox thrives on new technology. He calls his philosophy “Continuing Life Education”—embrace it or be left behind.

“We have shown that we can reproduce traditional learning methods,” Dr. Fox remarked. And he added the following caveat, “This is great, but it’s a first experiment. It is not meant to replace ‘live’ learning. TANGO complements other teaching methods.”

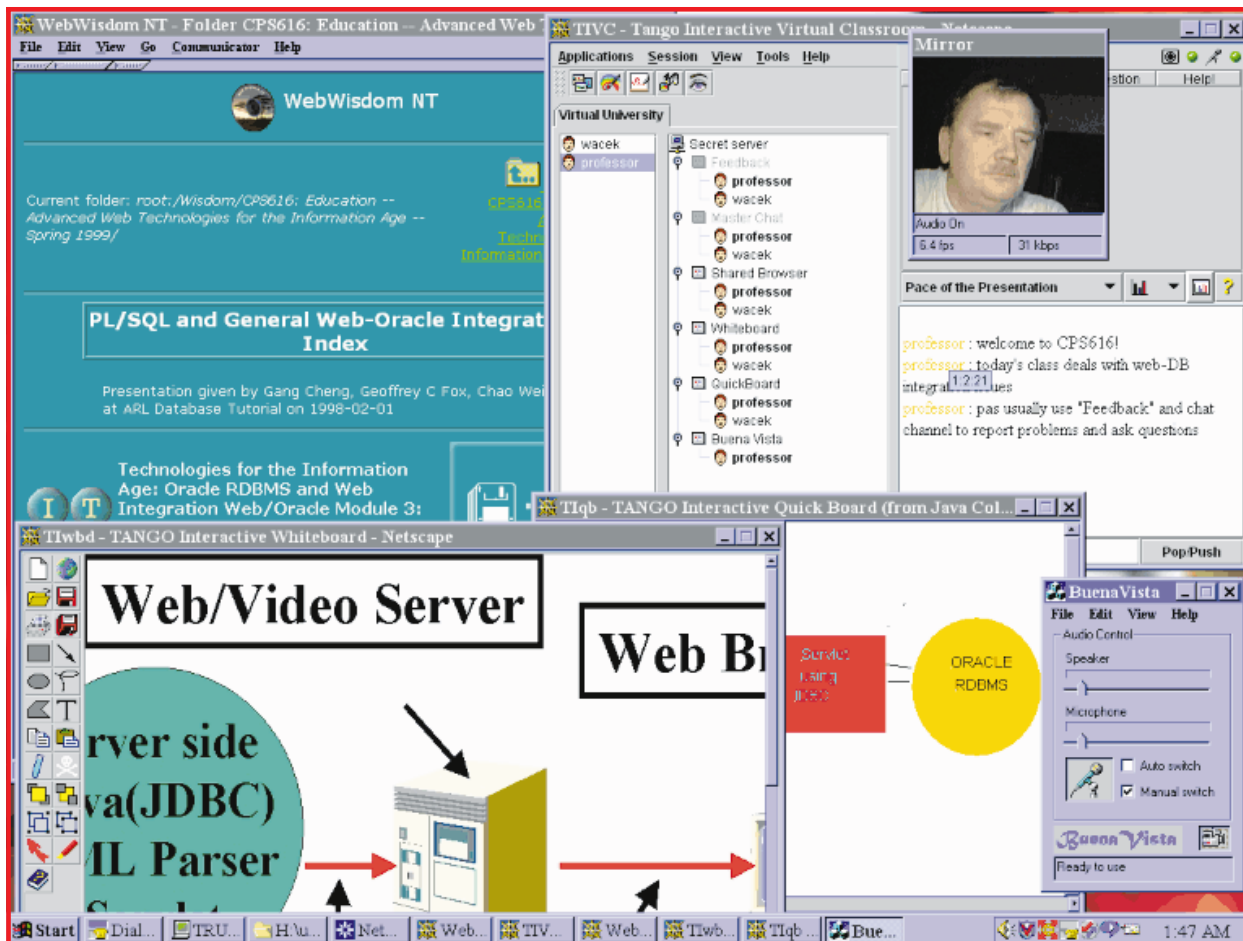



Figure 1. TANGO offers the sensory input of a live classroom: streaming audio/video, a white board, a drawing program, a live chat window, a shared HTML browser, shared XEmacs (a powerful UNIX text editor primarily used for code development), and, for PC platforms, MS Word, PowerPoint, Excel, and C++. For distance learning purposes, the “raise your hand” button allows students to “attract the lecturer’s attention or to signal utter confusion,” according to the TANGO Web site.

TANGO provides a cost- and time-effective method of delivering information—in academic, corporate, and government settings. Unlimited numbers of students can benefit from the classes offered by the CEWES MSRC, not just those students present in the classroom. For CEWES, TANGO represents a significant enhance-

ment to the return-on-investment made to deliver any single course or collaboration. It opens numerous avenues of on-line, real-time, two-way communication - with PET partners, DoD collaborators, and the user community. 

More information:

CEWES MSRC Training Information: www.wes.hpc.mil/msrc/training/

TANGO Web Site: www.npac.syr.edu/tango

ITL Dedicates New Addition

David Stinson

The day was not an ordinary one for the people of the Information Technology Laboratory (ITL). The huge tent covered much of a parking lot usually filled with cars. Underneath the tent stood rows of chairs waiting for the anticipated large audience. People, instead of normally going about their morning routine, were standing by ready to talk to visitors about the myriad of technology on display within the laboratory. The day was Wednesday, 9 December 1998. The event was the dedication of the new addition to the Information Technology Laboratory at the U.S. Army Engineer Waterways Experiment Station.

Since the founding of the Information Technology Laboratory in 1986, space has been in constant demand. ITL is the premier Department of Defense laboratory for the development and application of advanced information technology to military engineering and Army civil works mission areas. ITL's first real home was the Jamie L. Whitten Building, which was completed in November 1989. Over the next 5 years, ITL's dramatic program and capability growth led to the requirement of additional computing and laboratory space. The December dedication ceremony highlighted the completion of a \$5.3 million, 32,500 square-foot expansion to meet these growing demands.


The dedication ceremony started with a call to order by COL Robin R. Cababa, Commander of ERDC and Master of Ceremonies for the dedication. The National Anthem was led by the ITL Singers, and then welcomes and addresses were presented by local leaders and dignitaries. These included the Reverend David A.




Dr. Robert W. Whalin, Sen. Thad Cochran, Dr. N. Radhakrishnan, and COL Robin R. Cababa (foreground from left)

Elliot, III, pastor of a local church; Dr. Robert W. Whalin, then Director of the Waterways Experiment Station; the Honorable Gertrude A. Young, Mayor Pro Tem of the City of Vicksburg; and Dr. N. Radhakrishnan, Director, ITL. The event climaxed with an address from Senator Thad Cochran, the senior United States Senator from the State of Mississippi. The formal program ended with a high-tech ribbon cutting ceremony using laser-light "scissors."

Following the official ceremony, tours were conducted of the ITL building and its capabilities. Visitors received tours of the joint computing facility, home of the Corps of Engineers' CEAP Central Processing Center, the WES Computer Center, and the DoD High Performance Computing Major Shared Resource Center and Mass Storage Facility.

Additional tour stops included the Groupware facility, which allows scientists and engineers to collaborate using interactive electronic meeting technology; the Scientific Visualization Center, which demonstrated immersive technologies and DoD applications; the CADD/GIS technology center; technologies used for the production of visual aids; and library automation tools. Over 200 people attended the event — yet another milestone in ITL's rapidly growing infrastructure and mission. 



The ribbon-cutting ceremony was a high-tech affair. Cutters used laser-light "scissors" to dramatically cut the ribbon for the new expansion

Challenge Projects Up Close: Integrated Topside Design

Michelle Morgan Brown and
Charles W. Manry, Jr., Ph.D.

In fiscal year 1999 (FY99), the High Performance Computing Modernization Program (HPCMP) initiated Challenge projects as part of the ongoing MSRC mission to support the warfighter. These highly visible projects focus on complex DoD research problems that require large data sets and lots of computational power. One of the ongoing Challenge projects at the CEWES MSRC is the Integrated Topside Design (ITD) project, which was awarded 202,000 CPU-hours on the CEWES IBM SP system for FY99. The principal investigator on this project is Dr. Charles W. Manry, Jr., an electrical engineer with the Space and Naval Warfare Systems Center in San Diego, CA (SSCSD). The SSCSD supports the command, control, communications, computers, intelligence, surveillance, and reconnaissance missions of the U.S. Navy.

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

Designing and optimizing a complex electromagnetic environment is typically slow, expensive, and error-prone. It is only through the application of concurrent engineering and its associated simulation-based design environment that 21st Century Integrated Topside Designs can be implemented. Therefore, the need for advanced simulation, visualization, and optimization tools that exploit high performance computing (HPC) are critical to the required design tools. An Electromagnetic Interactions GenERalized (EIGER) framework has been developed. It incorporates a variety of numerical and analytical techniques into an efficient,



Figure 1. USS Radford with Advanced Enclosed Mast/Sensor (AEM/S) mast. (Courtesy U.S. Navy and Space and Naval Warfare Systems Center, San Diego)

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

Both SSCSD and CEWES MSRC personnel are working together to make demonstrations that show the capabilities of HPC and EIGER used to solve ITD problems (Figure 2). EIGER development is supported by multiple sponsors including the Department of Energy and is a DoD Common High Performance Computing Software Support Initiative (CHSSI) project. Advanced EM modeling has been, and will continue to be, critical to the success of these technology initiatives.

Please see related story in the CEWES MSRC technical journal, spring 1999. [\[link\]](#)




Figure 2. Artist conception of amphibious assault ship (LPD 17 class) with AEM/S masts. (Courtesy U.S. Navy and Space and Naval Warfare Systems Center, San Diego)

interview with... **Jamie Grantham**

Mary Gabb

Jamie Grantham is a Deputy Program Manager for Nichols Research Corporation, the primary integrator for the CEWES MSRC. He began his career as an electrical engineer working for Nichols on signal processing for strategic defense systems. He came to Vicksburg in 1991 as Office Manager of the new Vicksburg, Mississippi technical office. He was directly involved with the Joint Test Program at CEWES for signal processing and analysis of attack weapons systems. In 1996 he joined the CEWES MSRC, initially as the Director of Infrastructure, but his role has expanded to program management. He is also responsible for overseeing the CEWES Challenge Projects. We spoke with Mr. Grantham about the history of DoD Challenge projects at CEWES and our philosophy toward meeting Challenge user requirements.

Let's start with some background on the DoD Challenge projects. How long has the program existed and when did the CEWES MSRC join? HPCMP [The High Performance Computing Modernization Program] initiated the Challenge program in FY97. The MSRCs became involved at the same time because it's part of our charter. We can solve large, complex problems and the HPCMP recognized that. These are high priority projects and we support them through our supercomputing power as part of our support for the warfighter.

How are the Challenge projects defined? Each project has to fit into one of the 10 CTAs [Computational Technology Area]. DoD scientists and engineers submit proposals to the HPCMP. They are evaluated by a committee of peers and selected based on their contribution to DoD science as well as their computational requirements. The projects are solving complex problems that have large computational requirements. The target for the HPCMP is to use a minimum of 20% of the computational resources at each MSRC for Challenge-type problems. This year, the CEWES MSRC is delivering 31% of its capabilities for Challenge projects.

What kind of computing power does that mean for the CEWES MSRC? We currently handle 14 [Challenge] projects. There are 11 distinct projects, but some projects run on both the machines we've allocated for Challenge project use (the Cray T3E and the IBM SP). There are one million node-hours allotted for each machine. Last year, we had 650,000 hours on the T3E and 450,000 on the IBM SP. This year's demands have grown by a factor of two, but the minimum 20% total allocation of our resources hasn't changed. The MSRCs have been getting commensurate resource upgrades.


Have we been successful at meeting these computational requirements? How? Our record is stellar. We exceeded 100% of the computational requirements for the projects for the last two years. We had four or five projects that solved their problem prior to using all of their allocated cycles, but others needed a bit more and we provided it. How did we do it? One of the ways we've been successful is the establishment of Challenge support teams with a single point of contact (POC) for each Challenge project. Single POCs are important because they guarantee a quick response to the project PI [principal investigator] and users. The POC is the avenue into the MSRC, to address unique processing demands or support services.

Who constitutes the Challenge support teams? Each team is specially formed based on our Challenge project user survey, which is sent to each Challenge project PI. The survey assesses requirements for access, data archival, disk space, network bandwidth, and scientific visualization. Based on this survey, a support team is formed. For example, a project with a heavy scientific visualization requirement will have someone from the SciVis lab on the support team. Other functional support could include computational science, application support, systems analysts, and CAC [Customer Assistance Center]. Some of our team members have visited PIs at their home sites to assist with particular problems or to hold workshops.

Are there any logistical problems? It is sometimes difficult to balance the computational requirements of the Challenge users with our regular user base. We support over 700 users from the DoD and its related contractors. There are approximately 40-50 challenge users. We make people here aware of the importance of the Challenge projects. We also have dedicated people monitoring and manipulating the job queues for both machines. Comments from our users have been very positive. They are appreciative of the commitment we've made to their projects. Challenge users are sometimes not familiar with our environment. We recognize that and work hard to try and work out the unknowns before the project starts - working out methods for archiving data, establishing modified queue structures, moving data.

Conclusion

As Mr. Grantham stressed, "The Challenge projects are a high priority for us. We intend to continue this level of support for our Challenge users in the coming years. Challenge projects are the leading edge of our HPC demands and Challenge users are our 'gold card' customers."



More information:

SSCSD and its support of U.S. Naval missions: <http://www.spawar.navy.mil/sandiego/welcome.page>

Integrated Topside Design efforts: <http://bobcat.spawar.navy.mil/d85/index.htm>

DoD HPC CHallenge Projects: <http://www.hpcmo.hpc.mil/Thdocs/Challenge/index.html>

Scientific Illustration in HPC

Christine E. Cuicchi and Alex R. Carrillo

In today's computational environments, it is common for teams of specialists to focus on a given problem. These teams usually consist of professionals in the fields of engineering or science, computation, and visualization. An often overlooked, but invaluable member to such a team is the scientific illustrator.

The target audience of most research is wide and varied. The public, managers, nontechnical customers, the media, and even scientists from other disciplines can all be a part of a technical audience. Where visualization promotes understanding of computational data to the scientist, incorporating graphic arts into a visualization can help promote understanding of a visualization to such an audience. By precluding visualizations with animations of the phenomena being investigated, by incorporating visualizations into a real-world background, or even by supplementing visualizations with interacting components not normally seen, a link can be drawn between computational science and real-world circumstances. This can be especially important for a nontechnical audience.

At the U.S. Army Engineer Waterways Experiment Station, graphic arts have been incorporated into a range of computational projects during the last eight years. The impact of adding "artistic aides" can be seen in the following two examples - a military site cleanup project and a wave motion modeling project.

The military site cleanup project simulated the spread of a contaminant plume over time for a DoD storage site. Figure 1 shows an illustrator's rendering of the storage site. Figure 2 shows simulation results (the plume) with the addition of site buildings. The combination of Figures 1 and 2 gives the audience a direct link between the simulation results and the phenomena being investigated.

The wave motion modeling project computed the effects of 293 wave components on the sea surface at the Ponce Inlet in Florida. Figure 3 shows total sea surface amplitude. Colors indicate maximum vertical distance between wave crests and troughs. The same data with the conceptual beach and jetty added are shown in Figure 4. In both figures, the ocean "walls" are conceptual as well, giving the visual perception of water depth by joining the sea surface to the Ponce Inlet bathymetry. The addition of the jetty also





Figure 1. Artist illustration of a generic storage site. (Courtesy Randy Kleinman, CEWES MSRC scientific visualization illustrator.)

added perspective during a 3-dimensional stereo fly-through.

The Scientific Visualization Center at the CEWES MSRC supports an interdisciplinary team of engineers, computer scientists, and scientific illustrators who specialize in working with MSRC users to transform their raw data into scientific visualizations and broadcast-quality presentations. If you are in need of these services, please contact the CEWES MSRC Customer Assistance Center. 




Figure 2. Scientific visualization of contaminant plume. (Courtesy Richard Walters and Kent Eschenberg, scientific visualization specialists; Randy Kleinman, illustrator; and Jeff Holland, Tom McGhee, Jerry Lin, and David Richards, scientists.)




Figure 3. Scientific visualization of Ponce Inlet wave data. (Courtesy Christine Cuicchi, computer engineer; Randy Kleinman, illustrator; and Zeki Demirbilek, scientist.)




Figure 4. Ponce Inlet wave visualization with conceptual beach and jetty. (Courtesy Christine Cuicchi, computer engineer; Randy Kleinman, illustrator; and Zeki Demirbilek, scientist.)

UNIX tool tips

Alex R. Carrillo

GIMP

The GIMP (GNU Image Manipulation Program) is a freely distributed, powerful software tool for image composition, processing, and manipulation. Similar to Adobe's PhotoShop program, the GIMP's capabilities range from a simple paint program to an expert-quality image manipulation program.

The GIMP can be used as a quality photo-retouching program, an online batch processing system, a mass production image renderer, and even an image format converter. The GIMP is extremely expandable, as it was designed to be augmented with plugins (over 100 plugins are already available) and extensions. For enhancing or developing images for presentations and publications, most will find the GIMP a valuable tool.

Some of the GIMP's features include: a full suite of painting tools, including custom brushes and patterns; a complete set of transformation tools (rotate, scale, shear, flip, etc.); an assortment of image enhancement routines; support for most file formats; tile-based memory management so image size is limited only by available disk space; advanced scripting capabilities; layers and channels; and much more.

The GIMP was written by Peter Mattis and Spencer Kimball and was developed on X11 on UNIX platforms. Many other developers have contributed plugins, as well as support and testing. Currently, there is an OS/2 port and two preliminary win32 ports in development. GIMP can be downloaded from <http://www.gimp.org>.

XTar

XTar is a package for graphically viewing and manipulating files and directories that have been archived using the `tar` (Tape ARchive) command, and possibly compressed with either `gzip` or the standard UNIX compress utility. The main purpose of XTar is to allow the user to browse inside such a file without unpacking it. This allows a reduction not only in disk space utilization, but in inode count as well. This feature can also simplify the task of maintaining and moving large file structures.

The main window for XTar shows the contents of the open tar archive in a similar style to the UNIX long listing format (including file permissions, user ID/group ID, file size, date stamp, and filename). The entire tar archive can be extracted, or the extraction may be limited to individual files or directories by merely selecting them from the contents. In addition, double-clicking an individual file will open the file into an appropriate viewer, allowing the user to view the contents without having to extract the file. Users can use the default viewers, or configure XTar to launch their own preferred programs for viewing various file types (i.e., launching `xv` for viewing images or `showps` for viewing PostScript documents).

Additional XTar features include the ability to search for a regular expression to match against a filename in the archive, and the ability to set various attributes when files or directories are extracted, such as permissions, time stamp, and ownership.

XTar is freely available but requires an ANSI C compiler and the Motif toolkit to build. Xtar can be downloaded from <http://www.csc.liv.ac.uk/~rik/xtar/>.

*1999 training schedule**

July

Open MP and Pthreads
Ensignt for CFD and CSM Applications

August

Workshop on Parallel Algorithms
Distance Training Workshop

September

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

October

Grid Generation and Adaptive Grids
IBM POWER3 SP Parallelization Workshop

December

Using the SGI Origin2000 for Code Development and Analysis

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

CEWES MSRC Newsletter

Editorial Staff

Interdisciplinary Coordinator

Mary L. Hampton

Outreach Coordinator

Michelle Morgan Brown

Chief Editors

Mary Gabb

Michelle Morgan Brown



CEWES MSRC Customer Assistance Center

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

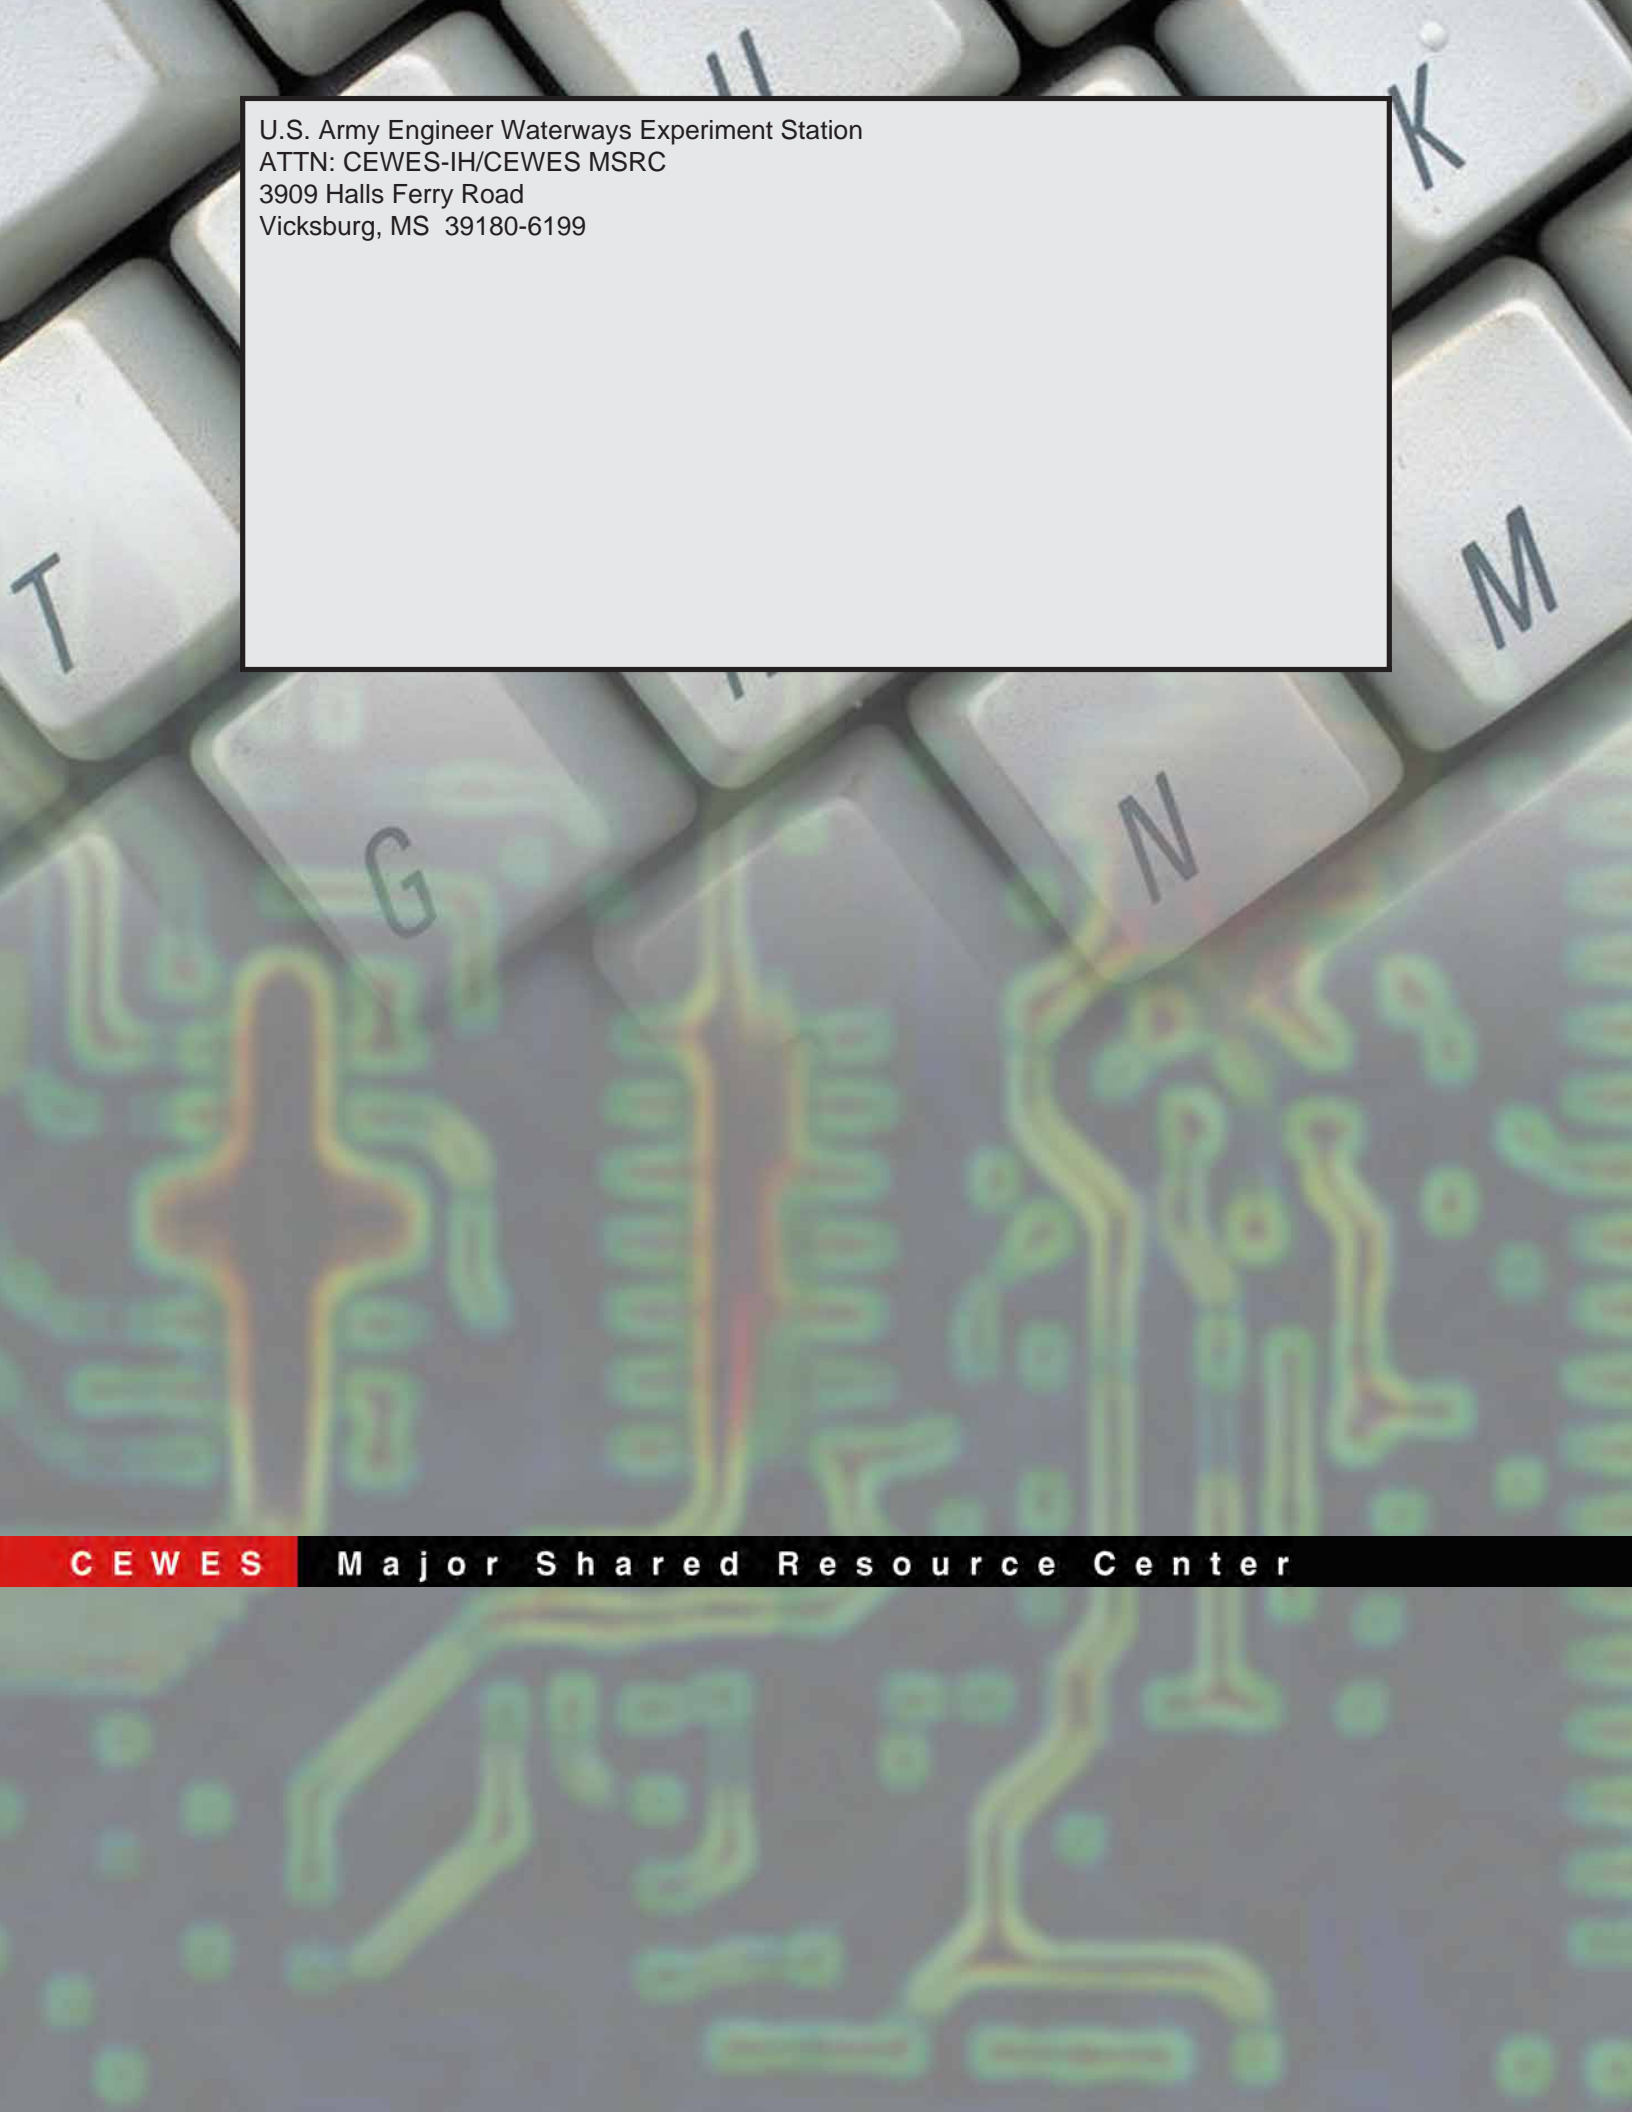
E-mail: info-hpc@wes.hpc.mil

Telephone: 1-800-500-4722

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

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

Approved for public release; distribution is unlimited.

The background of the entire page is a close-up photograph of a computer keyboard. The keys are white with dark lettering. Overlaid on the keyboard is a semi-transparent image of a green and yellow circuit board. In the upper center, there is a white rectangular box with a black border containing contact information.

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

CEWES

Major Shared Resource Center