



Navigator



NAVO MSRC

SPRING 2002



Image created by the NAVO MSRC Visualization Center staff derived from data generated by the Parallel Ocean Program (POP) model run on the NAVO MSRC IBM RS/6000 SP (HABU).

News and information from...

The Naval Oceanographic Office Major Shared Resource Center



The Director's Corner

Steve Adamec, NAVO MSRC Director

Sweeping Series of Enhancements Engulf NAVO MSRC

The next few months will be notable ones here at the NAVO MSRC. We are completing a sweeping series of center enhancements, designated as Technology Insertion for FY02 (TI-02), across several major technology areas within the MSRC. The centerpieces of TI-02 are the acquisition of a new IBM POWER4 HPC system and completion of the NAVO MSRC Remote Storage Facility (RSF). When complete, the TI-02 enhancements will provide almost 8.5 teraflops of aggregate peak computing capability with commensurately balanced storage and networking capabilities. The RSF will permit us to offer what is perhaps one of the most resilient and best performing HPC storage environments in the world today. This enormous computational capability, coupled with a sustained 11-year NAVOCEANO focus on supporting the largest and most demanding DoD computational applications, will continue to enable unparalleled advances in the DoD science and technology areas served by the HPC Modernization Program (HPCMP).

With all of this diverse computational capability that's been fielded across 20+ shared resource centers (SRCs) by the HPCMP, it has become crit-

ically important for us to redouble our efforts in assessing and implementing common user environments, practices, and tools within and across the SRCs. Your continuing individual and collective user feedback makes it clear that you consider this to be one of your highest priorities for the SRCs. In response, the SRCs have intensified and made significant progress toward strategic cross-cutting collaborative efforts in technical areas such as mass storage and archival, metacomputing, HPCMP-wide shared information environments, and security. Here at the NAVO MSRC, we are redoubling our efforts to supplement those activities by strengthening the links to the new HPCMP Programming Environment and Training (PET) program that's more tightly focused than ever on user environment, tools, and productivity.

My staff and I look forward to seeing you in June at the 2002 HPCMP Users' Conference in Austin, Texas. 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 was generated by the POP model running on HABU, an IBM RS/6000 SP located at the NAVO MSRC. The simulation represents a state-of-the-art eddy-resolving effort that will eventually be used as the ocean component of a coupled global air/ocean/ice prediction system for Navy needs as well as for short-term climate studies.

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
help@navo.hpc.mil

NAVO MSRC Navigator
www.navo.hpc.mil/navigator

NAVO MSRC Navigator is a bi-annual 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 address above.

EDITOR:
Gioia Furness 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 Sweeping Series of Enhancements Engulf
NAVO MSRC

Feature Articles

- 4 Active Swirl and Fuel Modulation to Control
Combustion Instability in Gas Turbine Engines
- 7 NAVO MSRC Visualization Center Supports
the Retrieval of the *Ehime Maru*
- 11 Knowledge Management: A Case for
Intelligent Data Archives

Scientific Visualization

- 13 Improvements in the Use of Color Wheel
Visualization of Global Modeling Data
- 14 Towards a High-Resolution Global Coupled
Navy Prediction System: The Ocean Component

Programming Environment and Training

- 18 VSIPL/ERI: An Enhanced Reference Implementation
of the Vector Signal and Image Processing Standard
- 20 NAVO MSRC PET Update

The Porthole

- 22 A Look Inside NAVO MSRC

Navigator Tools and Tips

- 24 File Transfer Techniques: Transferring Between HABU
and Other Systems

Upcoming Events

- 27 Conference Listings

Active Swirl and Fuel Modulation to Control Combustion Instability in Gas Turbine Engines

Chris Stone and Suresh Menon, Georgia Institute of Technology

Combustion-generated pollutants such as nitrous oxide (N_2O), carbon monoxide (CO), unburned hydrocarbons (UHC), and soot can be reduced significantly if lean burning combustion systems can be used in gas turbine engines. However, as the fuel-air mixture becomes lean, small perturbations in the flow can change the unsteady heat release pattern, and if the heat release is in-phase with the acoustic fluctuations in the combustor, the pressure (acoustic) fluctuations can grow into high-amplitude, low-frequency oscillations. These oscillations can cause flame extinction (often called Lean Blowout, or LBO), reduce engine life span, and/or cause catastrophic structural damage under extreme circumstances. If these oscillations can be controlled and even leaner mixtures can be burned stably, then not only can emissions be reduced but the over fuel consumption can also be decreased significantly. Since these oscillations occur in a dynamic manner, passive

control techniques (e.g., using structural changes) cannot deal with all possible changes, and therefore, in order to suppress these oscillations, active control strategies are needed. This article describes the study of two active control methods applied to premixed combustion in a gas turbine engine: modulation of the swirl imposed on the incoming fuel-air mixture (which is at a fixed equivalence ratio) and modulation of the fuel content in the premixed mixture for a fixed swirl condition. These studies have been conducted using a state-of-the-art Large-Eddy Simulations (LES) model developed at the Georgia Institute of Technology.

The first active control technique studied here employs modulation of the inlet swirl velocity of the incoming premixed mixture. In a typical gas turbine engine, swirl is introduced into the air from the compressor by swirl vanes (typically small airfoils at high angles of attack) in the pre-mixer located

upstream of the combustor. Fuel is then injected into this swirling airflow in the pre-mixer so that a swirling, premixed fuel-air mixture enters the combustor. Swirling inflow is used in dump combustors to exploit a fluid dynamic phenomenon called Vortex Breakdown (VB) that makes the flame zone compact and stabilizes the flame. As the swirling flow expands into the combustor, an adverse pressure gradient is formed in the flow, and this slows down the axial motion of the mixture. For a given geometry, when the swirl intensity exceeds a critical value, a recirculation bubble is formed (called VB) along the centerline of the combustor. This bubble acts as an effective bluff body in




Figure 1(a)




Figure 1(b)

Figures 1(a)/(b). Swirl modulation effect on flame-vortex dynamics: (a) high-swirl condition, (b) low-swirl condition. At the high-swirl state, the vortex rings (purple) rapidly break down, and the flame surface (gray) is compact. On the other hand, at low swirl, the vortex rings are more coherent, and they drag the flame with them, making the flame longer. No VB occurs when the swirl is low, and therefore, the flame in the low-swirl case oscillates and is unstable, whereas at high swirl, VB stabilizes the flame.

the flow, and along the centerline, the flow from the inlet actually stops (stagnates) just upstream of the bubble. This phenomenon helps to stabilize the flame just upstream of the stagnation point. However, in current gas turbine engines, since the swirl vanes are fixed, the swirl introduced into the flow is static and is typically optimized for only one operating condition. This study explores how swirl magnitude can be actively changed to stabilize the combustion as the fuel air mixture is changed from rich to lean.

The second technique studied here is the dynamic adjustment of the fuel content in the fuel-air mixture. For a fixed mass flow rate and a fixed inlet swirl intensity, we dynamically change the fuel-air mixture ratio by changing the equivalence ratio of the incoming mixture. The goal of this study is to determine how changing the fuel content (which, in turn, will affect the amount of heat released and peak temperature reached) will affect the combustion dynamics in the combustor.

The LES methodology is used here to simulate the combustor response to both control techniques. In LES, the large scales of motion are fully resolved in both space and time, and only the small scales (smaller than the grid resolution) are modeled. Premixed combustion is modeled using a thin-flame model in which a scalar field that represents the propagating flame front is evolved, along with the LES momentum and energy transport. To model the fuel-air mixture change (in the second control technique), an additional (conserved) scalar equation is modeled, which tracks the change in local equivalence ratio. The local equivalence ratio is then used to calculate the burning rate and combustion temperature.

The LES algorithm (LESLIE3D) is parallel (using Message Passing Interface) and performs very well on a variety of distributed (Cray T3E (SEYMOUR), RS/6000 SP3 (HABU), etc.) and

shared-memory parallel platforms (SGI Origin 2000/3800). In fact, due to its scalability, LESLIE3D is now a benchmark code for all Department of Defense (DoD) High Performance Computing (HPC) sites and is being used to evaluate new systems' performances. The current study employed 1.3 million grid cells and required approximately 1000 CPU hours (on HABU) for one flow-through time. Typically, 10-15 flow-through times are simulated to obtain sufficient data for statistical analysis.

The first investigation looked at changing the level of swirl imposed at the inlet. In an open-loop study conducted earlier, the pressure fluctuation was shown to reduce when swirl intensity

an increase in swirl stabilizes the flame and reduces the fluctuation in both the inflow mass flow rate and in the pressure fluctuations.

The governing physics of combustion instabilities is quite complicated owing to the fact that not one, but several physical processes are coupled and in contention. The dominant process for combustion instabilities is the interaction of acoustic waves (pressure fluctuations) with the unsteady heat source (in this case, the flame surface). When the pressure (p') and heat-release fluctuations (q') are "in-phase" (the product of the two being positive), energy is added to the instabilities. As the fuel concentration is decreased, two important controlling parameters, the com-

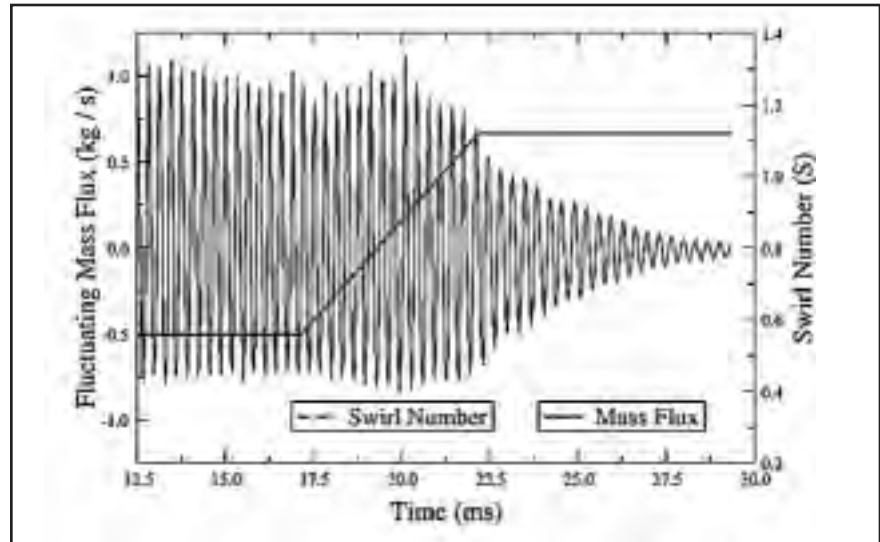


Figure 2. Mass flux measured at the inlet during swirl modulation. At low swirl, the fluctuating mass flux is nearly 100% but drops sharply as the swirl vane angle is increased (i.e., when the swirl velocity is increased). The pressure fluctuation in the combustor is also decreased as the swirl intensity is increased.

was increased. To dynamically reproduce this effect, a simulation was started with a low-swirl S (0.56), and then swirl was steadily increased so that S reached a final value of 1.12. Analysis shows that after a small delay, the pressure oscillation amplitude dropped by 6.6 decibels (dB). Figures 1 and 2, respectively, show the effect of swirl on vortex-flame dynamics and on the mass flow-rate fluctuations. Results show that for a fixed fuel-air mixture,

combustion temperature (flame temperature) and the burning velocity (flame speed), are subsequently reduced. The nearer one is to the LBO limit, the more pronounced is this effect (See Figure 3a). At a lower burning velocity, unsteady flames are more susceptible to large-scale flow perturbations since the recovery time (i.e., the time needed to return to the mean or stable

Article Continues...

Figure 3(a)

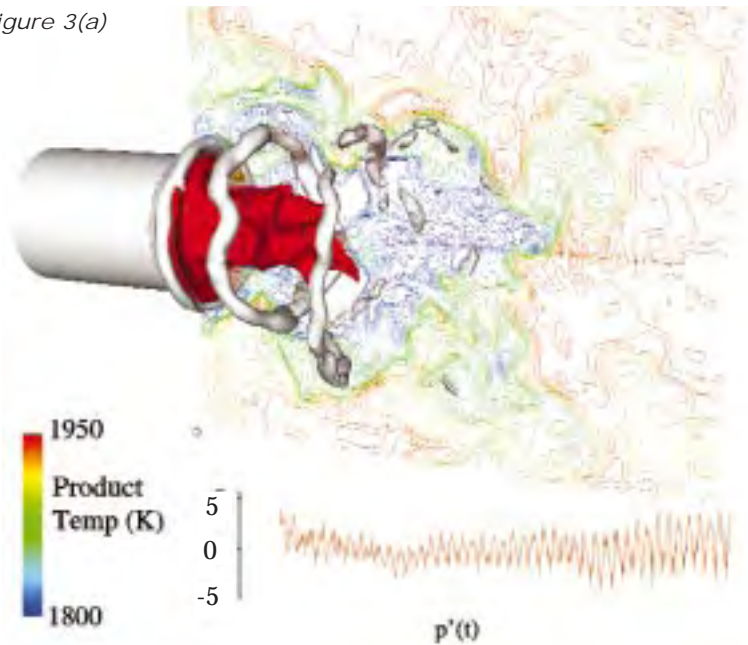



Figure 3(a)/(b). Instantaneous view of fuel modulation control. Shown are the flame surface (red isosurface), vortex rings (gray rings), product temperature (color contours), and fluctuating pressure signal (shown below the figure). In (a), the fuel mixture entering the combustor has become lean, and this results in a decrease in temperature (blue core region), and the pressure oscillation increases. This time corresponds to the far right of the $p'(t)$ signal. In (b), the inflow fuel mixture equivalence ratio is increased (rich mixture), and as a result, the flame temperature increases, resulting in a hotter core region (red temperature contours). At this phase of the control cycle, the pressure fluctuations begin to decrease (first quarter of the $p'(t)$ signal).

location) is lengthened. In the present simulation, the magnitude of flame fluctuations is reduced at higher burning velocities (See Figure 3b). Another contributing factor to the increased stability is the modification of the relevant time scales. As the combustion temperature is altered, the acoustic velocity (speed of sound) is also modified, resulting in a shift in the phase between the pressure and heat-release fluctuations. This effect can be significant in either increasing or decreasing the instabilities, depending on the phase between the p' and q' . In some combustion devices, instabilities are more intense for higher equivalence ratios due to this effect, while, as simulated here, the reverse may also occur. Simulations such as these can be used to understand the complex nature of combustion instability that is a result of nonlinear interactions between the three modes of wave motion in a compressible reacting

flow in a confined domain: acoustic waves, vortex motion, and unsteady heat release due to combustion. Active control of this instability, as demonstrated here, can then be used

not only to achieve stable combustion in a regime that is not possible otherwise, but also to understand how the dynamics have been changed due to control actuation.

Figure 3(b)



Acknowledgements

This work was supported in part by General Electric Power Systems and Army Research Office (ARO). Computational support was provided by the DoD High Performance Computing Modernization Program Office (HPCMO) at HPC Major Shared Resource Center (MSRC) sites (Naval Oceanographic Office (NAVO), U.S. Army Engineer Research and Development Center (ERDC), and Army Research Laboratory (ARL)), under an ARO Challenge Project.

NAVO MSRC Visualization Center Supports the Retrieval of the *Ehime Maru*

Pete Gruzinkas, NAVO MSRC Visualization Center

On 9 February 2001, a collision between a U.S. Navy submarine and a Japanese fishing trawler sent the fishing vessel to the bottom of the Pacific Ocean about 10 miles south of Diamond Head, Oahu, Hawaii. The decision was made to retrieve the vessel, which lay in approximately 1800 feet of water. A Crisis Action Team was formed, and the NAVO MSRC Visualization Center was asked to provide visualization support for the *Ehime Maru* Retrieval Operation. There are many support aspects in an operation like this, but the first challenge was to visualize the high-resolution model output generated by the Shallow-Water Analysis and Forecast System (SWAFS) a three-dimensional (3D) ocean circulation model that produces time-series images with correct

speed and direction, temperature, and salinity fields. Initially, a 2-km grid that covered the entire Hawaiian Islands was generated, but this model provided input to a high-

er resolution (500-m) nest, that bounded the operation area just south of Oahu. The model was run on the newest and most powerful supercomputer at the NAVO MSRC, the RS/6000 SP3 (HABU).


Part of the NAVO MSRC Visualization Center mission is the development of analytical software environments for the Department of Defense (DoD) research community. Some of these tools, designed to analyze ocean model output, were modified to analyze the Navy's operational model output. These software tools provided significant diagnostic capability, which assisted in the validation of the model output in a very complex environment.

While virtual environments may not provide all of the answers to data analysts, the




Screen capture of SWAFS probe compared to the ADCP current measurements.

Article Continues...



SWAFS surface currents in the operational area.



Vertical profile of SWAFS over the *Ehime Maru*.

fact of the matter is the real world environment is in 3D. Features within the environment have 3D structure that can be difficult, if not impossible, to realize in two-dimensional (2D) space. The same technology that was built to help research and development modelers scrutinize their model output was applied to an operational scenario where the speed and direction of the ocean currents were critical factors.

An additional feature added to the application, unique to the *Ehime Maru* retrieval operation, was the display of an Acoustic Doppler Current Profiler (ADCP) buoy. The display of the ADCP data provided analysts a direct comparison between model output and in-situ current measurements in near-real time. Another significant feature of the virtual environments built for


ocean model analysis is portability — the ability to run on a variety of hardware architectures, including laptop computers.

Portability is accomplished by savvy application of graphics techniques, as well as smart Input/Output and memory management. This ability


was critical to the provision of an analysis environment and data to forward-deployed personnel on-scene at the recovery site. Other support products were developed to render the recovery area in 3D from high-resolution bathymetry (water depth) data to delineate critical areas. Conceptual animations

were built to demonstrate the mechanics of an extremely difficult recovery operation.


The operation was successful, and all mission objectives were accomplished. This recovery operation demonstrates the excellent synergy that has developed between the operational Navy and the DoD research infrastructure built by the High Performance Computing Modernization Office (HPCMO), and is a stellar example of the NAVO MSRC realizing the HPCMO maxim of "Delivering Science to The Warfighter."




3D shaded relief of OPAREA using high-resolution bathymetry collected by the USNS SUMNER. Imagery shows where the ship was resting, the shallow-water retrieval area, and the final relocation site.




Screen capture showing ocean currents at depth over an exaggerated (3:1) model of the *Ehime Maru*.



Another snapshot of currents at depth. Note the legend has been "tuned" to accentuate faster currents.



Still images from conceptual animation showing retrieval of the *Ehime Maru* by the vessel *Rockwater*.



Still image from conceptual animation showing the lift procedure. A conceptual animation was generated depicting how the towing harness would be attached to the *Ehime Maru*, which rested in approximately 1800 feet of water.



Still image showing *Ehime Maru* retrieval over real bathymetry data. The animation was designed to conceptualize/visualize what the recovery would look like in the context of the real data collected by the USNS SUMNER.



Oceans 2002
Conference & Exhibition

October 29-31

Mississippi Coast Coliseum
& Convention Center

Biloxi, Mississippi

Proposed technical sessions at the OCEANS 2002 MTS/IEEE Conference will focus on the following areas:

- I. Advanced Marine Technology
- II. Marine Resources
- III. Ocean and Coastal Engineering
- IV. Marine Policy and Education
- V. Fisheries Technology
- VI. Information Technology
- VII. Ocean Modeling
- VIII. Integrated Ocean Observing Systems
- IX. Communications and Navigation
- X. Underwater Acoustics Technology
- XI. Non-acoustic Signal and Image Processing
- XII. Shallow Water Environmental Technology

The technical program will include a student poster competition and tutorial sessions. An exciting exhibition featuring innovative marine technology products and services will be held at the Mississippi Coast Coliseum & Convention Center. Exhibitors will be given passes to attend the technical sessions.

www.OCEANS2002.com



Knowledge Management: A Case for Intelligent Data Archives

Douglas D. Cline, Ph.D., Program Manager, Northrop Grumman Information Technology

There is little doubt that one of the most significant challenges facing the Department of Defense (DoD) High Performance Computing Modernization Program today is intelligently dealing with the sheer magnitude of data that are created, collected, analyzed, and manipulated

As the data archive approaches the petabyte range, it is clear that data management initiatives must be undertaken by the DoD to create structured, intelligent data archives that provide a broad range of functionality. Such initiatives will prove difficult to implement without first

Management (HSM) software, networking, and file servers — are all critical components of the system. Scalability is essential if future capacity demands are to be met through incremental expansion of the components. As data capacity expands with time, so do related storage costs. If the growth in data storage is exponential, then, arguably, storage costs will also grow exponentially. Over time, "data inertia" can become so great that the flexibility needed to migrate the active archive to newer, lower cost storage technologies is problematic at best.

From an end-user's perspective, the traditional archival system is inherently limited when supporting higher level access to data. Generally, access to data stored in the archive is low-level via HSM-specific commands. For example, `get /path/file`, `put/path/file`, and `copy /path/file` are all command-driven directives for accessing archived data. Each directive requires specialized knowledge of the HSM commands (e.g., `get`, `put`, `copy`) as well as a specific pathname that points to a particular file (e.g., `/path/file`). This knowledge is typically limited to a single user or small community where the specialized knowledge has been shared. An individual outside this group, not privileged with this special knowledge, would find it very difficult to locate the data set, much less extract useful information from it. Therefore, one of the primary limitations of the traditional storage archive is its inability to support "Information Discovery," a critical missing capability for long-term persistent data storage.

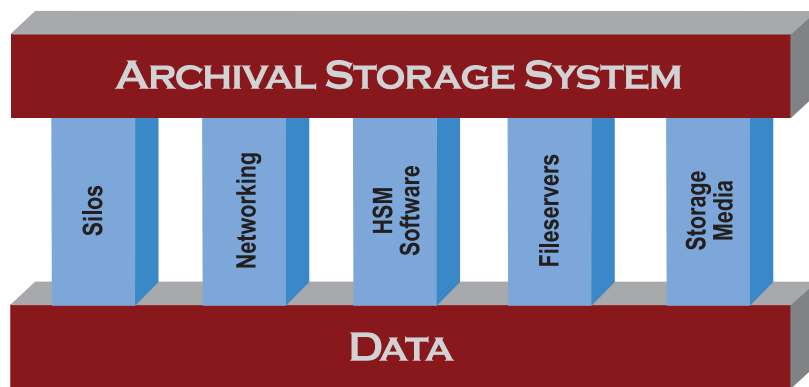


Figure 1. Conceptual Framework for the Current MSRC Archival Storage System.

throughout the MSRCs and Distributed Centers (DCs). Increasing computational capabilities, as demonstrated by faster HPC systems and larger system memories, enable high-fidelity simulations that provide critical information to support the nation's warfighters. In addition, through advances in sensor technologies, scientific data are being collected and archived at an unprecedented rate. These two factors create what is known as "infrastructure stress" — a fundamental stress applied to the current systems and data storage infrastructure. As of 1 March 2002, the NAVO MSRC data archive surpassed 375 terabytes (TB), with a projected growth rate of approximately 14 TB per month.

empowering end-users with a suite of high-level data management tools for organizing, browsing, and accessing long-term archival data. Development of these tools will require new perspectives on the MSRC program to leverage existing storage architectures with newer, more robust technologies.

TRADITIONAL ARCHIVAL STORAGE

Archival storage provides the means to preserve large-scale computational and experimental data indefinitely. A conceptual framework for a traditional archival storage system is shown in Figure 1, with "Data" representing the foundation of the system. The various infrastructure elements — tape silos, tape media, Hierarchical Storage

Article Continues...

KNOWLEDGE MANAGEMENT SYSTEMS

By leveraging the MSRC's current storage infrastructure, higher level functionality can be achieved with nominal investment to create a more robust, intelligent data archive. For the purposes of this article, the more gen-

software tools and applications to isolate them from the implementation details of the underlying storage architecture.

Middleware typically consists of Object Brokers such as Common Object Request Broker Architecture

such as Information Discovery are supported, by referencing data files through their data attributes rather than physical file location.

BEYOND DATA ARCHIVES - DIGITAL LIBRARIES

This article has examined the basic framework for two storage systems — a traditional storage archive representing the MSRC's current architecture and a more robust framework representing a Knowledge Management System. The Knowledge Management System supports a broad range of data access tools and applications to facilitate data manipulation from remote clients. The key aspect of the Knowledge Management System is that it leverages the MSRC's existing storage infrastructure (i.e., silos, file-servers, and networks) by introducing Middleware to deliver a new class of higher level data services to the end-user.

Beyond these simple frameworks lie more sophisticated storage architectures. Digital libraries offer the potential to create online knowledge centers for discipline-specific application areas that are important to the DoD. These libraries generally provide shared information resources, coupled with high-speed networks, to deliver high-level information content directly to the desktop. Client-based tools and applications can be specifically designed to enable data manipulations for analysis and interpretation.

The technical challenges posed by large-scale data management presents, in reality, new opportunities for infrastructure development within the DoD modernization program. In exploiting these opportunities, new requirements and capabilities will be defined that will ultimately deliver a new class of data services for users of the MSRCs and DCs.

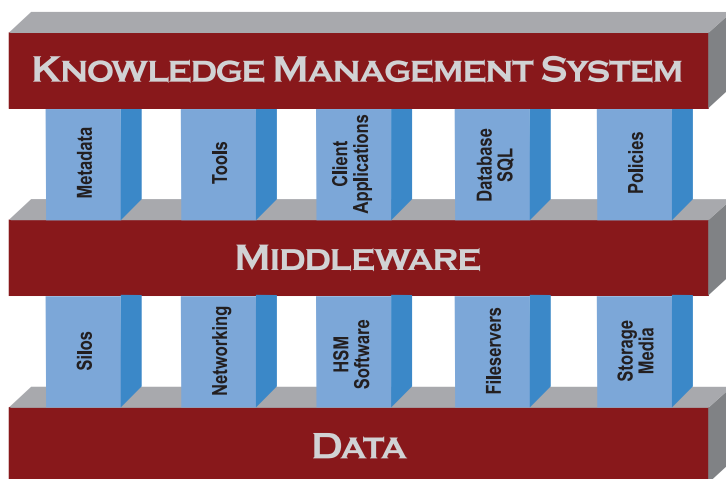


Figure 2. Conceptual Framework for the MSRC Knowledge Management System.

eral type of archival storage will be referred to as a Knowledge Management System.

In principle, the Knowledge Management System should be capable of supporting Information Discovery, including browsing, searching, and retrieving archived data. The conceptual framework of a Knowledge Management System is shown in Figure 2. All essential infrastructure elements of the archival storage system are present, as in Figure 1. However, a new foundation layer is introduced in our Knowledge framework, Middleware, which supports a set of higher level infrastructure elements.

The Middleware provides a software layer to interpret high-level data access functions into low-level HSM-specific commands. Moreover, it acts as an interface between sophisticated

(CORBA), or, more specific to storage architectures, the Storage Resource Broker (SRB), developed at the San Diego Supercomputing Center. In the Knowledge Management framework, changes to the storage architecture can be introduced without impacting the functionality of user tools and applications, thereby protecting software investments and preserving operational continuity.

Of critical importance to the Knowledge Management System is metadata, or "Data about the Data." The metadata represents a database of data attributes, with linkages to physical file locations in the archive. Data attributes may contain information such as data types, data format or representation, data descriptors, and date of creation. It is through the metadata that higher level functions

References

1. Chaitanya Baru, Reagan Moore, Arcot Rajasekar, and Michael Wan, "The SDSC Storage Resource Broker," Proc. CASCON '98 Conference, Nov. 30-Dec. 3, 1998, Toronto, Canada.

Improvements in the Use of Color Wheel Visualization of Global Modeling Data

Ludwig Goon, NAVO MSRC Visualization Center

The color wheel method of two-dimensional (2D) vector field visualization uses the Hue-Saturation-Value (HSV) color space to map vector direction to hue (color) and magnitude to saturation and value. This visualization technique is better than traditional graphics primitives (arrows) in high-resolution models because it offers less visual congestion. In addition, the color wheel technique more effectively displays eddy formation and vector magnitudes in circulation models.

Two major changes to the initial color wheel technique, developed at Mississippi State University as a component of the High Performance Computing Management Office Programming Environment and Training Program working within the Climate, Weather, and Ocean Modeling arena, are used in POPFlow (an interactive graphical viewer of Parallel Ocean Program (POP) modeling data) to achieve interactive graphics performance. First, the Alpha channel is linked to the vector magnitude. Normalizing vector magnitude gives a direct correlation to transparency or opacity. As the




Figure 1. Global image of POP model data displayed using the original color wheel method.

magnitude of the vector approaches the maximum, the pixel contains more color, and vice versa. Second, a simpler color lookup table is generated. The HSV colorspace complicates calculations with converting the saturation and value components to the Red-Green-Blue (RGB) colorspace, the de facto colorspace in OpenGL. The hue component has 360

degrees directly correlated to direction (i.e., red is 0 degrees and is due east). Previous color wheel techniques either converted the HSV to RGB or used simpler color wheel lookup tables. The POPFlow version uses 360-color entries converted from the HSV colorspace. POPFlow is built on the OpenGL and GLUT application programming interfaces (API) with

hardware alpha and blending. The POPFlow application using the APIs, and the color wheel optimizations easily show a year's worth of vector data, allowing the user to interact with ease.

Color wheel optimizations in POPFlow show impressive global animations of major effects like the Gulf Stream currents, the Kuroshio currents off the eastern coast of Japan, and equatorial striations. Figure 1 is a global image of the POP model using the original color wheel method. Figure 2 shows the enhanced color wheel technique in the POPFlow application. Notice that in POPFlow, the alpha blending of the color wheel layer does not occlude the bathymetry below. In addition, POPFlow has interactive adjustments for vector magnitudes to allow users to set threshold wheels that serve to filter these magnitudes.

Future efforts will include dynamic color map manipulation and incorporation of a similar technique to show convection in fluid flow models. For more information please contact the NAVO MSRC Visualization Center staff at viz@navo.hpc.mil.




Figure 2. Demonstration of the enhanced color wheel technique used in the POPFlow application.

Towards a High-Resolution Coupled Navy Prediction System

Julie McClean, Naval Postgraduate School

Mathew Maltrud, Los Alamos National Laboratory

The vision for Navy meteorological and oceanographic forecasting is that of a high-resolution global coupled air/ocean/ice prediction system with very-high-resolution regional air/ocean coupled models nested into the global system at key locations. The development of such a coupled prediction system is a computational Challenge problem requiring commensurate supercomputing resources. Towards this goal the authors are spinning up an eddy-resolving global ocean model; upon completion it will be delivered to the Fleet Numerical Meteorology and Oceanography Center (FLENUMMETOCEN) for future transition to the operational environment.

Any ocean model to be used in such a prediction system must be capable of simulating high-frequency (days to several months) and short-scale (10 to 1000 km) processes. In particular, it must be capable of producing realistic mean and varying surface and thermocline flows. High horizontal and vertical resolutions are required to produce not only mesoscale variability, but also the characteristics of strong, narrow currents such as the Gulf Stream. Additionally, the model must be capable of producing realistically varying thermal structures in the upper water column and therefore requires a mixed-layer parameterization. Satisfying these requirements is computationally demanding; for this simulation we have a Department of Defense (DoD) High Performance

Computing Modernization Office (HPCMO) Challenge Grant. Its execution is taking place on the Naval Oceanographic Office's RS 16000 SP3 (HABU).

OCEAN MODEL DESCRIPTION

The ocean model used in this study is the Los Alamos National Laboratory (LANL) Parallel Ocean Program (POP). It is a primitive equation z-level model with a free-surface boundary condition. Approximations to governing fluid dynamics equations permit a decoupling of the model solution into barotropic (vertically averaged) and baroclinic (deviations from vertically averaged) components; these are solved using an implicit elliptic scheme and an explicit parabolic equation system, respectively. It is written in Fortran90 and is designed to run on multi-processor machines using domain decomposition in latitude and longitude. Message Passing Interface (MPI) is used for inter-processor communications on distributed memory machines and Shared Message Interface (SHMEM) on shared memory machines. Benchmarking shows the code to be highly scalable onto a large number of processors, provided the processor sub-grid is large enough.¹ Further technical details and references regarding the code and its adaptation for massively parallel

Resolution Global Prediction System: The Ocean Component

computers can be obtained from
<http://climate.acl.lanl.gov>.

A global 0.1-degree, 40-level configuration of POP is being spun-up on a displaced pole grid (See Figure 1), whereby the North Pole is rotated into Hudson Bay to avoid a polar singularity. The grid consists of 3600x2400x40 grid points, and the spacing is about 11 kilometers (km) at the equator, decreasing to about 3 km in the Arctic Ocean. At mid-latitudes this spacing is 5-7 km. A blended bathymetry was created from Smith and Sandwell,² International Bathymetric Chart of the Arctic Ocean (IBCAO),³ and British Antarctic Survey (BEDMAP)⁴ products. All important channels and sills were checked and modified to facilitate correct flow.

The model was initialized from an ocean state depicted by temperature and salinity from the Navy's 1/8-degree January Modular Ocean Data Assimilation System (MODAS) climatology outside of the Arctic and the University of Washington's Polar Hydrography winter climatology in the Arctic (<http://psc.apl.washington.edu/Climatology.html>). To produce

an energetically realistic ocean, synoptic surface forcing is used whenever possible during a 20-year adjustment of the model to its initial state. This forcing was largely constructed from National Center for Environmental Prediction (NCEP) fluxes for 1979-1998.⁵ A mixed layer formulation is active.⁶ To date, about a decade of the simulation has been completed.

THE GLOBAL SIMULATION


The realistic simulation of sea surface temperature (SST) will be important both to the prediction of fronts and eddies and for coupling to the ice/atmosphere components of the future prediction system. A snapshot of temperature in the uppermost model level on 1 February 1987 is seen in Figure 2. To display the whole globe at the full model resolution, the field is plotted in model coordinates; hence, the distortion in the northern hemisphere. Well-defined frontal structures associated with the western boundary currents

Article Continues...

Figure 1. Spacing (km) of 0.1°, 40-level global POP grid.

and tropical instability waves are observed. The large meander of the Kuroshio Current seen to the south of Japan is one of its known meander states at this location; movie loops indicate that the model also reproduces the other observed states. Mesoscale features are apparent in all ocean basins; particularly striking are the Agulhas Current eddies that are shed south of Africa.

To more fully appreciate the mesoscale activity in the model, the flow in the upper level of the model is seen in Figure 3. Here the direction and speed of the flow are represented by color and hue, respectively.



Sea Surface Temperature (°c)




Figure 2. Snapshot of temperature in the uppermost model level (5 m) on 15 January 1987.

The Agulhas Retroflection Current system is enlarged to highlight the shedding of eddies, their sense of rotation, and their northwest movement across the South Atlantic Ocean. The characteristic reversing zonal flows of the tropics can be easily identified along with the associated standing meander patterns of the North Equatorial Countercurrents in the western Atlantic and Pacific oceans. Rich eddy fields are associated with the other boundary flows and the Antarctic Circumpolar Current (ACC) in the Southern Ocean.

Figure 3. Snapshot of surface velocity from the uppermost model level (5 m) on 1 January 1983. Color and hue represent current direction and intensity, respectively. The expanded region demonstrates eddy shedding by the Agulhas Retroflection.

To understand the interactions between the large-scale flows of the Western Pacific, the Asian marginal seas, and the Indonesian Throughflow, a snapshot of surface velocity from 1 June 1987 is displayed (See Figure 4). This is a challenging region for an ocean model to simulate realistically, and is therefore a good test of the model's capabilities. Again, direction and strength are indicated by color and color intensity, respectively. Grid distortion is apparent in the northern part of the domain, particularly in the Kuroshio Extension. The flow in this region is complex and punctuated by mesoscale activity; regardless, details are in qualitative agreement with observations.⁷

The North Equatorial Current (NEC), identified by the northwestward (blue) flow off-shore of the Philippines, bifurcates near the



coast to flow northwestward as the Philippines Current and southward as the Mindanao Current. The Kuroshio is identified as the strong northeastward current off Taiwan; to the north an intrusion of the current is seen in the East China Sea. The Tsushima Current enters the Japan Sea and continues in a northeastward direction along the coast of Japan. The Mindanao Current turns southeastward to enter the Celebes Sea, and southward flow is seen between Borneo and Sulawesi. Flow into the Indian Ocean is seen through all the Indonesian straits to the south of these islands. The circulation in this region is dominated by the seasonal monsoons, and southward flow into the Indian Ocean is consistent with observations at this time of year.⁸ In February, the surface flow reverses, and the net flow is away from the Indian Ocean. The Indonesian Throughflow is the major low-latitude conduit of properties between the Indian and Pacific oceans and is therefore important on both short-term and climatic time scales. Quantitative evaluations and analyses of the spin-up are under-

way using data sets with global or near-global coverage such as those obtained from satellite altimeters, surface drifting buoys, and vertical profilers. We are comparing modeled and observed mean flows, transports through key passages, energy levels, intrinsic scales, and mixed layer variability, among others.

In this manner we are able to assess the progress of the spin-up and

make adjustments as appropriate. Following the 20-year spin-up, a decade-long simulation starting in the 1990s will be performed using high-frequency Navy Operational Global Atmospheric Prediction System (NOGAPS) surface forcing. Further details of the model run, benchmarking, results, and movie loops of the spin-up are available at <http://www.oc.nps.navy.mil/navypop>.




Figure 4. Snapshot of surface flow on 1 June 1987 in the Western Pacific, the Asian marginal seas, and the Indonesian Throughflow.

References

1. McClean, J.L., W. Maslowski, and M.E. Maltrud, 2001: Towards a Coupled Environmental Prediction System, Computational Science - ICCS 2001 (Eds: Alexandrov, Dongorra, Juliano, Renner, and Tan), Lecture Notes in Computer Science 2073, Springer-Verlag, pp 1098-1107.
2. Smith, W.H.F., and D.T. Sandwell, 1997: Global seafloor topography from satellite altimetry and ship depth soundings, Science, 277, p. 1957-1962. http://topex.ucsd.edu/marine_topo
3. Jakobsson, M., N.Z. Cherkis, J. Woodward, R. Macnab, and B. Coakley, 2000: New grid of Arctic bathymetry aids scientists and mapmakers, EOS Trans., American Geophysical Union, 81 (9).
4. British Antarctic Survey (BEDMAP), <http://www.antarctica.ac.uk/bedmap>
5. Doney, S.C., S. Yeager, G. Danabasoglu, W.G. Large, and J.C. McWilliams, 2002: Modeling global oceanic interannual variability (1958-1997): Simulation design and model-data evaluation. "J. Climate," submitted.
6. Large, W.G., J.C. McWilliams, and S.C. Doney, 1994: Oceanic vertical mixing: a review and a model with a nonlocal boundary layer parameterization. Rev. Geophys., 32, 363-403.
7. Tomczak, M., and J.S. Godfrey, 1994: Regional Oceanography: An Introduction. Pergamon Press, New York, 422 pp.7. Figure 8.6; <http://www.es.flinders.au/~mattom/regoc/text/8circ.html>
8. Gordon, A.L., and J.L. McClean, 1999: Thermohaline stratification of the Indonesian Seas-model and observations. J. Phys. Oceanogr., 29, 198-216.

Acknowledgements

This work is supported by the Office of Naval Research, the National Science Foundation, and the Department of Energy (CCPP Program). Computer time is provided through the Department of Defense High Performance Computing Modernization Office at the Naval Oceanographic Office (NAVOCEANO) as part of a Grand Challenge Award. The NAVO MSRC Visualization Center developed the applications used for the representation of the velocity fields. Detelina Ivanova (NPS) created the graphics.

VSIPL/ERI: An Enhanced Reference Implementation of the Vector Signal and Image Processing Standard

Gary Boudreaux and Dr. Tony Skjellum, MPI Software Technology, Inc.

As a result of the NAVO MRSC Programming and Environmental Training Program (PET) Tiger Team Collaboration Final Report,¹ the creation of a new performance Vector Signal and Image Processing Library (VSIPL)² that interfaced with optimized signal processing and linear algebra libraries was assigned by PET to MPI Software Technology Inc. (MSTI). The goal of the four-month project was to begin with the Tactical Advanced Signal Processing (TASP) reference implementation of VSIPL and to "wrap" function calls to the Linear Algebra Package (LAPACK), Basic Linear Algebra Subroutines (BLAS), and Fastest Fourier Transform in the West (FFTW) optimized libraries. This new software was entitled VSIPL Enhanced Reference Implementation (VSIPL/ERI).

BACKGROUND

VSIPL, a community standard, consolidated and streamlined existing mathematical libraries and defined standard functions for the scientific and engineering community. VSIPL functions include Linear Algebra, Signal Processing (i.e., Fast Fourier Transform (FFT) Filters), and Scalar, Vector, and Matrix functionality. The main features are its portability, object-based description, and run-time modes for both development and production. The library also features opaque objects such as blocks, views on the blocks (i.e., vectors, matrices, and tensors), explicit memory/algorithm hints, and public and private data arrays. As an active member of the VSIPL forum, MSTI ensures that its current implementations of the libraries conform to the VSIPL standard.

As a baseline of what is available for embedded computing, MSTI's commercial implementation of VSIPL uses `ativec` and assembly language optimizations to achieve high performance on element wise and signal processing functions. It contains advanced C++ programming

to enhance the performance of the TASP⁴ VSIPL reference implementation by using commercial-off-the-shelf (COTS) lower level libraries, instead of the TASP kernels.

This was accomplished by replacing sections of the TASP VSIPL reference implementation with wrappers to

FIGURE 1. Comparative FFT Benchmarks

ERI
UDP, a dual 800MHz PIII with 256M ram. ***ALL TIMES milliseconds***

benchfft_f				benchfft_d			
	256	1024	2048		256	1024	2048
rcfftsp	0.011	0.058	0.128	rcfftsp	0.015	0.078	0.168
crfftsp	0.012	0.058	0.127	crfftsp	0.016	0.079	0.174
ccfftsp	0.019	0.089	0.218	ccfftsp	0.031	0.152	0.332
ccfftsp	0.024	0.12	0.258	ccfftsp	0.033	0.188	0.364

TASP

benchfft_f				benchfft_d			
	256	1024	2048		256	1024	2048
rcfftsp	0.037	0.148	0.371	rcfftsp	0.028	0.158	0.405
crfftsp	0.039	0.157	0.389	crfftsp	0.041	0.186	0.421
ccfftsp	0.057	0.31	0.623	ccfftsp	0.104	0.564	1.177
ccfftsp	0.054	0.294	0.591	ccfftsp	0.102	0.571	1.167

concepts, and the source is actually a C++ implementation with C APIs. MSTI supports the VSIPL Core Profile for several operating systems and the PowerPC (G4) and Intel (PIII) processors. However, such software has not, as yet, been made available for the MSRCs' current platforms.

Thus, an interim solution for higher performance VSIPL is required. VSIPL/ERI³ satisfies this requirement. In particular, VSIPL/ERI is freely available middleware for United States (U.S.) Government Department of Defense (DoD) MSRC sites. The project requirements were

BLAS, LAPACK, and FFTW. Such libraries are available on many sites. This work involved the identification of the appropriate functions in the libraries to wrap, the creation and extraction of pointers to the appropriate data from the VSIPL data structures, compliance with VSIPL data alignment restrictions (with regards to memory management), and the resolution of problems associated with the C-to-Fortran interface. The project required the port of the VSIPL/ERI software to four platforms: IBM AIX (Power 3); SOLARIS/SPARC; Linux (x86); and MIPS (SGI, R10K). Currently, the

library has been delivered for SPARC, Linux, and AIX. MSTI has been granted access to the NAVO SGI machine, and the entire project was completed before 15 September 2001.

RELATED TECHNOLOGIES

The VSIPL/ERI library utilizes the FFTW library for all FFT functions, LAPACK library for solving advanced systems of linear equations, and the BLAS library for basic linear algebra to enable it to exceed TASP VISPL performance.

FFTW4, developed by Dr. Matteo Frigo and Steven G. Johnson at the Massachusetts Institute of Technology (MIT), is a C subroutine library for computing the Discrete Fourier Transform (DFT) in one or more dimensions, of both real and complex data, and of arbitrary input size. The library is known for its fast performance and excellent portability between architectures. The latest official version is 2.1.3, and in 1999 FFTW received the J. H. Wilkinson Prize for Numerical Software.⁵

In addition to FFTW, VSIPL/ERI employs LAPACK⁵ routines. As a Fortran-77 subroutine

Figure 2: Comparative Linear Algebra Benchmarks
ERI
UDP, a dual 800MHz PIII with 256M ram.
Function_f
ALL TIMES milliseconds

size	Inputs	M=256, N=128, K=64	M=128, N=64, K=32
CCholesky_f	2	112.5	11.58
Cholesky_f	1	18.79	2.4
CLUD_f	1	175	21.4
LUD_f	2	535	10.31
Gemp_f	1	282.5	13.5
CGemp_f	1	3310	62.5
Mprod_f	1	252.5	13.04
MVprod_f	2	0.411	0.11
CMVprod_f	2	1.36	0.255
Vouter_f	2	0.47	0.12
CVouter_f	3	1.96	0.3
Covsol_f	3	127.5	20.07
CCovsol_f	3	188.95	20.78
LLSQ_f	3	l=3.8 40.88	l=0.95 5.87
QRD_f	2	17.5	2.59
CQRD_f	2	540	54.08

TASP

Function_f	size	Inputs	M=256, N=128, K=64	M=128, N=64, K=32
CCholesky_f	2		167.5	15.78
Cholesky_f	1		77.5	7.3
CLUD_f	1		357.5	28.44
LUD_f	2		660	11.25
Gemp_f	1		687.5	16.87
CGemp_f	1		3100	67.5
Mprod_f	1		350	20.94
MVprod_f	2		0.41	0.11
CMVprod_f	2		1.56	0.29
Vouter_f	2		0.5	0.13
CVouter_f	3		1.39	0.22
Covsol_f	3		77.5	10.23
CCovsol_f	3		166.88	20.7
LLSQ_f	3		l=5.06 75.57	l=1.29 9.96
QRD_f	2		61.88	7.77
CQRD_f	2		1125	70.31

library, the well-known LAPACK Application Programming Interface (API) routines provide solutions to systems of simultaneous linear equations, least-squares solutions of linear systems of equations, eigenvalue problems, and singular value problems. In addition, LAPACK also includes other matrix factorizations such as Logical Unit (LU), Cholesky, QR, and SVD, and its library is equipped to handle both dense and banded matrices.

For all routines, LAPACK provides functionality for real and complex matrices in both single and double precision. These routines are written so that the maximum amount of computation is handled through calls to the BLAS. LAPACK was designed to exploit the Level-3 BLAS⁶ — a set of specifications for Fortran subprograms that perform various types of matrix multiplication and the solution of triangular systems with multiples right-hand sides. The coarse granularity of the Level-3 BLAS operations causes it to promote high efficiency on many high-

Article Continues Page 26...

References

1. NAVO MRSC PET Tiger Team Collaboration Final Report http://www.sdsc.edu/PET/NAVO/reports/vsip/VSIPL_final_report.html
2. VSIPL <http://www.vsipl.org>
3. VSIPL/ERI <http://www.mpi-softtech.com/VsiplERI>
4. TASP <http://www.tasp.org>
5. FFTW <http://www.fftw.org>
6. LAPACK/BLAS <http://www.netlib.org>

Contacts

Gary Boudreaux, gary@mpi-softtech.com and Dr. Tony Skjellum, tony@mpi-softtech.com

NAVO MSRC PET Update

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

The new PET program is approaching the end of its first year. As with any new start, it has been a tremendous learning experience for all involved.

At the NAVO MSRC, as reported in the last issue of the *Navigator*, we are primarily responsible for 3 of 15 PET Functional Areas (FAs). Those three areas are the two Computational Technology Areas (CTAs) of Climate, Weather, and Oceanography Modeling (CWO) and Environmental Quality Modeling and Simulation (EQM). The NAVO MSRC PET program is also responsible for taking the lead on the cross-cutting functional area of Computational Environments (CEs). More information about these functional areas can be found on the High Performance Computing Modernization Program Office (HPCMO) web site under the PET link.

Most of our positions have been filled or are in the process of being filled. We have had a few changes, so the following is a who's who at the NAVO MSRC PET program:

Government Lead: Eleanor Schroeder

Component Point of Contact (POC): Dr. George Heburn
// Executive Assistant: Tammy Smith

// Technologists: Andrew Schatzle and Brian Tabor

CWO FA POC: Dr. Jay Boisseau

// CWO On-site (NAVO MSRC): Dr. Tim Campbell

// CWO On-site US Army Engineer Research and Development Center (ERDC): Dr. Phu Luong

// CWO On-site (Monterey): Applicant being interviewed

EQM FA POC: Dr. Mary Wheeler

// Environmental Quality Modeling and Simulation (EQM) On-site (ERDC): Drs. Lea Jenkins and Victor Parr (interim)

// EQM On-site (NAVO MSRC): Position to be filled


CE FA POC: Dr. Shirley Moore

CE On-site (NAVO MSRC): Dr. Tom Cortese

The three FA POCs have submitted the PET strategies for each of their functional areas. Those strategies can be found in the sidebar to this article.

One of the highlights of PET is the refurbishing (many thanks to the NAVO MSRC) of the equipment in our PET classrooms. We now have 20 new single-processor Pentium 4 workstations operating at 1.8 gigahertz with 1 gigabyte of memory and a 20-gigabyte removable hard drive, which gives us the capability of hosting either Linux-based or Windows-based classes. The equipment was given its test run at the HPCMO-sponsored Service/Agency Approval Authority meeting that we hosted here in February.

One final note, PET will undergo its first review in May 2002 in Vicksburg, MS.



CWO STRATEGY

DR. JAY BOISSEAU, UNIVERSITY OF TEXAS, AUSTIN,
CWO FA POC

The MOS team supports the DoD climate/weather/ocean (CWO) modeling community in the new PET program. The CWO support strategy was developed by the MOS team in the PET recompetition and has been refined during the program with input from many DoD users. The MOS team for CWO includes:

- George Heburn (Mississippi State University (MSU)), Component Lead for the MOS team overseeing CWO, EQM, and CE
- Jay Boisseau (Texas Advanced Computing Center (TACC)/University of Texas (UT) at Austin), coordinator for MOS CWO support activities
- Tim Campbell (MSU) and Phu Luong (ERDC), on-site CWO support staff
- DoD CWO User Advisory Panel: Alan Wallcraft (NRL-Stennis), Jane Smith (ERDC), Rich Hodur (NRL-Monterey), and Frank Ruggiero (AFRL), plus the CTA leader for CWO, Bill Burnett (CNMOC)

The PET CWO support strategy is focused on the development, implementation, and support of advanced computing technologies that are needed to enhance the current and next generations of strategic CWO applications. Since the DoD CWO community is among the world leaders in atmospheric and ocean physics, the PET CWO support strategy is focused on advanced computing technologies. Many of these technologies are critical for enhancing the performance, scalability, and modeling capabilities of strategic CWO applications and for analyzing the resulting data. The MOS team supports the development and direct application of these technologies to CWO research problems and the integration of these tools into CWO codes and research infrastructure. Some particular activities requested by DoD users include: (1) porting codes to the newest architectures and optimizing them for maximum performance and scalability, (2) evaluating various grid/mesh techniques (unstructured grids for coastline, different vertical coordinate systems) for accuracy and performance, and (3)

developing and implementing code coupling techniques for ocean codes with atmospheric and/or ice models, but also for integrating other kinds of models (aerosol transport, biology/chemistry, etc.)

CE STRATEGY

DR. SHIRLEY MOORE, UNIVERSITY OF TENNESSEE,
KNOXVILLE, CE FA POC

The CE team is led by the Innovative Computing Laboratory (ICL) at the University of Tennessee-Knoxville (UTK). ICL, under the leadership of Distinguished Professor Jack Dongarra, is a world-recognized leader in many areas of high performance computing (HPC). The CE FA POC is Dr. Shirley Moore, who is the Associate Director of Research for ICL. Dr. Moore is assisted by Dr. David Cronk, a research scientist with ICL, and by Dr. Thomas Cortese, the onsite CE lead at NAVO MSRC. Other highly qualified personnel at UTK and other PET team institutions are involved in CE projects.

The CE vision is to improve user productivity and application performance through the provision of effective parallel programming tools as part of a consistent well-documented CE. Achieving this goal involves a number of activities. CE staff are collaborating with Shared Resource Center (SRC) systems staff to implement a CE across the SRCs that is as consistent as possible. The CE includes compilers, message-passing libraries, debugging and performance analysis tools, and data management and visualization tools. CE staff are working to ensure that all components of the environment are operating correctly in the SRC batch-queuing environments and that they are adequately supported and documented. Cross-platform tools are being made available wherever possible so that users do not need to learn a different tool interface for each platform. These tools include the TotalView debugger, Vampir Message Passing Interface (MPI) performance analysis tool, and KAP/Pro Toolset for OpenMP programming, all of which are commercial tools. Also included is the cross-platform library Performance Application Programming Interface (PAPI) to hardware performance counters, as well as freely available performance analysis tools for sites and users who do not have access to commercial tools. The environment will be updated regularly to keep pace with changes in architectures and operating systems.

CE core support consists of ongoing activities in the areas of assessing and meeting user needs and delivering effective training in HPC architectures and parallel programming tools. Core support activities include organization of a CE user advisory panel to help determine CE strategic goals and user requirements and to help with evaluation of CE tools, documentation, and training materials. Core support personnel are also involved in evaluation and beta-testing of new tools and new versions of existing tools. The

CE training strategy involves development of a comprehensive training curriculum to give DoD users the knowledge and skills they need to effectively use HPC platforms at the SRCs. Training topics include specific programming languages, parallel programming models, performance optimization techniques, and debugging and performance analysis tools. Another CE activity is to anticipate and plan for CE needs in the context of emerging HPC trends and systems, such as cluster computing, meta-computing, and interactive and real-time requirements.

CE projects target specific areas where critical user needs have been identified. Areas that have been identified by the CE user advisory panel include application portability, application performance evaluation and tuning, a consistent well-documented CE, data management, and parallel and asynchronous input/output. Current CE projects are addressing a consistent CE and deployment of PAPI for the purpose of application performance analysis. Project proposals to address these and the other newly identified needs are currently being formulated.

EQM STRATEGY

DR. MARY WHEELER, UNIVERSITY OF TEXAS, AUSTIN,
EQM FA POC

The EQM FA under PET has defined three primary strategic efforts. The first effort involves the development and implementation of new, accurate parallel discretizations with adaptivity that will be based on a posteriori error estimators. In addition, this effort will include robust scalable parallel solvers.

A second effort involves code couplings of different physical processes (e.g., flow, transport, reactions, and mechanics) which occur within the same physical domain, or when different physical regimes (e.g., surface/subsurface, and fluid/structure) interact through interfaces. The issues to be considered include (1) appropriate physical quantities or transmission conditions that are dependent on the application for coupling across physical domains, (2) development of algorithms for modeling transmission conditions, (3) time-stepping selection for treating couplings of different physical models, (4) grid projection strategies, and (5) computer science tools for implementing coupled processes.

The third effort involves the integration of emerging computational tools for handling large data sets, grid/metacomputing, interactive steering, scientific visualization, mesh partitioning, and dynamic adaptive grids.

It is anticipated that the algorithms and tools developed under these efforts will be applied to EQM/CWO software such as Adaptive Hydraulics (ADH), CE-QUAL-ICM, ADCIRC, and CH3D.

THE PORTHOLE

Left:
NAVOCEANO
Change of
Command
Ceremony 8
March 2002.
Participants L-R,
CAPT Philip G.
Renaud (incoming),
RDML Thomas O.
Donaldson, V,
Commander, NMOC,
and CAPT Timothy
McGee

Right:
Steve Adamec,
Director, NAVO
MSRC, briefs visit-
ing members of
the U.S.
Commission on
Ocean Policy.

Left:
RDML Thomas O.
Donaldson, V, and
members of
the U.S. Commission
on Ocean Policy.



Right:
(L-R) CDR Raul Martinez Sanchez,
Director of Hydrographic
Department, Hydrographic and
Cartographic Office (DIGADHICAR),
LCDR Rafael Ponce, Director of
Tides and Maritime Publications,
DIGADHICAR, and RADM Anastacio F.
DeAbiega Gamez, Director General,
DIGADHICAR, Eric Villalobos,
NAVOCEANO International Division.

Left:
(L-R) CDR Augusto
Mourao Ezequiel,
Technical Director,
Portuguese Hydrographic
Institute, VADM Jose
Torres Sobral, General
Director, Portuguese
Hydrographic Institute and
Dave Cole, NAVO MSRC,
Carol Nichols, CNMOC.



Left:
(L-R) LCDR Erhan Gezgin,
Project Officer, Turkish
Navy Department of
Navigation, Hydrography
and Oceanography
(TNDHNO), CPT Ali
Kaplan, Head of
Technology Group,
TNDHNO, CDR Ahmet
Turker, Chief of
Oceanography Division,
Michael Jeffries,
NAVOCEANO
International
Division.



Right:
(L-R) Dave Cole, NAVO
MSRC, CPT Eric Fraser,
British Naval Staff,
British Embassy, and
LCDR Don Ventura.

Right:
Members of the Basic
Officer Assesion
Training course at the
Naval Meteorology and
Oceanography Professional
Development Center.





Left:
(L-R, facing) LTC Jeff Pritchard (USAF), LCDR Dave Manero (USN), MAJ Jim Hecker (USAF) of Senator Trent Lott's staff.



Left:
Members of the editorial and reporting staff of The Sun Herald (Biloxi, MS) attending a presentation by Pete Gruzinskas of the NAVO MSRC Visualization Center.

Right:
(L-R) Frank Lovato, NAVO MSRC Security Officer, Doug Lee, Director, Louisiana Research and Technology Park Baton Rouge, and companions.



Left:
Visit of Steve Hall, National Imagery and Mapping Agency Gulf of Mexico director.

Right:
(L-R) Dave Cole, LCDR Henry Howell, United Kingdom (UK) Personal Exchange Program (PEP) Officer, FNMOC, and Tom Crew, Allied Cooperation/Technology Transfer Division, NAVOCEANO.



Right:
(L-R) ADM V.E. Clark, Chief of Naval Operations and Steve Adamec, Director NAVO MSRC.



Left:
Members of the PET MOS Team are familiarized with the NAVO MSRC by Dave Cole.



File Transfer Techniques: Transferring Between HABU and Other Systems

Sheila Carbonette, NAVO MSRC User Support

This article describes a method for transferring files between HABU and other systems via the transfer queue using a stream of job steps in a single job command file. The transfer queue was created as a shared usage queue for users to transfer input and/or output files before/after the executable program as serial jobs.

This means the transfer queue shares the interactive nodes with other users. In order to submit a job to this queue, you must specify the following in the job step:

```
#@ class = transfer
#@ node_usage = shared
```

To specify a stream of job steps, you need to list each job step in the job command file. Each job step must be followed by the "queue" statement. The "dependency" keyword is used to let LoadLeveler know that the job steps are not independent and that the next job will run based on the exit status of the previously run job. In the example below, there are three job steps:

1. A serial job step to transfer the input file via the transfer queue.
2. A parallel job step to run a sample POE script via the batch queue.
3. A serial job step to transfer the output file via the transfer queue.

Note: Each job step must have a unique name for all output and error files.

Example job command file:

```
f18n13e% more multistepjobscript
#!/bin/ksh
# ***** STEP 1 - SERIAL *****
#@ account_no=NA0101
#@ class = transfer
#@ step_name = step1
#@ executable = get_file.ksh
#@ input = /dev/null
#@ output = $(jobid).step$(stepid).log
#@ error = $(jobid).step$(stepid).log
#@ environment = COPY_ALL
#@ initialdir = /scr/shecar/multisteps
#@ job_name = multistep
#@ job_type = serial
#@ node_usage = shared
#@ notification = never
#@ notify_user = shecar
#@ shell = /bin/ksh
#@ wall_clock_limit = 300
#@ queue
# ***** STEP 2 - PARALLEL *****
#@ account_no=NA0101
#@ class = batch
```

```
#@ dependency = (step1 == 0)
#@ step_name = step2
#@ executable = poe_script.ksh
#@ input = /dev/null
#@ output = $(jobid).step$(stepid).log
#@ error = $(jobid).step$(stepid).log
#@ environment = COPY_ALL; MP_LABELIO=yes
#@ initialdir = /scr/shecar/multisteps
#@ job_name = multistep
#@ job_type = parallel
#@ node = 4
#@ tasks_per_node = 4
#@ node_usage = not_shared
#@ notification = never
#@ notify_user = shecar
#@ shell = /bin/ksh
#@ wall_clock_limit = 600
#@ queue
# ***** JOB STEP 3 - SERIAL *****
#@ account_no=NA0101
#@ class = transfer
#@ node_usage = shared
#@ dependency = (step2 == 0)
#@ step_name = step3
#@ executable = put_file.ksh
#@ input = /dev/null
#@ output = $(jobid).step$(stepid).log
#@ error = $(jobid).step$(stepid).log
#@ environment = COPY_ALL
#@ initialdir = /scr/shecar/multisteps
#@ job_name = multistep
#@ error = $(jobid).step$(stepid).log
#@ environment = COPY_ALL
#@ initialdir = /scr/shecar/multisteps
#@ job_name = multistep
#@ job_type = serial
#@ notification = always
#@ notify_user = shecar
#@ shell = /bin/ksh
#@ wall_clock_limit = 300
#@ queue
# ***** END OF multistepjobscript SCRIPT *****
```

JOB STEP 1 EXECUTABLE

```
f18n13e% more get_file.ksh
#!/bin/ksh
let RC=0
# *** IDENTIFY NODE RUNNING SERIAL JOB STEP ***
test_script.ksh
# *** REMOVE FILE FROM PREVIOUS RUN ***
rm test_file
let RC=$?
if (($RC != 0))
then /bin/echo "$0-ERROR: Removing test_file; RC=$RC" exit
```



```

$RC
fi
# *** GET FILE FROM JULES ***
/usr/bin/rcp jules-hip0:/u/a/shecar/test/test_file test_file
let RC=$?
if (($RC != 0))
then
  /bin/echo "$0-ERROR: RCPing test_file FROM Jules; RC=$RC"
else
  /bin/echo "File Staging FROM Jules Completed SUCCESSFULLY!"
fi
exit $RC
# *** END of get_file.ksh script ***

```

JOB STEP 2 EXECUTABLE

```

f18n13e% more poe_script.ksh
#!/bin/ksh
# *** Run POE job ***
poe test_script.ksh
# *** END of poe_script.ksh script ***

```

Which runs the following:

```

f18n13e% more put_file.ksh
f18n13e% more test_script.ksh
#!/bin/ksh
let RC=0
/bin/echo "Hostname = `hostname -s`"
let RC=$?
if (($RC != 0))
then
  exit $RC
fi
/bin/echo "Date = `date`"
let RC=$?
exit $RC
# *** END of test_script.ksh ***

```

JOB STEP 3 EXECUTABLE

```

f18n13e% more put_file.ksh
#!/bin/ksh
let RC=0
#
# *** IDENTIFY NODE RUNNING SERIAL JOB STEP ***
#
test_script.ksh
#
# *** COPY FILE TO JULES ***
#
/usr/bin/rcp test_file jules-hip0:/u/a/shecar/test/test_file
let RC=$?
if (($RC != 0))
then\
  /bin/echo "$0-ERROR: RCPing test_file TO Jules; RC=$RC"
else
  /bin/echo "File Staging TO Jules Completed SUCCESSFULLY!"
fi
exit $RC
# *** END of put_file.ksh script ***

```

Use the "llsubmit" command to submit the multistepjob:

```

f18n13e% llsubmit multistepjobscript
llsubmit: Processed command file through Submit Filter:
"/loadldir/local/bin/ll_submit_filter".
llsubmit: The job "f18n13s.navo.hpc.mil.41673" with 3 job steps has
been submitted.

```

When this job command file is submitted to LoadLeveler, three jobs will be placed in the queue. This can be verified by using the "llq" command.

```

f18n13e% llq | egrep "^|shecar"

```

ID	Owner	Submitted	ST	PR	Class
f18n13s.416	shecar	3/21 09:11	I	50	transfer
73.0					
f18n13s.416	shecar	3/21 09:11	NQ	50	transfer
73.2					
f18n13s.416	shecar	3/21 09:11	NQ	50	batch
73.1					

Initially the first job is queued, and the second two will wait until the previous one finishes.

JOB STEP 1 OUTPUT

```

f18n13e% more 41673.step0.log
Hostname = f18n13e
Date = Thu Mar 21 09:37:16 CST 2002
File Staging FROM Jules Completed SUCCESSFULLY!

```

JOB STEP 2 OUTPUT

```

f18n13e% more 41673.step1.log
ATTENTION: 0031-408 16 tasks allocated by LoadLeveler, continuing
1:Hostname = f21n03e
7:Hostname = f21n04e
11:Hostname = f21n05e
12:Hostname = f21n06e
0:Hostname = f21n03e
5:Hostname = f21n04e
8:Hostname = f21n05e
14:Hostname = f21n06e
2:Hostname = f21n03e
6:Hostname = f21n04e
9:Hostname = f21n05e
13:Hostname = f21n06e
3:Hostname = f21n03e
4:Hostname = f21n04e
10:Hostname = f21n05e
15:Hostname = f21n06e
2:Date = Thu Mar 21 09:51:24 CST 2002
7:Date = Thu Mar 21 09:51:24 CST 2002
8:Date = Thu Mar 21 09:51:24 CST 2002
14:Date = Thu Mar 21 09:51:24 CST 2002
11:Date = Thu Mar 21 09:51:24 CST 2002
15:Date = Thu Mar 21 09:51:24 CST 2002
10:Date = Thu Mar 21 09:51:24 CST 2002
3:Date = Thu Mar 21 09:51:24 CST 2002
5:Date = Thu Mar 21 09:51:24 CST 2002
12:Date = Thu Mar 21 09:51:24 CST 2002
0:Date = Thu Mar 21 09:51:24 CST 2002
6:Date = Thu Mar 21 09:51:24 CST 2002
1:Date = Thu Mar 21 09:51:24 CST 2002
13:Date = Thu Mar 21 09:51:24 CST 2002
4:Date = Thu Mar 21 09:51:24 CST 2002
9:Date = Thu Mar 21 09:51:24 CST 2002

```

JOB STEP 3 OUTPUT

```

f18n13e% more 41673.step2.log
Hostname = f15n13e
Date = Thu Mar 21 09:53:22 CST 2002
File Staging TO Jules Completed SUCCESSFULLY!

```

VSIPL...continued

performance computers. Efficient machine-specific implementations of BLAS are available for several high-performance computers.

APPROACH AND RESULT

The VSIPL TASP library was modified to call all three libraries (i.e., FFTW, LAPACK, and BLAS). Another feature is a header file that allows the user to switch between building the ERI library and the standard TASP library. The final delivered library will include a perl script so that the end-user can download each of the reference libraries and run the script to add the necessary modifications. The user will then build each library, and MSTI will provide some example codes and benchmarking routines.

The initial results for VSIPL/ERI are encouraging: Figure 1 shows some comparative benchmarks for four FFT functions using float and double data types. The top table is the VSIPL/ERI numbers, and the lower table is the TASP numbers. The four functions are real-to-complex FFT

out-of-place (rcfft0p), complex-to-real FFT out-of-place (crfft0p), complex-to-complex FFT out-of-place (ccfft0p), and complex-to-complex FFT in-place (ccfft1p). The numbers were generated using the Linux VSIPL/ERI implementation on an 800-MHz PIII with 256M RAM (mentioned since processor speed and RAM size affect the performance). Figure 2 shows some comparative Linear Algebra benchmarks. The six functions are as follows: Cholesky decomposition (symmetric (Hermitian) positive definite linear system) for real (Cholesky_f) and complex (Ccholesky_f); LU decomposition (Gaussian decomposition) for real (LUD_f) and complex (CLUD_f); general matrix product (compute the general product of two matrices and accumulate) Gemp_f; and matrix product (product of two matrices) Mprod_f. The upper table represents the VSIPL/ERI numbers, and the lower table shows the TASP VSIPL numbers. Most of the functions require two input values, and for these cases, the same number (column heading number) was used

for both inputs. These benchmarks were generated on the same 800-MHz PIII with 256M RAM.

SUMMARY AND CONCLUSION

It is important to note that the LAPACK routines are called only when the matrix majors are column and the stride is equal to one, otherwise the TASP VSIPL code is executed. It was found that the MPROD and GEMP functions yielded no performance improvement over the TASP VSIPL implementation. The other Linear Algebra functions showed increased performance, especially for the larger matrix sizes. Each of the VSIPL/ERI FFT functions executed in half the time of the TASP VSIPL FFTs for the sizes shown. In all, 56 linear algebra functions in VSIPL and LAPACK/BLAS were wrapped in VSIPL/ERI. VSIPL/ERI will be free to all U.S. sites (Government-purposes rights) and will also be open source for the U.S. Government. MSTI will offer a free license for academic institutions. A commercial implementation will also be available.

The Updated NAVO MSRC User Support Team



UPCOMING EVENTS

JUN

SIGGRAPH 2002

July 21 | 26, 2002

San Antonio, TX
www.siggraph.org/s2002

JUL

2002 USENIX Annual
Technical Conference

June 9 | 14, 2002

Monterey, CA
www.usenix.org

AUG

2002
Western Pacific
Geophysics Meeting

July 9 | 12, 2002

Wellington, New Zealand
www.agu.org

SEP

IBM SP SCICOMP 6

August 19 | 23, 2002

San Francisco, CA
www.spcicomp.org

OCT

Global
Grid Forum 5
July 21 | 25, 2002
Edinburgh, Scotland
www.gridforum.org

USENIX
Security
Symposium
August 5 | 9, 2002
San Francisco, CA
www.usenix.org

NOV

PACT 2002
September 22 | 25, 2002
Charlottesville, VA
www.pactconf.org

DEC

Oceans 2002
October 29 | 31, 2002
Biloxi, MS
www.oceans2002.com

APS DCIOMP
CCP 2002
August 25 | 28, 2002
San Diego, CA
www.aps.org/DCOMP/Meetings/meeting/meet02/CPCP2002.htm

JAN


FEB

MAR

APR

Global Grid Forum 6
October 15 | 17, 2002
Chicago, IL
www.gridforum.org

MAY



Naval Oceanographic Office * MAJOR SHARED RESOURCE CENTER

1002 Balch Boulevard . Stennis Space Center, Mississippi . 39522