




Navigator



NAVO MSRC

SPRING 2003



News and information from...

The Naval Oceanographic Office Major Shared Resource Center

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 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE NAVO MSRC Navigator. Spring 2003				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) Naval Oceanographic Office (NAVO),Major Shared Resource Center (MSRC),1002 Balch Boulevard,Stennis Space Center,MS,39522				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 Same as Report (SAR)	18. NUMBER OF PAGES 32	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			



The Director's Corner

Steve Adamec, NAVO MSRC Director

Changing to Better Serve You

In the past several months we have seen substantial change and progress here at the NAVO MSRC. We've brought on board our new MSRC technical services provider, Lockheed Martin (more in the article that follows on page 4), who brings substantial expertise and enthusiasm to their support of this MSRC and the DoD HPC Modernization Program (HPCMP). We're also completing several major Center enhancements, designated as Technology Insertion for Fiscal Year 2003 (TI-03), across several major technology areas within the MSRC. These include substantial upgrades to the IBM POWER4 HPC (MARCELLUS) system, the Remote Storage Facility (RSF), and the internal MSRC networking capability. When complete, these enhancements will provide almost 10 teraflops of aggregate peak computing capability with commensurately balanced storage and networking capabilities. This enormous computational capability will continue to enable unparalleled advances in the DoD science and technology areas served by the HPCMP.

As I've mentioned in past issues of the *Navigator*, we recognize that it is critically important for us to redouble our efforts in assessing and implementing common user environments, practices, and tools within and across the Centers. Your individual and collective user feedback through the User Advocacy Group makes it clear that you consider this to be one of your highest priorities for us. In response, we've formed new internal teams whose primary goal is to strengthen our linkage and participation with the HPCMP Programming Environment and Training (PET) program elements that emphasize user environment, tools, and productivity.

My staff and I look forward to seeing you in June at the 2003 HPCMP Users' Conference in Bellevue, Washington. As always, please take every opportunity to let us know how we can better serve you. Your feedback is critically important to us and to the HPCMP.

ABOUT THE COVER:

This image shows the temperature variable in a dataset created by a computational model run on the IBM POWER4 (MARCELLUS) in support of the Airborne Laser Challenge Project II. The data were visualized using Alias|Wavefront Maya 4.5 on a Windows 2000 workstation. A volumetric surface rendering technique was used for data elements where temperature was in the top 20 percent of the data range, 0.8 to 1.0. Temperature data values from 0.65 to 0.8 were rendered using a volume cloud technique. See "Data Visualization with Alias|Wavefront Maya 4.5," page 22, for further information about how this image was created.

**The Naval Oceanographic Office (NAVO)
Major Shared Resource Center (MSRC):
Delivering Science to the Warfighter**

The NAVO MSRC provides Department of Defense (DoD) scientists and engineers with high performance computing (HPC) resources, including leading edge computational systems, large-scale data storage and archiving, scientific visualization resources and training, and expertise in specific computational technology areas (CTAs). These CTAs include Computational Fluid Dynamics (CFD), Climate/Weather/Ocean Modeling and Simulation (CWO), Environmental Quality Modeling and Simulation (EQM), Computational Electromagnetics and Acoustics (CEA), and Signal/Image Processing (SIP).

NAVO MSRC
Code N7
1002 Balch Boulevard
Stennis Space Center, MS 39522
1-800-993-7677 or
msrchelp@navo.hpc.mil

NAVO MSRC Navigator

www.navo.hpc.mil/Navigator

NAVO MSRC Navigator is a biannual technical publication designed to inform users of the news, events, people, accomplishments, and activities of the Center. For a free subscription or to make address changes, contact NAVO MSRC at the above address.

EDITOR:
Gioia Fumess Petro, petrogio@navo.hpc.mil

DESIGNERS:
Cynthia Millaudon, cynmill@navo.hpc.mil
Kerry Townson, ktownson@navo.hpc.mil
Lynn Yott, lynn@navo.hpc.mil

Any opinions, conclusions, or recommendations in this publication are those of the author(s) and do not necessarily reflect those of the Navy or NAVO MSRC. All brand names and product names are trademarks or registered trademarks of their respective holders. These names are for information purposes only and do not imply endorsement by the Navy or NAVO MSRC.

Approved for Public Release
Distribution Unlimited

Contents

The Director's Corner

- 2 Changing to Better Serve You
- 4 A New Teammate Joins the NAVO MSRC

Feature Articles

- 5 Lattice-Boltzmann Large-Eddy Simulation of Turbulent Jet Flows
- 9 High Performance Computing and Simulation for Advanced Armament Propulsion
- 16 Clear Air and Optical Turbulence in a Jet Stream in the Airborne Laser Context

High Performance Computing

- 11 Largest NAVO MSRC System Becomes Even Bigger and Better
- 11 Using the smp Queue on MARCELLUS

Programming Environment and Training

- 13 Environmental Quality Modeling Activities Under PET
- 24 NAVO MSRC PET Update
- 24 PET Distance Learning: Ready to Serve 24/7
- 25 A Consistent, Well-Documented Computational Environment for the DoD HPC Centers
- 27 PET Climate/Weather/Ocean (CWO) Modeling and Simulation—A Brief Review

Scientific Visualization

- 22 Data Visualization with Alias|Wavefront Maya 4.5

The Porthole

- 23 Visitors to the Naval Oceanographic Office Major Shared Resource Center

Navigator Tools and Tips

- 29 Programming TotalView and Vampir

Upcoming Events

- 31 Conference Listings

A New Teammate Joins the NAVO MSRC

Linda Wise Pyfrom, Program Manager, Lockheed Martin at NAVO MSRC

On 15 January 2003, Lockheed Martin (LM) became the newest member of the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) team. LM is excited about providing Technical Operations and User Support Services to the NAVO MSRC and supporting the NAVOCEANO High Performance Computing (HPC) team.

LM brings a wealth of HPC expertise to the NAVO MSRC through leading research into, and participation in the development of, next-generation HPC systems. LM provides HPC hardware and software services for large-scale computational users and utilizes HPC resources in the design of advanced technology products. Whether producing HPC hardware with a five order-of-magnitude performance improvement at Sandia National Laboratories or linking the LM team with customer laboratories in real-time simulations for the Joint Strike Fighter Program, LM is an innovative member of the HPC community.

LM brings a highly experienced group to the NAVO MSRC team that is honored to work with the Operational and Research users of NAVO MSRC services. The LM team is committed to continuing the evolution of the NAVO MSRC HPC capabilities in the 21st century.

Key management team members include:

LINDA WISE PYFROM - LM PROGRAM MANAGER

Linda Wise Pyfrom brings to the NAVO MSRC team extensive experience in large-scale Information Technology (IT), strategic direction, program management, engineering, and process implementation for commercial, defense, and civil government customers.

She has served as Director of Information Technology and Director of Information Systems for LM in support of NASA, the Navy, and other government customers. As Program Manager, Ms. Pyfrom led a 186-person team in the development of the Naval Standard Integrated Personnel System (NSIPS), the newly operational system supporting Navy active service personnel and retirees. Ms. Pyfrom has

provided hardware and software management and technical direction to global Fortune 500 companies.

Ms. Pyfrom authored and instructed the Trusted Software Methodology for the National Security Agency Cryptologic Center and assisted in the transition of this methodology to the Carnegie Mellon University Software Engineering Institute. She was Principal Systems Engineer in support of the Strategic Defense Initiative for Martin Marietta and General Electric. Ms. Pyfrom began her career as a Test Engineer for Ford Aerospace serving the U.S. Air Force at its Cheyenne Mountain Space Defense Operations Center.

CHARLIE ROBERTSON - MANAGER, TECHNICAL OPERATIONS

Charlie Robertson has more than 40 years of management, technical, and supervisory experience. He has served as site manager for HPC facility management services and site manager for military command and control software development projects. Most recently, he was Program Director for Technical Operations for the NAVO MSRC. Prior to that position, he served as the Program Director of the U.S. Navy Primary Oceanographic Prediction System (POPS).

JEFF GOSCINIAK - MANAGER, USER SERVICES

Jeff Gosciniak has 20 years of experience in the leadership of software development efforts, including the development of Highly Available Enterprise System Architectures. Most recently, he served as Manager, Information Systems Engineering and Security for LM on the Consolidated Space Operations Contract (CSOC) for NASA. In this position he served as the Chief Systems Architect for the CSOOnline computing infrastructure. Mr. Gosciniak also served more than 10 years for LM Aeronautics, where he played an instrumental role in the IT efforts that supported the LM win of the Joint Strike Fighter contract. Prior to joining LM, Mr. Gosciniak worked in the Technology Laboratories for the Eveready Battery Company.



**The new LM team
leaders (L-R):
Jeff Gosciniak,
Linda Wise Pyfrom, and
Charlie Robertson.**

Lattice-Boltzmann Large-Eddy Simulation of Turbulent Jet Flows

S. Menon and H. Feiz, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA

Sponsored by Army Research Office

Active control using fuel modulation (by employing embedded micro-synthetic jets inside fuel injectors) has been experimentally shown to be an effective approach to control combustion instability in gas turbine engines. Numerical simulations can help in the design cycle if the dynamics of the interaction between the actuators and the combustion process can be properly modeled. However, this involves resolving motion over a wide range of scales. For example, a typical fuel injector orifice can be as small as 1-5 millimeters (mm), and the embedded microscale actuators are even smaller. On the other hand, the typical scale of a combustor is around 30 centimeters (cm). The resolution requirement to resolve the microjets and the flow outside in the combustor is too severe for any single numerical method.

The Lattice-Boltzmann (LB) method, when combined with the conventional Finite-Volume Large-Eddy Simulation (FV-LES), has the potential to provide a collaborative resolution to this multiscale problem.

In this approach, the LB-LES approach is employed to resolve regions inside the microjets and fuel injector while FV-LES is employed elsewhere in the combustor. This article reports on the ability of the LB-LES approach to capture complex dynamics in jet flows in a computationally efficient manner. The coupled LB-LES and FV-LES multi-scale approach is currently being validated and will be described in the near future.

The LB method is considered an attractive alternative to conventional





Figure 1. Computational domain for the synthetic and forced jet simulations. In both cases, the flow at the entrance of the inlet pipe is forced at the same frequency. But, in the jet case only, a mean flow is also added to the inflow condition. The flow at the exit plane of the orifice evolves naturally in both cases.

finite-difference schemes because it recovers the Navier-Stokes equations and is computationally very efficient, more stable, and easily parallelizable. Additionally, the LB method solves a single continuous particle distribution (which is analogous to the particle distribution function in kinetic theory) in a lattice (or grid).

The introduction of the Bhatnager-Gross-Krook (BGK) single relaxation time model for the collision operator further simplifies the algorithm and eliminates the lack of Galilean invariance and the dependence of pressure on velocity. Solving the LB Equation (LBE) instead of the Navier-Stokes equation has three distinct advantages: First, due to the kinetic nature of the LB method, the convection operator is linear. Simple convection in conjunction with a

collision process allows the recovery of the nonlinear macroscopic advection through multi-scale expansions. Second, because the macroscopic properties of the flow field are not solved directly, the LB method avoids solving the Poisson equation, which is numerically difficult in most finite difference methods. Third, the macroscopic properties are obtained from the microscopic particle distributions through simple arithmetic integration.

In this model, a new, second-order, accurate Three Dimensional (3D) LB method has been developed using a 3D cubic lattice model with the 19-bit velocity discretization (used here to recover the Navier-Stokes equation). This model has been extended to deal with complex geometries and to include a variable grid without loss of accuracy.



Additionally, to enhance its applicability to high Reynolds number flow, an LES version of this model has been developed whereby a localized dynamic subgrid model is employed to compute an additional subgrid relaxation time in the BGK model of the LBE. The dynamic evaluation eliminates the need to specify any ad hoc parameters since all model coefficients evolve naturally as a part of the simulation.

To expedite the turnaround time, the LBE-LES solver is implemented in parallel using the Message Passing Interface (MPI). The computational efficiency of the LB-LES solver is considerable and achieves 4.42×10^{-9} Central Processing Unit (CPU) seconds per time step, per grid point, per processor on the IBM SP4. For a typical simulation of 20 forcing cycles, using 11 million grid points, approximately 2,000 single-processor hours are needed on the IBM POWER 4 machine (MARCELLUS).

A key feature of all the studies reported here is that the inlet pipe is fully resolved so that the flow at the jet exit plane evolves naturally. This is in contrast to many past studies where the jet exit plane profile is typically specified as a boundary condition.

THE SQUARE SYNTHETIC AND FORCED JETS

The dimensions of the square jet computational domain are shown in Figure 1. The grid is stretched from the high resolution in the orifice region, but the stretching is maintained below 10 percent to ensure accuracy is not compromised. The inlet region is resolved using $170 \times 170 \times 52$, the nozzle is resolved using $66 \times 66 \times 7$, and the outflow region is resolved using $202 \times 202 \times 234$.

Figures 2a and 2b show, respectively, typical visualization of the vortex flow generated by the synthetic jet and forced square jet. The forcing frequency for both cases is the same, and the main difference between these two cases is that there is no mean flow in the synthetic jet case.

A key feature observed in both is the effect of vortex stretching and

breakdown as the flow expands downstream of the orifice. Dissipation is maximum in the high strain regions that typically reside in the braid regions and in the regions surrounding the vortices.

The square jet also shows an axis-switching behavior seen in the experiments as well. Axis switching is indicated by the crossover of the spreading rate of the jet in the two planes. In the near-field region of the jet exit, the vortex structures at the corners are formed farther downstream with respect to the sides. This triggers the axis switching since it results in the formation of nonplanar vortex structure. Comparisons with data from Feiz et al.¹ show reasonable agreement with past experiments.

SQUARE JET IN CROSS-FLOW

The computational domain for the test case shown in Figure 3 is resolved using $200 \times 150 \times 100$ for the cross-flow domain and $50 \times 50 \times 100$ for the jet section. The Reynolds number is 4700, based on the jet velocity and the nozzle width D , and the jet cross-flow velocity ratio is 0.5. The cross-flow velocity profile is initialized with a boundary layer thickness of $2D$. Figure 3 also shows a comparison of predicted mean velocity and total kinetic energy with data at a specified





Figure 2. Flow visualization of (a) forced square jet and (b) synthetic jet. The color iso-surfaces indicate values of constant vorticity. Green indicates azimuthal vorticity, and red/blue indicates streamwise vorticity of equal and opposite sign. Initially, azimuthally coherent vortices are shed from the orifice, but undergo vortex switching and stretching, eventually leading to breakdown in more randomly oriented streamwise vortices.

Figure 4. Flow visualization of the jet in cross flow. The formation of the hanging vortices and the formation of the counter-rotating pair in the downstream direction is clearly seen. Recirculation downstream of these structures also forms, as indicated by the streamlines.



location. Very good agreement is obtained here and also at other locations.¹

A jet in cross-flow generates a complex flow topology due to the highly 3D nature of this flow. Past studies have identified two structural features in this flow: a horseshoe (or kidney-shaped) structure and a Counter-Rotating Vortex Pair (CRVP) form in this flow.

The current LB-LES captures both these features and also explains the dynamics of the formation of these structures and their subsequent breakdown. Figure 4 shows these features quite clearly.

The horseshoe vortices are tubelike structures that form directly above the exit on the lateral edges of the jet and extend around the jet body and up along the lee side of the jet, approximately matching the path of the jet. These tubes coincide with the location where the jet shear layer folds and eventually contribute to the circulation of the CRVP.

Figure 4 also shows how the jet rolls up and creates the recirculation region: an important mechanism for the mixing of jet and the cross-flow. Finally, Figure 5 shows a time sequence of the formation of these flow features as the jet exits from the orifice and is turned downstream by the cross-flow.

In summary, a new LES implementation of the LBE method has been developed and used to simulate a 3D square jet and a 3D square jet in cross-flow. A localized dynamic subgrid closure is used to close the LES version of the LBE model. In these simulations the inflow is applied far upstream of the jet exit plane,

which allows the jet exit profile to evolve naturally.

Good agreement with established data was obtained in the present study. These results establish LBE-LES as an alternate method for simulating turbulent shear flows.

Future application of this LBE-LES approach will be in a hierarchical simulation approach whereby conventional finite-volume LES methods will be used to resolve the large-scale flow features in the combustor, while the LBE-LES approach will be used to resolve the finer scale features as in the embedded synthetic jet and/or the flow inside the fuel injector.




Figure 3. Schematic of the jet in cross-flow and comparison with experimental data.




Figure 5. Time sequence of the formation of the jet in cross-flow and the shedding of the hanging vortices as the flow propagates downstream.

References

1. Feiz, H., H. Soo, and S. Menon, "LES of Turbulent Jets Using the Lattice Boltzmann Approach," AIAA Paper No. 2003-0780, 41st AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, January 2003.

Contact Information

For further information contact Suresh Menon at suresh.menon@ae.gatech.edu

SUPERCOMPUTING 2003

PHOENIX | AZ | November 15-21

IGNITING INNOVATION

SHOWCASING:

- ★ Current DoD Research Projects
- ★ Visual Analysis Techniques

- ★ Implementations of leading-edge HPC Technologies including Hardware, Software, Storage, & Networking
- ★ Tool Development

HPCMP Booth Hosted by NAVO MSRC

HPCMP Booth contact: Pete Gruzinskas NAVO MSRC, Visualization Center, 228.688.4027, SC03@navo.hpc.mil, www.navo.hpc.mil/SC_2003.html

High Performance Computing and Simulation for Advanced Armament Propulsion

Michael J. Nusca, U.S. Army Research Laboratory (ARL), Aberdeen Proving Ground, MD

The Army is exploring a variety of armament propulsion options for indirect- and direct-fire weapons (guns) for the legacy force and Future Combat Systems (FCS).

As it transforms, the Army has identified requirements for hypervelocity projectile launch systems for strategic Army missions. Among these systems are those that use solid propellant—granular form loaded in modules (indirect-fire) or disk and strip form for High-Loading-Density (HDL) cartridges (direct-fire)—augmented by ElectroThermal-Chemical (ETC) technology. Two such armament propulsion systems are the Army's Modular Artillery Charge System (MACS) and HLD charges for the FCS.

The MACS is being developed for indirect fire cannon on current 155 millimeter (mm) systems (e.g., the M109A6 Paladin and M198 Towed Howitzer). The efficiency of the MACS charge is dependent on proper flamespreading through the propellant modules, a process that has been repeatedly demonstrated in gun firings, successfully photographed using the Army Research Laboratory (ARL) 155 mm ballistics simulator, and numerically modeled using the ARL Next Generation Three Dimensional (3D) interior ballistics code (NGEN3). The FCS requires weapons systems exhibiting increased range, accuracy, and highly repeatable projectile launch performance.

One of the technologies under investigation to achieve these goals is the ETC concept, in which electrically generated plasma is injected into the gun chamber in order to efficiently

and repeatedly ignite the high-energy and HLD solid propellant charge.

As modular and HLD propelling charges are being developed, optimized, and ultimately mated to systems such as indirect-fire cannon and the continually evolving FCS, there is a critical need to have a single, validated, maintainable computer code based on state-of-the-art Computational Fluid Dynamics (CFD) as an evaluation and performance analysis tool.

It has long been recognized that the availability of such a tool would provide the Army with the unique capability to simulate current and emerging gun propulsion systems using computer simulations. These simulations would serve to both streamline testing and aid in the optimization of weapon performance.

Indeed, such a tool would dovetail nicely with the Army's initiative in the creation of national High Performance

Computing (HPC) facilities. However, the gun propulsion-modeling environment has historically been one in which separate codes (some one-dimensional, some two-dimensional (2D)) are used, with no single multidimensional code able to address the truly 3D details of all these weapons systems.

This unfortunate situation renders comparison of ballistic performance cumbersome and inconclusive. In contrast, the multiphase continuum equations that represent the physics of gun propulsion comprise a set of general equations universally applicable to all solid propellant armament propulsion systems.

In direct response to this situation the ARL began a development program about eight years ago to revolutionize the Army's ability to use HPC to simulate propelling charges. The current author at ARL, with consultation from noted industry/academic experts, has worked on this project. The result is the Army's "next-generation," computer scaleable, 3D, multiphase, CFD code for armament propulsion modeling.

The ARL NGEN3 code represents the sole Department of Defense (DoD) computer tool that is able to simulate the highly complex physics associated with indirect- and direct-fire guns. NGEN3 code development and application to the FCS is a DoD HPC Challenge Project (FY01-03) and is being exercised regularly with priority access to the Cray SV1ex at the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC).




Figure 1. (a) Porosity Contours (red is dense material) at Initial Time, (b) Porosity Contours at 6 ms, and (c) Propellant Temperature Contours (red is fully ignited propellant at 440K) at 6 ms.

A NEW COMPUTING TOOL FOR THE ARMY

The Army's NGEN3 code is a multidimensional, multiphase CFD code that incorporates 3D continuum equations along with auxiliary relations into a modular code structure. Since armament propulsion involves flowfield components of both a continuous and a discrete nature, a coupled Eulerian Lagrangian approach is utilized. On a sufficiently small scale of resolution in both space and time, the components of the flow are represented by the balance equations for a multicomponent reacting mixture describing the conservation of mass, momentum, and energy. A macroscopic representation of the flow is adopted using these equations derived by a formal averaging technique applied to the microscopic flow. These equations require a number of constitutive laws for closure including state equations, intergranular stresses, and interphase transfer.

The numerical representation of these equations, as well as the numerical solution thereof, is based on a finite-volume discretization and high-order accurate, conservative numerical solution schemes. Further details are supplied elsewhere.¹

Multidimensional, multiphase flow modeling of a single armament launch scenario proceeds from propellant and projectile loading (initial conditions) to propellant consumption and projectile launch. Each detailed simulation necessarily requires large amounts of computer memory (10-50 gigabit) and time (10-130 Central Processing Unit hours) on the Cray SV1ex. As a result,




Figure 2. NGEN3 simulation of a notional high-loading density charge consisting of propellant disks (breach to 35 cm) and granular propellant (35 cm to projectile base) at 6 ms since igniter functioning. Nonsequential propellant ignition and the establishment of an adverse pressure gradient are demonstrated: (a) Porosity Contours (green is dense material), (b) Igniter Gas Mass Fraction Contours, (c) Propellant Temperature Contours (red is fully ignited propellant) and Selected Velocity Vectors, and (d) Gas Pressure Contours (red is high pressure).

the HPC resources at the NAVO MSRC that were made available through a DoD Challenge project are being utilized by the ARL when employing the NGEN3 code, as discussed below.

RESULTS FOR TYPICAL GUN CHARGE

Figure 1, as a whole, shows the computed porosity contours (blue to red: open space to nearly solid material) and propellant temperature contours (blue to red: ambient to 440K) for a 120 mm HLD charge consisting of separate regions of

disk and granular propellant. The propellant disks are stacked axially in the chamber; each disk has an inner radius that provides space for the igniter and the projectile afterbody and outer radius that is smaller than the radius of the chamber (shown from centerline to chamber wall).

Figure 1a shows the initial condition of the charge (i.e., before the igniter is activated). Figures 1b and 1c show the condition after 6 milliseconds (ms). (Note that the projectile has moved into the gun tube and out of view.) It can be noted that the granular propellant has been consumed and that the stack of disks has been pushed forward but is not fully ignited (See Figure 1c). This compressed region of disks (from 50-55 cm in Figure 1b) now lacks the interstitial gaps that have been closed, preventing convective heat transfer.

When this simulation is repeated using an ETC igniter, plasma convection is accomplished between all disks, before significant disk compression, and thus the entire stack of disk propellant is efficiently ignited. Thus, the low molecular weight plasma circumvents this problem by establishing a convectively driven flame that propagates faster than the material compression wave in the disk propellant, thereby permitting even ignition.

A basic design tenet for using HLD charges was established: namely, pressure waves in the chamber generated by the conventional igniter, when paired with the disk propellant charge, were avoided by using an ETC igniter. This conclusion is of critical importance to the design process for the FCS weapon.

Acknowledgements

The ARL mission program has supported development of the NGEN3 code. Support was also provided via contracts with PGA, Inc. Computer resources at the NAVO MSRC were provided via the DoD HPCMO under the HPC Challenge program. Application of the NGEN3 code to the 155 mm MACS propelling charge has been supported by PM-Crusader, Project Manager - Combat Ammunition Systems (PM-CAS), and the Army Armament, Research, Development, and Engineering Center (ARDEC).

References

1. Nusca, M., "High Performance Computing and Simulation for Advanced Armament Propulsion," Proceedings of the 23rd Army Science Conference, 2002.

Largest NAVO MSRC System Becomes Even Bigger and Better

David K. Magee, HPC Systems Engineer, NAVO MSRC

The NAVO MSRC is pleased to announce that the IBM POWER4 system (MARCELLUS) is being upgraded with seven additional compute frames as part of the Technology Insertion for Fiscal-Year 2003 (FY03) (TI-03) process within the Department of Defense (DoD) High Performance Computing (HPC) Modernization Program. This upgrade provides 224 additional processors and 224 gigabytes (GB) of memory for the POWER4 system.

With these upgrades, the POWER4 system features 1,408 1.3 gigahertz POWER4 processors, 1,632 GB of aggregate high-speed memory, and 14+ terabytes of resilient, high-performance disk arrays. All 1,408

processors within the POWER4 system reside within 44 compute frames, each containing 32 processors. Forty-two of these frames are each configured with 32 GB of memory and are logically partitioned into four eight-processor nodes with 8 GB of memory per node. The remaining two compute frames will be configured as larger 32-processor Symmetric Multiprocessing (SMP) nodes, one with 256 GB of shared memory and the other with 32 GB of shared memory. These two frames will be specially scheduled within Loadleveler to provide enhanced large memory and SMP processing support to users of the POWER4.

The system will continue to be configured with resilient interactive login and Global Parallel Filesystem (GPFS) capabilities, all interconnected with the current IBM Colony Switch fabric.

These upgrades will provide an approximate 20-percent increase in the peak computing capacity of the POWER4 system with improved processing capacity, resilience, and new capabilities for accommodation of very-large-memory, small-processor-count jobs. We believe the DoD HPC user community will be well-served by these enhancements.

Using the "smp" Queue on MARCELLUS

John Cazes, NAVO MSRC

Due to the recent addition of new frames to MARCELLUS as part of Technology Insertion 03 (TI03) (See article above), the NAVO MSRC has replaced the bigmem queue with a new queue, "smp", to offer shared memory users more performance.

The previous configuration of MARCELLUS was composed purely of 8-way shared memory nodes, which is optimal for Message Passing Interface (MPI) programs on MARCELLUS. Unfortunately, the 8-way nodes limited shared memory programs using OpenMP or pthreads to at most 8 processors. The MARCELLUS TI-03 expansion of 7 new IBM p690 frames with 32 processors each has given the NAVO MSRC the flexibility to offer shared

memory users more power. Two of the new frames will be configured as 32-way shared memory nodes, allowing shared memory applications to run on up to 32 processors. Also, one node will be configured with the maximum amount of memory, 256 gigabytes (GB) for shared memory or serial applications requiring large amounts of memory. Both nodes will be accessible via the "smp" queue.

Since these nodes are intended for use by shared memory jobs which may use fewer than 32 processors, the default setting for the LoadLeveler directive, "node_usage", will be "shared". This will allow jobs to be

added to the node as long as resources remain available. If the user wants to guarantee that there are no competing jobs on this node, possibly for benchmarking purposes, "node_usage" may be set to "not_shared". Of course, then the user's allocation will be charged for using 32 processors regardless of the actual number used.

Unlike an MPI job, for pure OpenMP and pthreads applications, the "tasks_per_node" or "total_tasks" LoadLeveler directive should be set to



1. Otherwise, if "total_tasks" is set to N, N identical shared memory jobs will run when started with poe. This would oversubscribe the node and overcharge the user's allocation. Also, users will be required to state what resources are needed, such as number of processors and amount of memory. This task may be done by using the "resources" directive as shown in the example below. The "resources" directive will set the "ConsumableCpus" and "ConsumableMemory" LoadLeveler variables. "ConsumableCpus" refers to the number of processors needed per task, and "ConsumableMemory" refers to the amount of memory needed per task. For a pure OpenMP or pthreads application running as one task, this would be the total amount of memory needed for the application.

LoadLeveler will automatically route the job to the appropriate node based on the stated memory and Central Processing Unit (CPU) requirements. One node will have slightly less than 32 GB available, some memory is set aside for the operating system, and

the other will have slightly less than 256 GB available. By default, "ConsumableMemory" is specified in megabytes, although LoadLeveler will also accept units of gigabytes (gb) or kilobytes (kb).

For example, the submit script in the box at right is for a 16-thread OpenMP application to be run through the "smp" queue. Since this is a pure OpenMP job, the total number of tasks is set to 1, and "ConsumableCpus" is set to "16". This job requires at most 512 MB per thread, so "ConsumableMemory" is defined as "8192" or "8 gb".

In summary, to use the "smp" queue, users must add the LoadLeveler "resources" directive to their scripts with the appropriate processor and memory requirements. Also, users should remember to set the number of tasks to 1 and the "node_usage" directive to "shared" to avoid being overcharged. As usual, all problems and questions may be directed to User Support via email, msrchelp@navo.hpc.mil, or via telephone at 1-800-993-7677.

```
#!/bin/csh
smp queue
@ output = my_OpenMP.out
@ error = my_OpenMP.err
@ account_no = Project_Name
@ wall_clock_limit = 6:00:00
@ class = smp
@ job_type = parallel
@ node_usage = shared
@ node = 1
@ tasks_per_node = 1
@ resources =
ConsumableCpus(16)
ConsumableMemory(8 gb)
@ queue

setenv OMP_NUM_THREADS
16
poe ./my_OpenMP

End of script
```

Environmental Quality Modeling Activities Under PET

Mary F. Wheeler and Clint Dawson, Center for Subsurface Modeling, Texas Institute for Computational and Applied Mathematics, The University of Texas at Austin

The Programming Environment and Training Program (PET) Environmental Quality Modeling (EQM) Computational Technology Area (CTA) team at the University of Texas at Austin (UT) includes Professor Mary F. Wheeler, the Functional Area Point of Contact (FAPOC) for EQM, and Professor Clint Dawson; Dr. Jeff Hensley, the EQM on-site representative at the United States Army Engineering Research and Development Center (ERDC) in Vicksburg, MS; Drs. Victor Parr and Malgorzata Peszynska, and several graduate students. Drs. Lea Jenkins and Beatrice Riviere were also involved in these efforts.

The UT PET EQM team activities include the development and implementation of accurate, locally mass-conservative numerical methods for flow; the assessment and comparison of multiphase flow; models for the subsurface; the implementation of stable and accurate solution methods for reactive transport models; the development of error estimation techniques for reliability and for driving adaptive numerical strategies; the development of robust

linear and nonlinear solvers; and the development of multialgorithmic and coupling strategies for multiphysics applications.

The migration of software to High Performance Computing (HPC) platforms through the development of parallel algorithms and domain decomposition methods has also been a major objective of this effort. Much of this work has been directed at enhancements to several Department of Defense (DoD) EQM codes, including the Adaptive Hydraulics (ADH) code, CE-QUAL-ICM, a water quality model, and the DoD Advanced Circulation Model (ADCIRC).

BACKGROUND

EQM encompasses the simulation of flow and reactive transport in environmental systems. These systems include the atmosphere and subsurface (below ground) and surface water environments. Specific EQM applications within the DoD include containment and mitigation of contamination for environmental cleanup, water quality, and noise and air pollution.

Flow models developed under EQM range from multiphase, multi-

component flow within the subsurface to hydrodynamics modeled by shallow water and Navier-Stokes equations. Reactive transport of chemical species is pervasive throughout EQM; thus the solution of systems of advection-diffusion-reaction equations is extremely important. DoD simulators for these systems include finite element and finite difference flow models for the subsurface and surface water as well as finite volume schemes for reactive transport.

ENHANCEMENTS TO ADH AND CE-QUAL-ICM

Two recent PET focused efforts have involved the study and implementation of improved simulation capabilities in ADH and CE-QUAL-ICM. Dr. Hensley's on-site activities at the ERDC MSRC PET facilities have included helping DoD users with issues of parallel code maintenance, testing and verification of codes, and providing users with the latest in parallel software enhancements. As part of this effort, Dr. Hensley worked with ADH code developers Dr. Stacy Howington and Dr. Charlie Berger at the ERDC to




Figure 1. Experimental Schematic of Vauclin Problem — Vauclin problem models a vadose zone water table recharge. The test domain is a sand-filled tank (600 cm x 180 cm x 10 cm) with trenches on either side. The tank is initially filled with water and allowed to drain to a water height of 60 cm. After the reservoir has stabilized, the tank is infiltrated with water through 100 cm along the top edge while a pump mechanism maintains a water height of 60 cm. After 8 hours of simulation, the water saturation profile resembles the figure on left.

consolidate recent additions and alterations to ADH and to provide technical assistance in verifying/ debugging the code. Dr. Hensley developed an OpenMP version of the CE-QUAL-ICM code that allows EQM users to use OpenMP to parallelize similar codes.

ADH ACTIVITIES

ADH is an adaptive, finite element flow simulator for subsurface and surface flows that models single-phase flow in saturated porous media and uses a simplified two-phase flow model (Richards Equation) for unsaturated porous media. The use of Richards Equation is restricted to an infinitely mobile air phase, which is not valid in some flow situations. A more accurate, but possibly more expensive, alternative would be to model the unsaturated zone as a compressible two-phase system. A standard test problem to compare the two formulations is the so-called Vauclin Problem, which models water flowing into a variably saturated porous medium.

In the Vauclin Problem, water enters through a narrow trench, as seen in Figure 1. As the water flows through

the air-water region, it displaces air and causes a rise in the water table. To simulate this phenomenon, the Integrated Parallel Accurate Reservoir Simulator (IPARS) was used. IPARS is a multiphase flow simulator developed at UT. Results of these simulations can be found in a recent technical report (L. Jenkins, The IPARSv2 air-water model, TICAM Report 02-27, Institute for Computational Engineering and Science, The University of Texas at Austin, 2002).

By using a full two-phase simulator, there were no difficulties in modeling this problem. The Richards Equation model, on the other hand, had severe numerical difficulties, especially in solving the nonlinear system of equations that arises at each time step. Further comparisons between the models on benchmark problems are the subject of current and future work.

CE-QUAL-ICM ACTIVITIES

CE-QUAL-ICM is a finite volume water quality model. Though the data structures in the code allow for very general types of quadrilateral elements, the underlying numerical methods in the code assumed a

structured, rectangular grid. However, current and future DoD applications require the use of unstructured, even nonconforming grids. At the request of DoD users, a discontinuous Galerkin Method was implemented into the code, with both slope and flux-limiting stabilization procedures, which allows for unstructured quadrilateral elements. To test this new formulation, it was applied it to a model of Chesapeake Bay. This model uses a structured grid so results could be compared to the finite volume discretization previously in the code. The two methods gave very similar results, which would be expected for this problem, as can be seen in the contour plots of temperature and salinity generated by the Discontinuous Galerkin (DG) and Finite Volume (FV) methods in Figure 2.

MULTIALGORITHMIC COUPLING

A PET-focused effort in 2001-2002 involved the investigation of multialgorithmic strategies for multiphysics couplings. One motivation for this research is in coupling flow and transport. For example, the transport model CE-QUAL-ICM requires a hydrodynamic flow field as input. In order for the




Figure 2. Temperature and salinity contours from the Chesapeake Bay simulation generated from CE-QUAL-ICM with a finite volume scheme (left) and CE-QUAL-ICM with DG enhancements (right).

finite volume discretization in the code to conserve mass and operate stably, this flow field must be locally conservative and flux continuous. That is, it must conserve mass element by element, and the normal flux across element interfaces must be continuous.

Many hydrodynamics codes do not satisfy this requirement, such as in the case of the ADCIRC. In fact, ADCIRC replaces the primitive continuity equation by a wave continuity equation. ADCIRC also has difficulties in modeling highly advective flow regimes. Replacement of the wave continuity formulation in ADCIRC with a discontinuous Galerkin Method applied to the primitive continuity equation was investigated with some success. This formulation provides

local mass conservation, flux continuity, and because it is suited for highly advective flows, improved capabilities in modeling steep water gradients.

ERROR ESTIMATORS/INDICATORS FOR EQM

A new effort for 2003 is in the area of a posteriori error estimation. The motivation for this continuing effort is to develop a means to be able to compute a reliably accurate solution with minimal degrees of freedom, e.g., to be able to use as coarse a mesh as possible and to know where to refine/de-refine the mesh. Work in this area has primarily focused on stationary models, and future efforts will concentrate on extending these ideas to time-dependent flow and transport systems.

Preliminary results for this project are shown in Figure 3, which shows the evolution of a pulse of concentration (e.g., contaminant). As the pulse is transported, the finite element mesh is adapted through the use of an error estimator. The result is a highly refined mesh in the neighborhood of the pulse, with much coarser mesh away from the solution. The degrees of freedom needed in this calculation are much less than if we used a refined mesh over the entire domain.




Figure 3. Evolution of a contaminant pulse using an adaptive discontinuous Galerkin Method. The mesh is adapted dynamically during the simulation based on an error estimator.

Clear Air and Optical Turbulence in an Airborne Laser Context

Alex Mahalov and Basil Nicolaenko,

Departments of Mathematics and Mechanical & Aerospace Engineering, Arizona State University



Michael Adams, NAVO MSRC Visualization Center

The Airborne Laser Challenge Project supports the Airborne Laser (ABL) and the associated technology program by performing computer simulations of pertinent optical, clear air turbulence phenomenology and forecasting. The problems addressed include (1) the development, through Direct Numerical Simulation (DNS) microscale turbulence modeling, of improved turbulence parameterization models for use in mesoscale atmospheric codes such as MM5 for optical turbulence forecasting; and (2) an investigation and evaluation of the efficacy of existing turbulence parameterizations and improved closure relations in mesoscale weather prediction models in support of the ABL Atmospheric Decision Aid (ADA).

The intent of the project is that this improved forecasting capability will eventually be incorporated into the atmospheric decision aid currently being developed for operational use in deploying ABL assets. The main focus of the present study is to characterize variability and statistics of tropopausal turbulence at vertical scales $O(10m)$ - $O(100m)$ in the ABL context. These scales are unresolved in mesoscale meteorological codes, and the turbulent dynamics on these scales is poorly understood.

The simulations discussed here are being conducted at the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) on the IBM POWER4, MARCELLUS.

Though poorly understood, vertical scales $O(10m)$ - $O(100m)$ are crucial for forecasting both Clear Air Turbulence (CAT and optical turbulence. The wind shear) associated with jet streams acts as a major source of clear air turbulence on these scales. Also, it is now well-known that the lower stratosphere, with its characteristic stably stratified conditions, can support persistent layering (sheets of strong optical turbulence) even up to vertical scales $O(1m - 10m)$ and relatively large horizontal extent,¹⁻³ the implication



Three-dimensional images of the absolute vorticity field.

Jet Stream in the

being that one should use very high vertical resolution in numerical studies. Vertical scales controlling the size of "sheets" in the atmospheric temperature field have been evidenced in the analysis of field measurements; the main dynamical properties of such layers, whether strongly mixed or not, can be characterized not only by Ri_{gr} , but also by various outer scales of turbulence as shown by Alisse and Sidi.⁴ In Joseph et al.⁵ the existence of multiple branches in the scaling of turbulence quantities across the jet stream is demonstrated.

The robustness of multiple branches, with respect to DNS at doubled resolution (1,024 vertical levels), is also demonstrated. These issues of multiple scalings can be expected to be crucial for motivating appropriate parameterizations of turbulence associated with tropopause jet stream for the ABL-ADA.

Atmospheric optical turbulence is defined as temporal and spatial fluctuations of the index of refraction. While it is most obviously manifested by the twinkling of stars, it also is a major source of performance degradation for optical system.⁶ Optical turbulence is not identical to CAT, but it is intimately related since temperature and velocity fluctuations are strongly coupled.⁷ The most important quantification of optical turbulence for optical propagation calculations is the refractive index structure constant C_n^2 , a crucial parameter in electromagnetic wave propagation studies. For the upper troposphere and above, C_n^2 is usually parameterized as:⁸⁻⁹

$$C_n^2 = a M^{-2} L_{out}^{4/3}, \quad (1.1)$$

where a is a constant (generally taken as 2.8), L_{out} is an outer length scale for turbulence (this article collectively refers to all length scales above the Kolmogorov viscous dissipation (inner) scale as outer scales), and M is the gradient of the

CONTINUED NEXT PAGE

generalized potential refractive index given by:

$$M^2 = \left(\frac{79 \times 10^{-6} P}{gT} \right)^2 N^4 \quad (1.2)$$

In Equation 1.2, T is the absolute temperature in $^{\circ}\text{K}$, P is the pressure in mb, g is the acceleration due to gravity, and N is the Brunt-Väisälä (buoyancy) frequency given by $N^2 = (g/\theta)(d\theta/dz)$, where the potential temperature θ is defined for the atmosphere as $\theta = T(1000/P)^{0.286}$. Radiosondes or mesoscale numerical models, such as MM5,¹⁰ can give P , T , and N^2 directly, but outer scale L_{out} , appearing in Equation 1.1, requires improved parameterization. One of the major challenges of ABL-ADA is to characterize the variability of L_{out} and its dependence on the background shear (jet stream) and stratification profiles using high resolution, three-dimensional (3D) DNS results.

The midlatitude tropopause is characterized by a thin stratified shear layer across which the buoyancy frequency (N) jumps approximately by a factor of two, even in a climatological sense, within a distance of about ± 2 km surrounding the jet wind maxima.¹¹ With the microscale computational box centered at the tropopause, the

tropopausal turbulence at the $O(10\text{m})$ - $O(100\text{m})$ vertical scales are studied, for which wind shear associated with the jet stream acts as the major source.¹²

Tropopause jet regions are also strongly influenced by stable background stratification, which allows gravity waves, patchy clear air, and optical turbulence.

The presence of (vertical) inhomogeneities in shear and

stratification causes the simulation of tropopause jet to have properties distinct from many previously considered DNS of uniform shear-stratified flows.

Tse et al.¹³ performed the first nonlinear, 3D DNS of such an inhomogeneous model tropopause jet using a spectral domain decomposition method in the vertical.

For this, the vertical domain is divided into subdomains. The spacing of subdomains is nonuniform so that a high resolution is maintained in regions with high shear. Any variable in each subdomain is interpolated using fourth-order Lagrange polynomials at Gauss-Lobatto-Legendre points. The vertical derivatives are evaluated by differentiating the interpolants. The time discretization follows the usual splitting procedure: with nonlinear and buoyancy terms advanced in the first substep, pressure-projection in the second substep, and dissipation term in the third substep.




Figure 1. Vertical cross-sections of instantaneous fields: (a) spanwise vorticity on central Y - Z plane, (b) spanwise vorticity on central X - Z plane, (c) total temperature in the central Y - Z plane, and (d) total temperature in the central X - Z plane.




Figure 2. A 3D image of the absolute vorticity field.

The code is parallelized using a Message Passing Interface (MPI) and runs on the NAVO MSRC IBM POWER4 (P4) MARCELLUS. Quasi-equilibrium solutions (i.e., when budget terms statistically showed an approximate balance and remain quasi-stationary in time) were obtained for a wide range of parameters, allowing mean jet stream and stratification profiles to evolve during the simulations. In order to demonstrate resolution independence of our results, 3D DNS was also performed with doubled resolution (1024 vertical levels).

SIMULATED FIELDS

Turbulent mean state is characterized by an equilibrium Brunt-Väisälä (buoyancy) frequency (N^2) profile and the absolute value of mean wind shear ($|S|$). A dramatic decrease of N^2 in the jet core, through turbulent dynamics and mixing, is clearly noticed. One may identify locally small buoyancy frequencies to be associated with relatively small values of $Ri_g = N^2/S^2 (< 0.25)$.

An N^2 profile of this kind, with a localized minimum in the mixing region, is sometimes referred to as possessing an “ N^2 notch” and is demonstrated to be a mean state configuration favoring emission of gravity waves from tropospheric jet streams. Nastrom and Eaton¹⁴ found a localized decrease of N^2 at tropopause levels in several winter seasonal profiles, which add credence to this project’s simulation results. The quasi-equilibrium mean velocity profile has inflection points on either side of the jet maximum, and an instability structure may be expected as two interacting trains of Kelvin-Helmholtz (KH) billows on either side of the jet maximum.

The time-varying mean part of the potential temperature tends to drive the total mean potential temperature profile toward a quasi-neutral state, characteristic of the formation of a mixing layer centered around the level of mean jet maximum. It has been verified in field measurements of Luce et al.³ that such quasi-neutral layers can result from strong turbulent mixing generated by shear instabilities near the tropopause.

In order to appreciate the nature of simulated fields, a few cross-sections of the spanwise vorticity and total temperature fields are presented in Figure 1. It is clear from spanwise vorticity plots (on the central Y - Z and X - Z planes), seen in Figures 1a and b, that the focus of this study is a turbulent field, after the occurrence of several instabilities, with a fully 3D structure.

From Figure 1b, it is possible to identify remnants of KH billow structures generated through instabilities.

The generation of an intense turbulent mixing region, with characteristic localized over-turning regions in the high-shear jet-core region, is better evident in the plots of (total) temperature (See Figures 1c and d). This leads to rapid homogenization of temperature and any impurities released into this region and provides a mechanism for




Figure 3. Gravity waves. Hovmöller plots of horizontal divergence along selected lines: (a) X - t cross-section, (b) Y - t cross-section, (c) Z - t cross-section.




Figure 4. Scalings of w'^2 (circle), $\langle \theta'^2 \rangle$ (diamond), $\langle w'w' \rangle$ (triangle), and $\langle \theta'w' \rangle$ (plus) with Ri_g : (a) for $Z > 0$ (b) for $Z < 0$, (c) for $N^2 > N_{cr}^2$, and (d) for $N^2 < N_{cr}^2$.

MULTIPLE SCALINGS

This section includes more on the nature of multiple branches in scalings, omnipresent in the Airborne Laser Challenge Project results. For convenience, only scalings $\langle w'^2 \rangle$, $\langle \theta'^2 \rangle$, $\langle w'w' \rangle$ and $\langle \theta'w' \rangle$ (primes denoting deviations from the mean) are considered here. Since the asymmetry in the mean state stratification, about the level of jet maxima, is a plausible suspect, scalings of above quantities, with Ri_g , are separately considered for the region above ($Z > 0$) (See Figure 4a) and below ($Z < 0$) (See Figure 4b) the jet.

Nonuniformity in stratification (doubling of N across the atmospheric tropopause) plays an instrumental role in generating multiple scaling curves. However, as evident from Figures 4a and b, there are still two distinct branches on each of the delineated scaling curves. Therefore, some parameter other than Ri_g may be playing a critical role.

In Figures 4c and d, the data are classified into one group for which $N^2 > N_{cr}^2$, and another group with $N^2 < N_{cr}^2$. Now a clearer division between the two different branches on each of the scaling curves can be identified. The scaling branches of $\langle w'^2 \rangle$, $\langle \theta'^2 \rangle$, and $\langle w'w' \rangle$, which remain stationary with Ri_g , correspond to data points with $N^2 > N_{cr}^2$ (i.e., those within the turbulent jet core). It is also of interest to note that scaling of these quantities with Ri_g , under stronger stratification (away from jet), are well-defined (See Figure 4c). These results imply a cautionary note to experimenters and modelers who may be inclined to fit a unique curve passing through their data points, which in fact may lie on several distinct branches (See Figure 4).

mixing of stratospheric and tropospheric air in the jet core region. The existence of a well-mixed quasi-neutral layer (See Figures 1b, c, and d), with temperature variance peaking along its edges (See Figure 1c), is in qualitative agreement with observations by Luce et al.³ in shear layers $O(10 \text{ m} - 100 \text{ m})$ above the jet maximum. The quantitative difference being that this project models a stronger turbulent event in which the mixing layer spans the whole jet core (accomplished by transport from shear production regions). A 3D image of the turbulent absolute vorticity field is shown in Figure 2.

Figures 3a, b, and c present $X - t$, $Y - t$, and $Z - t$ Hovmöller plots of horizontal divergence, a variable directly related to gravity wave dynamics in the Craya decomposition of stratified flows.^{15, 16} There is some evidence of downstream propagation of features in the streamwise x -direction, but not in the spanwise y -direction (See Figures 3a and b). The origin of wave propagation in the

x -direction could be related to the mechanism whereby a turbulent mixed region with an N^2 notch stratification in a jetlike shear flow acts as a source for gravity wave generation. Thus, propagating gravity wave modes may indeed play an important role in the late development of shear-stratified mixing layer associated with a tropopause jet. Such occurrences of gravity wave emission from tropopause jet streams have also been demonstrated through observations in Bedard et al.¹² The $Z - t$ plot of horizontal divergence (Figure 3c) shows that there is not any significant vertical propagation of features away from the core region of the jet. This is plausibly due to the mean jet profile that vanishes at $|Z| \approx 4$, which will cause a wide spectrum of propagating waves to find their critical level (where mean velocity equals the wave phase speed), before they could propagate any further than $|Z| \approx 4$, and consequently results in nonlinear breakdown.

Acknowledgements

This work is supported by the Air Force Office of Scientific Research (AFOSR) Contract FG9620-02-1-0026 and the Department of Defense (DoD) High Performance Computing (HPC) Challenge Program. The computations were performed at the NAVO MSRC on the IBM P4, MARCELLUS.

References

1. Dalaudier, F., C. Sidi, M. Crochet, and J. Vernin, "Direct Evidence of "Sheets" in the Atmospheric Temperature Field," *Journal of the Atmospheric Sciences*, 51, 237-248, 1994.
2. Nastrom, G.D., and F. D. Eaton, "Persistent Layers of Enhanced C₂n in the Lower Stratosphere from VHF Radar Observations," *Radio Science*, 36, 137-149, 2001.
3. Luce, H., S. Fukao, F. Dalaudier, and M. Crochet, "Strong Mixing Events Observed Near the Tropopause with the MU Radar and High-Resolution Balloon Techniques," *Journal of the Atmospheric Sciences*, 59, 2885-2896, 2002.
4. Alisse, J. R., and C. Sidi, "Experimental Probability Density Functions of Small-Scale Fluctuations in the Stably Stratified Atmosphere," *J. Fluid Mech.*, 402, 137-162, 2000.
5. Joseph, B., A. Mahalov, B. Nicolaenko, and K.L. Tse, "Variability of Turbulence and Its Outer Scales in a Nonuniformly Stratified Tropopause Jet," *Journal of the Atmospheric Sciences*, submitted for publication, 2002.
6. Jumper, G.Y., and R.R. Beland, "Progress in the Understanding and Modeling Atmospheric Optical Turbulence," 31st American Institute of Aeronautics and Astronautics (AIAA) Plasmadynamics and Lasers Conference AIAA 2000-2355, 1-9, 2000.
7. Cote, O. R., J. R. Roadcap, D. E. Wroblewski, R. J. Dobosy, and T. L. Crawford, "2002: Aircraft Measurements of Turbulence: Length Scales, Spectra, Budgets, and the Prediction Problem." Extended abstract in 12th Symposium Meteor. Observations and Instrumentation, 83rd American Meteorological Society Meeting, Feb. 2003, Long Beach, CA, 2002.
8. Beland, R. R., "Propagation Through Atmospheric Optical Turbulence," *Atmospheric Propagation of Radiation*, F. G. Smith, Ed., International Society for Optical Engineering (SPIE) Optical Engineering Press, Washington, DC, 157-232, 1993.
9. Eaton, F.D., and G. D. Nastrom, "Preliminary Estimates of the Vertical Profiles of Inner and Outer Scales from White Sands Missile Range, New Mexico, VHF Radar Observations," *Radio Science*, 33, 895-903, 1998.
10. Dudhia, J., "A Nonhydrostatic Version of the Penn State/NCAR Mesoscale Model: Validation Tests and Simulation of an Atlantic Cyclone and Cold Front," *The Monthly Weather Review*, 121, 1493-1513, 1993.
11. Birner, T., A. Dombrack, and U. Schumann, "How Sharp is the Tropopause at Midlatitudes?" *Geophysical Research Letters*, 14, (45-1)-(45-4), 2002.
12. Bedard Jr., A.J., F. Canavero, and F. Einaudi, "Atmospheric Gravity Waves and Aircraft Turbulence Encounters," *Journal of the Atmospheric Sciences*, 43, 2838-2844, 1986.
13. Tse, K.L., A. Mahalov, B. Nicolaenko, and H.J.S. Fernando, "A Spectral Domain Turbulence," *Lecture Notes in Physics*, 566, Springer-Verlag, Berlin, 353-378, 2001.
14. Nastrom, G.D., and F. D. Eaton, "Turbulence Eddy Dissipation Rates from Radar Observations at 5-20 km at White Sands Missile Range, New Mexico," *Journal of Geophysical Research*, 102(D16), 19495-19505, 1997.
15. Mahalov, A., B. Nicolaenko, and Y. Zhou, "Energy Spectra of Strongly Stratified and Rotating Turbulence," *Physical Review*, E 57, 6187-6190, 1998.
16. Cullen, M. J. P., "New Mathematical Developments in Atmosphere and Ocean Dynamics and Their Application to Computer Simulation," *Large-Scale Atmosphere-Ocean Dynamics*, Vol. 1, J. Norbury and I. Roulstone, Eds., Cambridge University Press, Cambridge, 202-287, 2002.

Data Visualization with Alias|Wavefront™ Maya® 4.5

Michael Adams, NAVO MSRC Visualization Center

Over the last decade, software packages used to make visual effects in blockbuster movies have become more and more sophisticated as those effects have become more and more complex. Today, a standard PC running Alias|Wavefront™ Maya® software can handle amounts of data that required very expensive workstations just a few years ago.

Scientific data from computational models run at the NAVO MSRC are not that different from the underlying data structures Maya creates during dynamic simulations of its own. The images that accompany the article "Clear Air and Optical Turbulence in a Jet Stream in the Airborne Laser Context" (pp. 14-19) and this article were created using Maya 4.5 on a Windows 2000 PC using two 1-Gigahertz (GHz) processors and 1 Gigabyte (GB) of RAM. After setup, the files were transferred to a Silicon Graphics Onyx2 system with eight processors for final computation.

The data for the figures that accompany this article were created on MARCELLUS (IBM SV1). The first step was to process the data into a format that could be used in Maya. A combination of both Fortran and C was used to translate the data into little-endian 8-byte floats, known as doubles. After this was completed, the data were ready to import into Maya.

Using dynamics within Maya, a particle object was created with a per-particle user scalar attribute. The data were read according to a regular grid, which was copied into each particle's XYZ position information. Then the temperature or absolute vorticity data were read into each particle's user scalar attribute. Color gradations were attached to the data attribute using array mappers and particle sampler nodes.

To ease control of the final "look" of the image, the data were split into several different particle objects, based on the range of data within the variable being visualized. After normalization, the temperature data, for example, were split into three different sections.

Data records ranging from 0 to 0.6 were not used. Values between 0.6 and 0.875 were rendered as clouds, with transparency also mapped to the data value. The highest data values, from 0.875 to

1.0, were read into a particle object using a "blobby surface" rendering style that volumetrically links elements together to create a surface from orange to red.







Figure 1. (Top) Maya 4.5 showing particle object shaded according to temperature data.

Figure 2. (Middle) Maya 4.5 showing particle object geometry.

Figure 3. (Bottom) Maya 4.5 Hypershade showing shading network used to color individual particles according to imported data.

Figure 4. (Left) Completed Maya 4.5 image.





Above: Dave Cole, NAVO MSRC; Sharon Samples, Chief Information Officer, NASA Office of Earth Science (ESA); Dr. Nathan Sovik, Stennis Space Center (SSC) ESA.



Above: RADM Evelyn J. Fields, Director, National Oceanic and Atmospheric Administration (NOAA) Marine and Aviation Operations and NOAA Commissioned Officer Corps; Steve Adamec, NAVO MSRC.

Above: CDRE Nick Harris, Royal Navy (RN) Attache, British Embassy to the United States; Dave Magee, NAVO MSRC; and LCDR Don Ventura, Royal Navy, Personal Exchange Program (PEP) Officer.

Right: CDR Jamie McMichael-Phillips (RN), Directorate of Naval Surveying, Oceanography and Meteorology (DNSOM); CDR Bob Stewart, DNSOM; Gerald M. Leone, NAVOCEANO.



Above: LCDR David M. Furr, NAVOCEANO Supply Officer; Dave Cole, NAVO MSRC.



Above: Representatives of the Tupelo, MS, Mermaids Outreach Program.

Above: Attendees of the Oceans 2002 Marine Frontiers Conference in Biloxi, MS.

NAVO MSRC PET Update

Eleanor Schroeder, NAVO MSRC Programming Environment and Training Program (PET) Government Lead

Programming Environment and Training (PET) will be entering its third year of the new contract in June 2003. In February 2003, PET underwent its first technical review. Each of the 15 Functional Area Points of Contact (FAPOCs) provided updates on accomplishments made to date, on work being done, and on plans for the near future. The presentations can be viewed through the PET Online Knowledge Center (OKC) hosted by the U.S. Army Engineer Research and Development Center (ERDC) Major Shared Resource Center (MSRC).

This particular issue of the Navigator will include reports by the three FAPOCs for Component 1 of PET, namely for Climate/Weather/Ocean (CWO), Environmental Quality Management (EQM), and Computational Environment (CE). If you have any further questions or comments about the work being performed, please contact the pertinent FAPOC, me, or the contractor lead, Dr. George Heburn. As



always, PET is here to serve you, the user, in the best way that we can.

Also, if you haven't registered to be a user on the PET OKC, I highly recommend that you do that soon. This is your "one-stop shop" for information about what is going on in the PET programs. You'll find presentations that have been made by our PET partners, information about currently funded projects, the final reports from FY01 projects, as well as training information. It doesn't take long to register, and it is free!

PET Distance Learning: Ready to Serve 24/7

Andrew Schatzle, PET Component I

The PET Component I Distance Learning Program can trace its roots back to the pioneering days of streaming video over the Internet. When the PET program was initiated in 1996, one of the tasks assigned was to provide training to Department of Defense (DoD) scientists on how to best exploit the resources provided by the High Performance Computing Modernization Program (HPCMP) Major Shared Resource Center (MSRC) for their research.

With users spread around the globe, PET turned to video streaming to meet this task. The PET distance-learning program was designed to take classroom-training courses taught by leaders in their fields to the users' desktops. Training would be made available 24 hours a day, 7 days a week, and could be taken at the users' convenience. Starting with a dual 200-Megahertz (MHz) Personal Computer (PC), a \$200 video capture card, and a promising new software from a company called RealNetworks, Inc., PET initially produced full-length, 7 frames per second (fps), HPC video courses with synchronized instructors' slides via the Internet. Today the video streams are encoded in the NAVO MSRC Video Production Studio and now run at full-motion 30 fps.

Taking advantage of this increased capacity, online training courses continue to be available around the clock, and instructors include such recognized leaders as Paul

Messina, Director, Center for Advanced Computing Research, Assistant Vice President for Scientific Computing, and Faculty Associate in Scientific Computing; Walt Brainerd of Fortran, Inc.; The Cornell Theory Center; and partners from San Diego Supercomputer Center (SDSC).

Since its inception, the distance-learning program has covered a wide range of HPC topics including Fortran 90, Message Passing Interfaces (MPI), OpenMP, visualization techniques, and a variety of others. Another major training initiative was the development of the "Path to Becoming a Parallel Programmer" series of courses. This series was designed for DoD MSRC users who would like to either make the transition from single Central Processing Unit (CPU) (serial) processing to multiprocessor (parallel) processing or optimize preexisting parallel code. This distance-learning series, which is still available online, covers the multiprocessor programming styles associated with two current programming paradigms, MPI and OpenMP. In addition to the training videos, the PET video library includes the last three years of the annual Defense Research and Engineering Network (DREN) Conference sessions. These sessions have all been produced and hosted within the PET distance-learning program.

Currently, the NAVO PET distance-learning program includes the following courses:

- ↳ Collaborative and Distance Learning Technologies
- ↳ An Introduction to Perl Programming
- ↳ The Legion Metacomputing System: A User's Perspective
- ↳ Introduction to the IBM SP
- ↳ Porting Codes from the C90 to the Origin 2000
- ↳ Path to Becoming a Parallel Programmer Series
 - o An Overview of Parallel Computing Hardware
 - o An Overview of Parallel Computing Software
 - o Fortran 90
 - o Introduction to MPI for Finite Difference Models
 - o Introduction to MPI: The Complete Standard
 - o Introduction to Open MP for Finite Difference Models

These courses can be accessed, via a desktop computer that has a Web browser and an installed copy of the "Real Player" streaming media viewer by going to the NAVOCEANO PET home page at <http://www.navo.hpc.mil/pet/>. A free copy of Real Player can be downloaded via a link to the RealNetworks Web site on the NAVOCEANO PET video library Web page at <http://www.navo.hpc.mil/pet/Video/>. The lecture notes

for the courses, in PDF format, can also be downloaded from that page.

The reorganization of the PET Program last year gave the Collaborative Distance Learning Technologies (CDLT) directive to the PET component located at the Army Research Laboratory (ARL). The PET Component I group at the NAVO MSRC is working very closely with the CDLT group to ensure that a successful HPC distance-learning program continues to grow within the HPCMP. In the past year, two courses have been added to the program: "An Introduction to Perl Programming," taught by Dr. David Ennis from the Ohio Supercomputer Center and "Collaborative and Distance Learning Technologies (CDLT)," taught by the East Carolina University.

PET Component I is very interested in satisfying your training needs. If you desire training on any HPC subject, either online or live, please contact Brian Tabor at taborb@navo.hpc.mil. We would appreciate receiving any other feedback you might offer.

A Consistent, Well-Documented Computational Environment for the DoD HPC Centers

Shirley Moore, Innovative Computing Laboratory/University of Tennessee-Knoxville, PET Computational Environments FAPOC

In order to be maximally productive, Department of Defense (DoD) users need a consistent well-documented computational environment across the High Performance Computing (HPC) Centers. Although individual centers may have special requirements due to specific architectures or application areas, the general computational environment should be as consistent as possible across the centers.

The computational environment includes compilers, message passing libraries, numerical libraries, and debugging and performance analysis tools, as well as data management and visualization tools. All components of this environment need to be working properly in the batch queuing environments and be adequately supported and documented. Cross-platform tools should be made available wherever possible so that users do not need to learn a different tool interface for each platform. The environment should be updated on a regular basis to keep pace with changes in architectures and operating systems, and new and emerging computational technologies should be evaluated for possible adoption. Comprehensive information about architectures, systems software, and available libraries and tools should be available to users via the PET Online Knowledge Center (OKC). (For more detailed information about PET OKC offerings, see the preceding article.)

To achieve these goals, PET is working with the HPC Center systems and user support staff to implement and support such a consistent environment. For each major architecture, a concise set of documentation is being developed that describes the architectural features, operating system, compilers, parallel programming models, and available tools for that platform. A set of cross-platform performance analysis tools has been evaluated and selected for deployment at the HPC Centers. All deployed software and tools are being documented in a repository that shows what software is installed on what machines and provides quick-start guides.

PLATFORM DOCUMENTATION

The documentation being developed for each platform is intended to convey in a concise manner what most users need to know to successfully develop and run applications on that platform. Pointers to more detailed information are also provided. Several users have requested a concise summary of the most useful compiler options for each platform. The information provided in the documentation describes the various programming language compilers available and explains options for debugging, for checking conformance with standards, and for improving performance. In addition, the documentation describes how to develop and execute Message Passing Interface

(MPI), OpenMP, and mixed-mode parallel programs on that platform, while emphasizing how to ensure portability to other platforms.

To supplement the written documentation, training classes are offered by the OKC on HPC architectures, programming languages, basic and advanced MPI, OpenMP, mixed-mode programming, and performance optimization. Materials from previous versions of these classes are also available on the OKC.

MATHEMATICAL SUBROUTINE LIBRARIES

A number of portable subroutine libraries are available that provide efficient implementations of key numerical algorithms for scientific applications. In some cases, these libraries can provide near-peak performance for core linear algebra operations. The Basic Linear Algebra Subroutines (BLAS) are high-quality "building block" routines for performing basic vector and matrix operations. The Linear Algebra Package (LAPACK) provides routines for dense linear algebra, including solving systems of simultaneous linear equations, least-squares solutions of linear systems of equations, eigenvalue problems, and singular value problems. The BLAS and LAPACK have become industry standards and are implemented in most vendor-provided math libraries, including Engineering and Scientific Subroutine Library (ESSL) for IBM AIX, Commerce Extensible Markup Language (CXML) for Compaq Alpha, and Scientific Computing Software Library (SCSL) for SGI IRIX. In addition, the following portable math libraries are being installed on HPC Center platforms: (1) SuperLU, for direct solution of large sparse linear systems; (2) Portable, Extensible Toolkit for Scientific Computation (PETSc), for solution of large-scale problems modeled by partial differential equations including iterative methods for solution of resulting large sparse systems; (3) and the Arnoldi Package (ARPACK), for solution of large-scale eigenvalue problems. Shared and distributed memory parallel versions of most of these libraries are also available. As with the other tools discussed in this article, the PET OKC offers a training class on math libraries, and the materials from a previously taught version are also available.

DEBUGGING AND PERFORMANCE ANALYSIS TOOLS

TotalView is a commercial cross-platform source level multithread/multiprocess parallel debugger that has both command-line and graphical user interfaces and is available at some HPC Centers. TotalView supports debugging OpenMP, Pthreads, MPI, and mixed-mode parallel programs written in Fortran 77/90/95, C, or C⁺⁺. The Performance Application Programming Interface (PAPI) library, developed with PET support, is a cross-platform interface to the hardware performance counters

available on most modern microprocessors. The counters exist as a small set of registers that count events that are occurrences of specific signals and states related to the processor's function. Monitoring these events has a number of uses in application benchmarking, performance analysis, and optimization. In addition to routines for accessing the counters, PAPI specifies a standard set of performance metrics considered most relevant to application performance tuning (e.g., cache and memory hierarchy accesses, cycle and instruction counts, and pipeline and functional unit status), and the PAPI reference implementation attempts to map as many of these standard events as possible to native events on the underlying platform. PAPI has been implemented for and is being deployed on all HPC Center platforms. PAPI may be used either directly or through a tool interface.

Tuning and Analysis Utilities (TAU) is a freely available profiling and tracing toolkit for parallel OpenMP, Pthreads, MPI, and mixed-mode parallel programs written in Fortran 77/90/95, C, or C⁺⁺. TAU provides automatic instrumentation for all programming languages and parallel programming models and uses PAPI to access hardware counter data. TAU provides both command-line and graphical user interfaces for viewing and analysis of profile data.

In addition, TAU can generate trace files of parallel executions that can be viewed by the commercial Vampir MPI performance analysis tool. TAU is being deployed on all HPC Center platforms, and Vampir is available at several sites. The Vampirtrace library can also generate Vampir trace files. Together PAPI, TAU, and Vampir provide a comprehensive suite of cross-platform performance analysis tools. The OKC offers a training course on this tool suite, and training materials from previously taught versions of the course are also available.

FOR MORE INFORMATION

For more information, visit the Computational Environments section of the OKC at <https://okc.ercd.hpc.mil/index.jsp>, or contact pet-ce@cs.utk.edu.

PET Climate/Weather/Ocean (CWO) Modeling and Simulation—A Brief Review

Dr. Jay Boisseau, PET CWO Team Lead, University of Texas at Austin

The PET CWO on-site support staff serve an important role in the PET program. They provide the most direct interactions—support and collaboration—for Department of Defense (DoD) users at that location (and also other potential sites) to enhance their research. The PET CWO on-site staff has collaborated with the DoD CWO users to improve the capabilities of several important CWO codes.

CWO ACCOMPLISHMENTS

Some of the major accomplishments from the past year by the CWO on-site staff include:

- ↳ Development of an OpenMP implementation of the Dartmouth College circulation model, QUODDY, three-dimensional (3D) baroclinic finite-element circulation model with excellent scalability, enabling much larger simulations for use by Naval Research Laboratory - Stennis Space Center (NRL-SSC) researchers Cheryl Ann Blain and Catherine Edwards.
- ↳ Development of an OpenMP implementation of the Simulating Waves Nearshore (SWAN) third-generation wave model for use by Rick Allard, Erick Rogers, and Larry Hsu of NRL-SSC. This has significantly reduced turnaround time for stationary and non-stationary cases, and the parallel code is now transitioning into operational use at the NAVO MSRC.
- ↳ Merging and updating the two-dimensional (2D) and 3D versions of the Advanced Circulation Model For Oceanic, Coastal, and Estuarine Waters (ADCIRC) Finite Element Model (FEM) circulation code successfully into a single, Fortran 90/Message Passing Interface (MPI) parallel version, now making it possible to consider dynamic problems that were previously not tractable. The ADCIRC is now well positioned to accept the future advances that await development.
- ↳ Development of an improved finite-volume, locally conservative, monotonic, interpolation scheme for remapping between vertical grids in ocean models, which is now implemented at NRL for nesting a sigma-Z coastal ocean model inside a basin-scale Hybrid Coordinate Ocean Model (HYCOM).
- ↳ Enhancing and optimizing the Wave Information Study (WISWAVE) code used by the Coastal and Hydraulics Laboratory (CHL) at the U.S. Army Engineer Research and Development Center (ERDC) for wave hindcasts. This activity included completion of a unified version of the WISWAVE model,

optimizing it using OpenMP, and validating it for the Northern Gulf of Mexico and Gulf of Mexico one-month simulation and for the Atlantic, South Atlantic, and North Atlantic one-year simulation.

- ↳ Adding variable grid resolution and local time-step capabilities to the hydrodynamic, salinity and temperature model CH3D-Z.
- ↳ Integration of the Coupled Ocean/Atmosphere Mesoscale Prediction System (COAMPS) dynamics engine into the Weather Research Framework (WRF) to take advantage of the modularization and other modern programming techniques implemented in WRF. The second is working with Dr. Wieslaw Maslowski of the Naval Postgraduate School to update the customized, coupled versions of the Parallel Ocean Program (POP) and Los Alamos sea ice model (CICE) that Dr. Maslowski uses for arctic region ocean and ice modeling.

PET FUNDED PROJECTS

The PET program also funds focused projects. In fiscal year 2002, the new PET program sponsored two CWO related projects:

CWO-007: Coupling of COAMPS and Wavewatch (WW) With Improved Wave Physics

Principal Investigator (PI): Pat Fitzpatrick; Team: Drs. Matt Bettencourt, Shahrdad Sajjadi, and Gueorgui Mostovoi, University of Southern Mississippi (USM)

Goal: Develop a coupled COAMPS-WW system using Model Coupling Environmental Library (MCEL) Training, and incorporate current parameterization methodologies used by NRL as well as the USM parameterization schemes.

Accomplishments: MCEL was successfully used to couple COAMPS and WW. It was successfully applied to model Hurricane Gordon, showing that two-way coupling results affect boundary layer physics differently than one-way coupling. The team also derived a new analytical expression for the wind/wave growth factor based on normal modes of analysis and a rapid distortion theory.

CWO-011 Functional Extension and Application of a Model Coupling Executable Library

PI: Keith Bedford, Ohio State University (OSU); Team: David Welsh, OSU; Matthew Bettencourt, USM

Goal: Extend MCEL validation and functionality. MCEL

reliability and ease-of-use will be validated by reformulating the coupling operations in COAMPS. Also, construct a generalized wave-current-sediment marine bottom boundary layer coupling filter for MCEL.

Accomplishments: A beta-test version of MCEL was deployed in COAMPS. This effort permitted a comparison of MPI and MCEL coupling and helped guide the development of MCEL 1.0. MCEL 1.0 was deployed in the SWAN model, in support of model coupling activities in the Common High Performance Computing Software Support Initiative (CHSSI) High Fidelity Simulation of Littoral High Fidelity Simulation of Littoral Environments (HFSole) project portfolio.

In this, the second year, two additional CWO projects have been initiated:

CWO-002: Infrastructure Development for Regional Coupled Modeling Environments

PI: John Michalakes, National Center for Atmospheric Research (NCAR); Team: Matthew Bettencourt, USM; Robert Jacob, Argonne National Laboratory (ANL); Jerry Wegiel, Air Force Weather Agency (AFWA); Joe Klemp, NCAR

Goals: Develop a software infrastructure for coupled modeling systems that abstracts details and mechanics of High Performance Computing (HPC) at two levels: (1) managing shared and distributed-memory parallelism within individual component models (L1), and (2) efficiently translating and transferring forcing data between coupled component grids each running in parallel (L2).

CWO-008: Enhancing the Capabilities of a 3D Nearshore Ocean Circulation Model System (SHORECIRC)

PI: Chandrasekher Narayanan, USM; Team: Matthew Bettencourt, USM; Tim Campbell, Mississippi State University (MSU); John Cazes, NAVO MSRC; Moinuddin Shalam, USM

Goal: Develop operational capabilities of SHORECIRC so that all users, including naval operations planners for defense-related activities, can use it quickly and efficiently. Development will include parallelization, coupling with wave models, and implementation of more realistic boundary conditions.

CWO TRAINING

Additionally, two computing classes have been taught specifically to address stated CWO user needs:

- ↳ "Portable Performance Programming," taught at the NAVO MSRC and ERDC by Dr. Kent Milfeld and Dr. Chona Guiang, Texas Advanced Computing Center, University of Texas (TACC/UT) and at the Fleet Numerical Meteorology and Oceanography Center

(FNMOC), Monterey, CA, by Dr. Avijit Purkayastha, TACC/UT, and John Michalakes, NCAR.

- ↳ "Comprehensive Overview of Scientific Visualization," taught at the NAVO MSRC, ERDC, and FNMOC by Dr. Kelly Gaither, Greg Johnson, and Reuben Reyes, TACC/UT.

In addition, the CWO project work on MCEL produced a training class to instruct CWO researchers in how to use it: "Model Coupling Environmental Library Training," taught at the NAVO MSRC by Dr. Matthew Bettencourt of USM. Training requests for this and other topics are taken throughout the year and classes are formulated as necessary by PET team members or through contracts to external parties to ensure that CWO user needs are met.

The PET program, which is in its second year under the new organization, has already transitioned from its previous incarnation and is now producing high-impact results in all Computational Technology Areas (CTAs) and in the development of useful crosscutting technologies. Funding is limited, but feedback from DoD users has been instrumental in developing an effective PET CWO support strategy and its implementation through on-site support staff activities, projects, and training. The PET CWO team is pleased with past results and current activities and looks forward to increased interaction and impact over the years. Please help the PET CWO team help you by providing feedback and requirements on your HPC-related CWO needs to the PET CWO team lead, Dr. Jay Boisseau (boisseau@tacc.utexas.edu) or any member of the PET CWO team.

CWO ON-SITE PERSONNEL

The CWO on-site personnel are:

- ↳ NAVO MSRC: Dr. Tim Campbell
tjcamp@navo.hpc.mil
- ↳ ERDC: Dr. Phu Luong
Phu.V.Luong@erdc.usace.army.mil
- ↳ FNMOC: Dr. John Romo
jromo@tacc.utexas.edu

Programming TotalView and Vampir

Sheila Carbonette, NAVO MSRC User Support

The following programming tools are now available on the NAVO MSRC IBM systems HABU and MARCELLUS. TotalView, a full featured, source-level graphical debugger is available on both systems. Vampir, a graphical tool for viewing trace files generated by Vampirtrace, is available on MARCELLUS only.

TOTALVIEW

The TotalView debugger from Etnus, Inc., can be used to debug applications written in Fortran, C, C++, and assembler. It is a multiprocess, multithread debugger that can be used on MPI, PVM, or OpenMP parallel codes. Before the debugger can be invoked, several environment variables must be set. These include PATH, MANPATH, and LM_LICENSE_FILE. CSH/TCOSH users should run the following command to set up the environment:

```
% source /site/totalview/setup.csh
```

and SH/KSH users should run the following:

```
$ ./site/totalview/setup.sh
```

In order to use the debugger, programs must be compiled with the “-g” option. Users should be aware that using this option will produce a larger executable that may run relatively slowly.

The following example demonstrates the compilation of the Fortran program, example1.f:

```
% mpixlf_r -o example1.exe -g example1.f
```

There are several ways TotalView can be started. A few of the more common ones are:

- * Start the debugger and then load a program or corefile:

```
% totalview
```

- * Start the debugger and load the specified program.

```
% totalview filename
```

where filename specifies the name of the executable

- * Start the debugger and load the program specified by filename and its core file.

```
% totalview filename corefile
```

where filename specifies the name of the executable
where corefile specifies the name of the core file

Next is an example of a LoadLeveler script that can be submitted to remotely display an interactive debugging session of a parallel program.

Start of Script

```
@ shell = /bin/csh
@ output = $(jobid).out
@ error = $(jobid).error
@ network.MPI = css0,shared,US
@ job_type = parallel
@ job_name = totalex      # Your Job Name
@ account_no = NA0101    # Your Project Name
@ node = 1
@ tasks_per_node = 4
@ node_usage = not_shared
@ wall_clock_limit = 05:00
@ class = batch
@ queue
# Setup TotalView environment
source /site/totalview/setup.csh
# Set DISPLAY to value set by your SSH login to
# MARCELLUS
setenv DISPLAY f27p4e.navo.hpc.mil:12.0
# Compile the program
mpixlf_r -o example1.exe -g example1.f
# Start an xterm in the background
xterm &
# Start TotalView with poe for parallel executable
totalview poe -a example1.exe -nodes 1 -procs 4
```

End of Script

Job Submission

```
% llsubmit totalex
```

INFORMATION LINKS

For detailed information on TotalView, refer to the following guides:

- * Getting Started: http://www.etnus.com/Products/TotalView/started/getting_started.html

- * User Guide: http://www.etnus.com/Support/docs/rel6/html/user_guide/index.html
- * Reference Guide: http://www.etnus.com/Support/docs/rel6/html/ref_guide/index.html

VAMPIR/VAMPIRTRACE

Vampir, from Pallas, Inc., is an interactive visualization tool designed to analyze the performance and message passing characteristics of Fortran, C, or C++ parallel programs that use the MPI library. It interprets and visualizes a tracefile, generated with the use of an instrumented MPI library, Vampirtrace.

Before the Vampir tool can be invoked, several environment variables must be set. Variables include PAL_ROOT, PATH, and MANPATH.

CSH/TCSH users should run the following command to set up the environment:

```
% source /site/vampir/setup.csh
```

SH/KSH users should run the following:

```
$ ./site/vampir/setup.sh
```

Once the environment is set up, a tracefile can be generated by compiling, linking the Vampirtrace library, and then executing the program. The default library is 32-bit and is located in /site/vampirtrace/lib. A 64-bit library is also available and can be found at /site/vampirtrace/lib64. After executing the code, a tracefile is created that has a ".bvt" extension. The tracefile can then be examined using Vampir.

The following example demonstrates compiling and linking:

```
% mpixlf_r -I $PAL_ROOT/include -c example1.f
% mpixlf_r -L $PAL_ROOT/lib -IVT -lld -o
example1.exe example1.o
```

Below is an example of a LoadLeveler script that can be submitted to remotely display an interactive session of a parallel program.

Script

```
@ shell = /bin/csh
@ output = $(jobid).out
@ error = $(jobid).error
@ network.MPI = css0,shared,US
@ job_type = parallel
@ job_name = vampirex      # Your Job Name
@ account_no = NA0101     # Your Project Name
@ node = 1
```

```
@ tasks_per_node = 4
@ node_usage = not_shared
@ wall_clock_limit = 05:00
@ class = batch
@ queue
# Set up the Vampir/Vampirtrace Environment
source /site/vampir/setup.csh
# Compile/Link Program
mpixlf_r -I $PAL_ROOT/include -c example1.f
mpixlf_r -L $PAL_ROOT/lib -IVT -lld -o example1.exe
example1.o
# Set DISPLAY to value set by your SSH login to
# MARCELLUS
setenv DISPLAY f27p4e.navo.hpc.mil:12.0
# Start an xterm from this script and send it to your
# local workstation and use it like a normal login
# session.
xterm &
#Start Vampir
vampir example1.bvt
```

End of Script

Job Submission

```
% llsubmit vampirex
```

INFORMATION LINKS

For detailed information on Vampir/Vampirtrace refer to the following URL: <http://www.pallas.com/e/products/vampir/documents.htm>.

For more information on running a Window application under LoadLeveler, refer to the following link: http://www.navo.hpc.mil/Navigator/Fall01_Tips.html

UPCOMING EVENTS

USENIX 2003 | 9-14 June 2003 | San Antonio, TX | www.usenix.org

HPDC-12 | 22-24 June 2003 | Seattle, WA | www.csag.ucsd.edu/HPDC-12/

GGF8 | 24-27 June 2003 | Seattle, WA | www.gridforum.org

APS DCOMP CCP 2003 | 20-25 July 2003 | Xi'an, China | www.aps.org/DCOMP/Meetings/meetings.html

SIGGRAPH 2003 | 27-31 July 2003 | San Diego, CA | www.siggraph.org/s2003/conference/

SCICOMP 8 | 5-8 Aug. 2003 | Minneapolis, MN | www.msi.umn.edu/general/Symposia/ScicomP8.html

Oceans 2003 | 22-26 Sept. 2003 | San Diego, CA | www.oceans2003.com

HPEC 2003 | 23-25 Sept. 2003 | MIT, Cambridge, MA | www.ll.mit.edu/HPEC/

PACT 03 | 27 Sept. - 1 Oct. 2003 | New Orleans, LA | www.pactconf.org

Naval Oceanographic Office * MAJOR SHARED RESOURCE CENTER

1002 Balch Boulevard . Stennis Space Center, Mississippi . 39522