Navigator

NAVO MSRC SPRING 2007

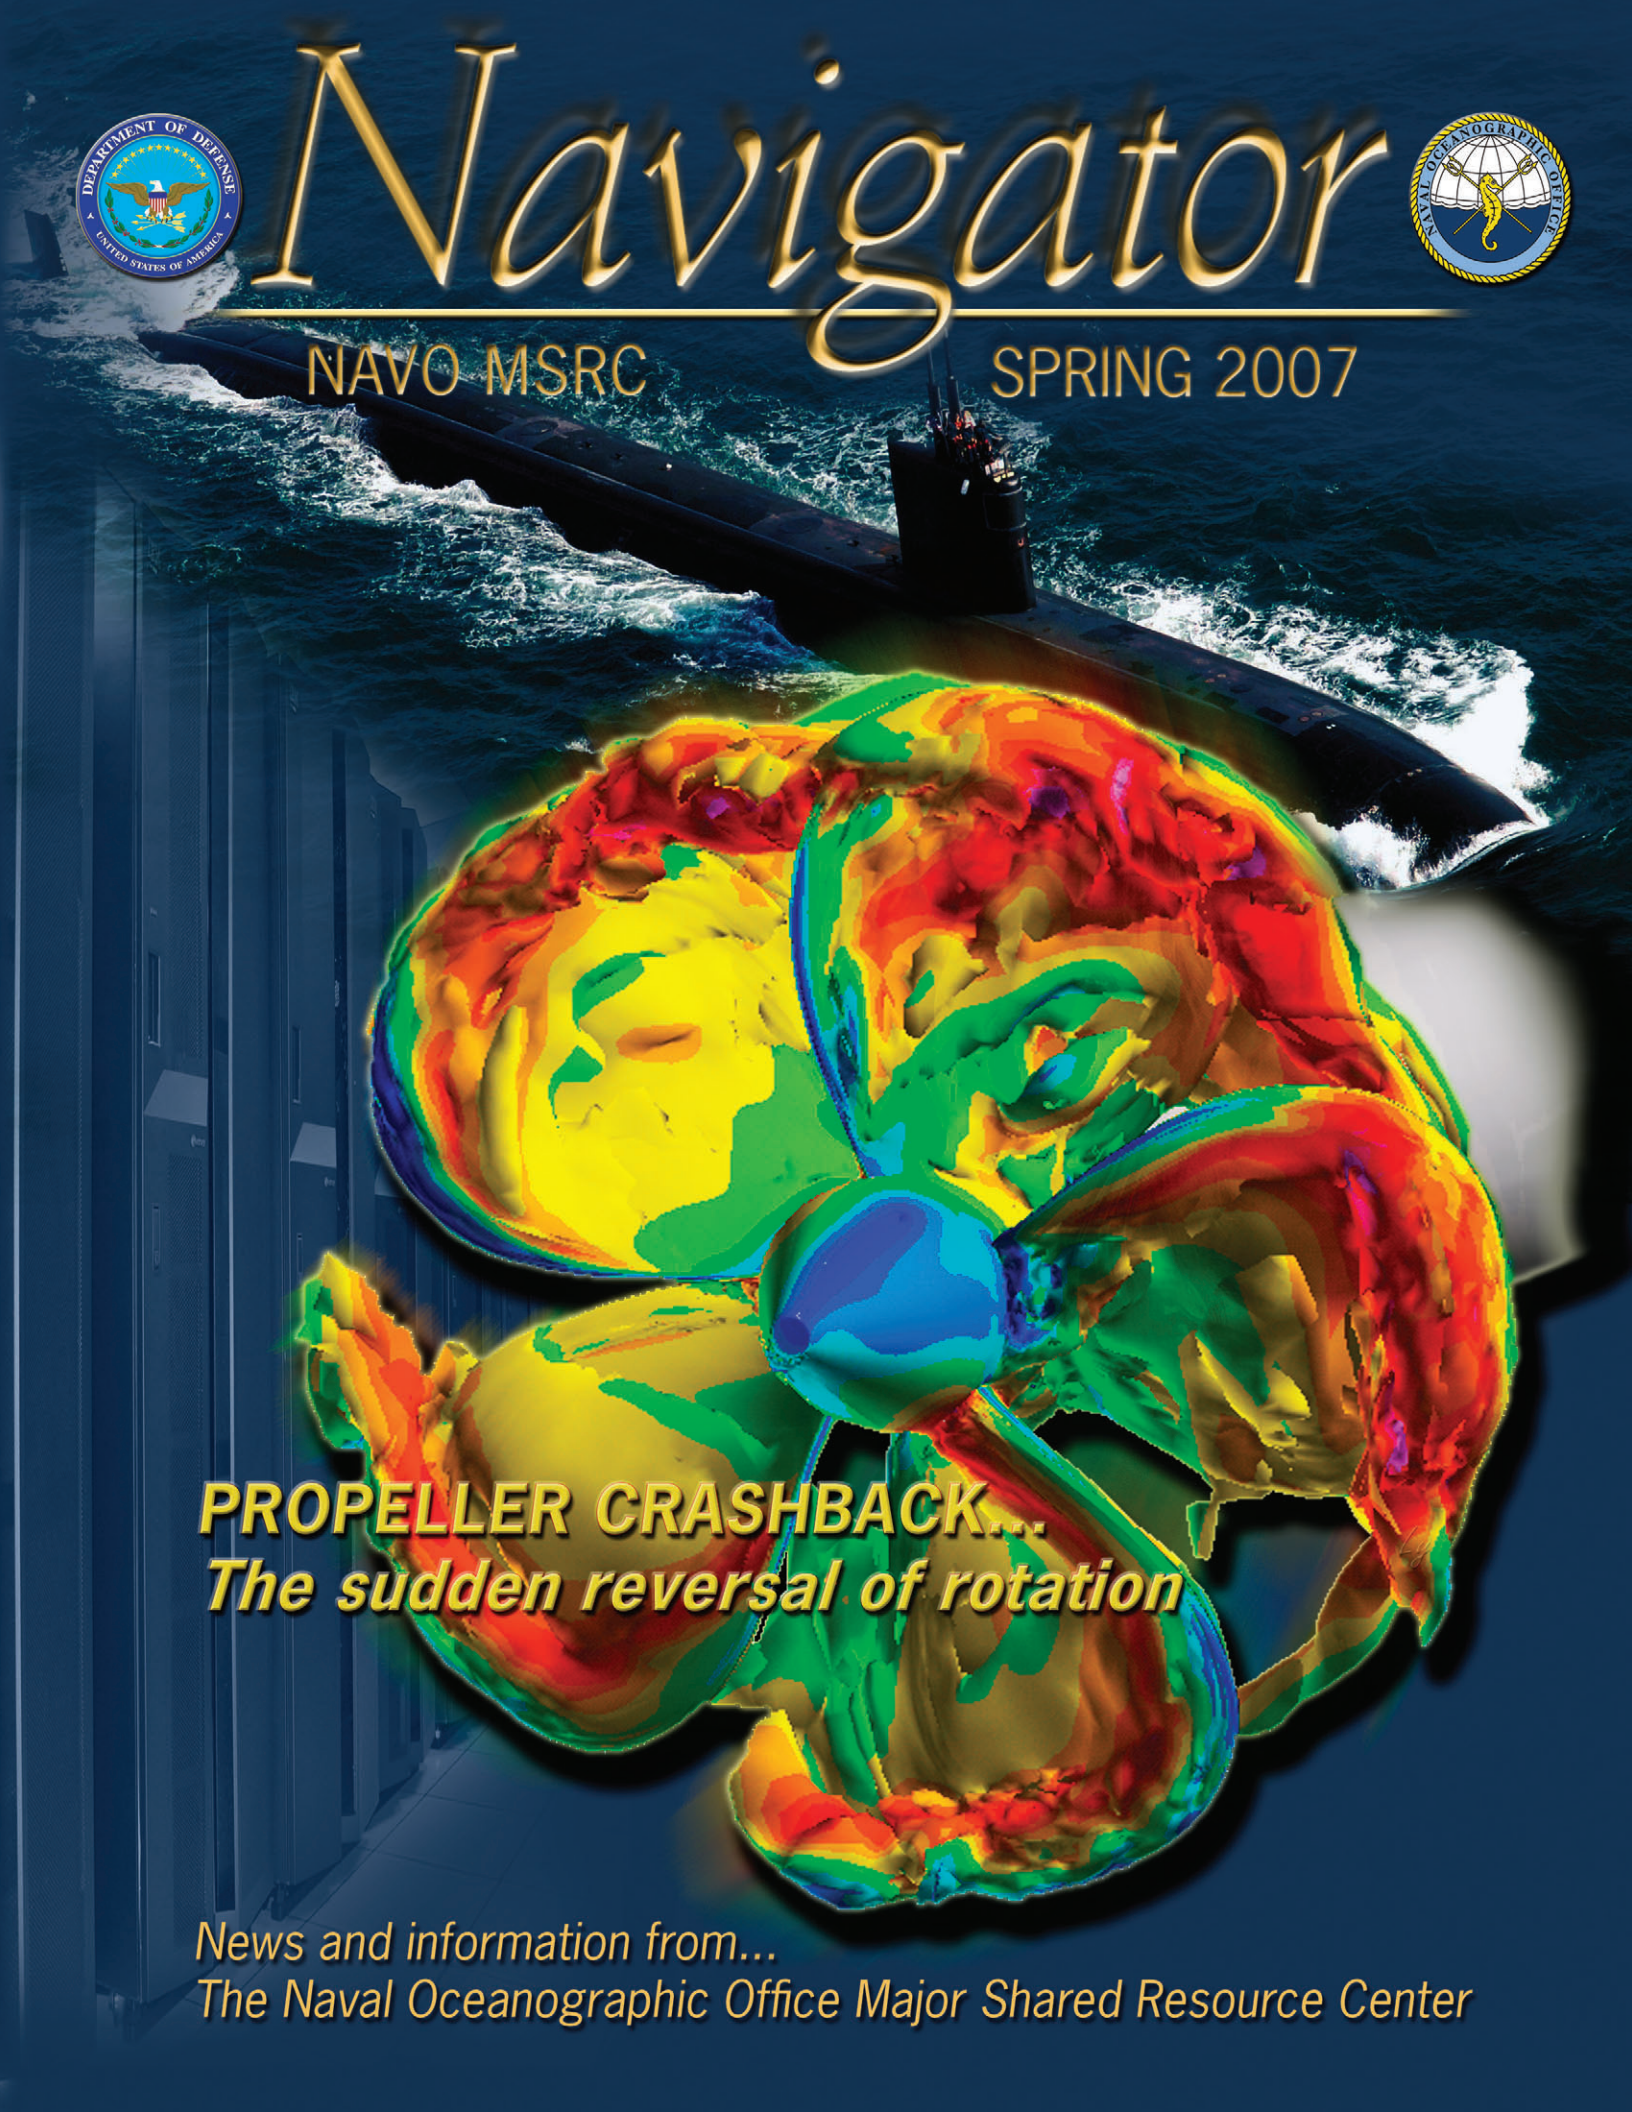

PROPELLER CRASHBACK...
The sudden reversal of rotation

News and information from...
The Naval Oceanographic Office Major Shared Resource Center

Changing, Growing, and Expanding at the NAVO MSRC

In this, my first Navigator article as the new Director, I am pleased to report that the NAVO MSRC has significantly increased both the computational and storage capability of the center. My predecessor, Steve Adamec, retired from federal service last October. Steve's federal career includes a long list of significant accomplishments, and we wish him well in future endeavors.

The most significant enhancement was the addition of two large IBM High Performance Computing (HPC) systems, both of which are based upon IBM's POWER 5+ processor technology. BABBAGE, a 3000-processor unclassified system, and PASCAL, a 2000-processor classified system, were delivered in October and completed system acceptance in December.

These systems are configured with two Gigabytes (GB) of memory per processor, IBM's Federation interprocessor switch fabric, and IBM's Global Parallel File System (GPFS), all of which will facilitate the execution of tremendously large applications. These new additions increased the center's computational capacity from 32 teraflops (10**12 operations per second) to 69 teraflops, a 115-percent increase.

The NAVO MSRC data storage growth rate has doubled within the past twelve months to approximately 12-15 Terabytes (TB) per week. As a result, we have initiated and executed a plan to introduce StorageTek (STK) T10000 tape drive/cartridge technology into the NAVO MSRC storage environment.

The STK PowderHorn silos that currently house NAVOCEANO MSRC first copy data have been upgraded to add a total of 14 T10000 drives. The T10000 cartridge capacity (500 GB) represents a 150-percent increase over the older 200-GB cartridges.

The High Performance Computing Modernization Program (HPCMP) Remote Storage Facility (RSF) site, operated and managed by the NAVO MSRC, has experienced significant data storage growth as well. In January the RSF growth was approximately 30 TB per week from the NAVO MSRC, Engineering Research and Development Center MSRC, Aeronautical Systems Center (ASC) MSRC, U.S. Army Research Laboratory (ARL) MSRC, Arctic Region Supercomputing Center

Continued Page 22

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 Electromagnetic 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 Furness Petro, petrogio@navo.hpc.mil

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

Contents

Two New IBM POWER5+ Systems Replace the Venerable IBM POWER4 System

As the adage goes, "out with the old, in with the new!" The Naval Oceanographic Office (NAVO) Major Shared Resource Center (MSRC) recently completed installation and acceptance of two IBM POWER5+ High Performance Computing (HPC) systems as a culmination of the Technology Insertion for Fiscal Year 2006 (TI-06). The NAVO MSRC also began activities to retire the venerable IBM POWER4 system MARCELLUS. With the addition of approximately 5000 POWER5+ processors with a peak computational capacity of 37 teraflops, the TI-06 enhancements will increase the NAVO MSRC aggregate in excess of 61 teraflops.

The larger of the two systems (BABBAGE) has a total of 3072 processors (192 nodes) of which 2912 (182 nodes) are available for unclassified computational use. BABBAGE is configured with sixteen 1.9 Gigahertz (GHz) POWER5+ processors per node that are interconnected by an IBM High Performance Switch (HPS). Two of the computational nodes have 128 Gigabytes (GB) of memory to support large memory applications while the remaining nodes have 32 GB. Approximately 160 Terabytes (TB) of usable disk space are available with the Global Parallel File System (GPFS).

PASCAL, which replaced the 1408 processor IBM POWER 4 MARCELLUS in the NAVO MSRC classified HPC environment, has a total of 1920 processors (120 nodes) of which 1792 (112 nodes) are available for computational use. Like BABBAGE, two of the computational nodes have 128 GB of memory to support large memory applications while the remaining nodes have 32 GB. Approximately 120 TB of usable disk space is available via the GPFS file system.

Both systems sailed smoothly through systems acceptance testing with a very low number of operational interrupts, demonstrating once again the resilience and reliability that are hallmarks of the IBM AIX HPC systems. This notable system stability enabled early system access to Capability Applications Project (CAP) users and enabled the NAVO MSRC to provide over 1,266,425 computational hours in support of CAP Phase I and 2,305,016 for CAP Phase II.

TI-06 IBM Power5+ Systems on the NAVO MSRC Computer Floor.

Crashback Maneuvers

Scott A. Slimon and Craig A. Wagner, Electric Boat Corporation

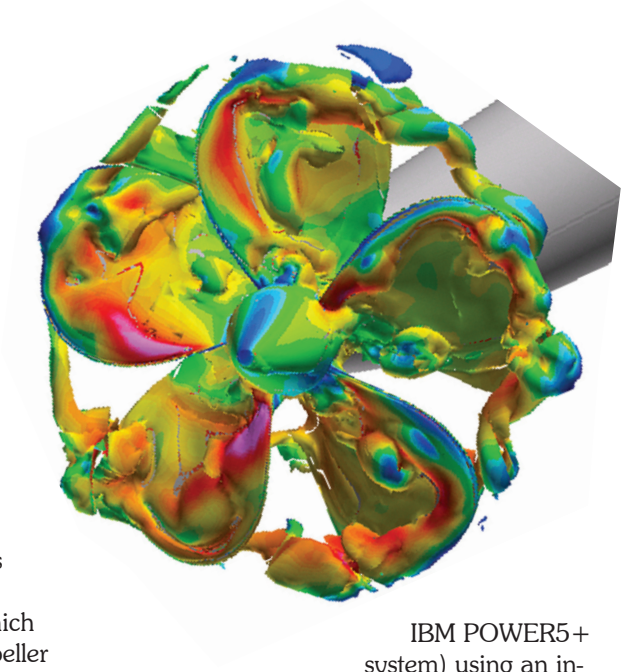

Propeller crashback is the sudden reversal of rotation by marine propellers, usually performed under emergency conditions. Large-scale unsteady flow structures are created when the reversing propeller generates a thrust force opposing the advance motion of the vessel.

These flow structures interact with the propeller and generate large side loads that result in high peak propeller blade stresses and have a considerable impact on the maneuvering performance of the marine vehicle. The study discussed in this article was performed to further the understanding of the unsteady flow structures and associated unsteady propeller forces generated during crashback maneuvers.

Application of computational methods in predicting these forces has been limited due in part to the substantial computational resources required for accurate crashback flow simulations as a result of the disparity in time scales that characterize these flows. The crashback simulations performed in this study are for Propeller 4381, which is a five-bladed open propeller recently investigated in high-fidelity experiments at the Naval Surface Warfare Center-Carderock Division (NSWC-CD).[1]

Computational Approach

The results reported in this Capability Applications Project (CAP) project were generated on BABBAGE (the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) 3072-processor unclassified IBM POWER5+ system) using an in-house Reynolds-averaged Navier-Stokes (RANS) solver. The solver is a finite difference code employing a pentadiagonal linearized block implicit solution algorithm. Spatial derivatives are approximated using central differencing with implicit and explicit fourth order numerical dissipation models. The flow field is discretized using structured Chimera and hybrid multi-block grids.

Pseudo-compressibility is used for steady flow simulations, and a dual time step approach is employed for unsteady simulations. A zonal variant of the Detached Eddy Simulation[2] (DES) turbulence model was used for the simulations. In this model,[3] regions in which the Large Eddy Simulation (LES) mode of the model is activated are identified based on local turbulence and flow quantities. This approach eliminates a failure mode of DES by ensuring that the LES-mode of the model is only applied to detached shear layers.

CAP Parallel Performance Tests

Parallel operation of the solver is achieved using unblocked Message Passing Interface (MPI) library calls

Continued Next Page...

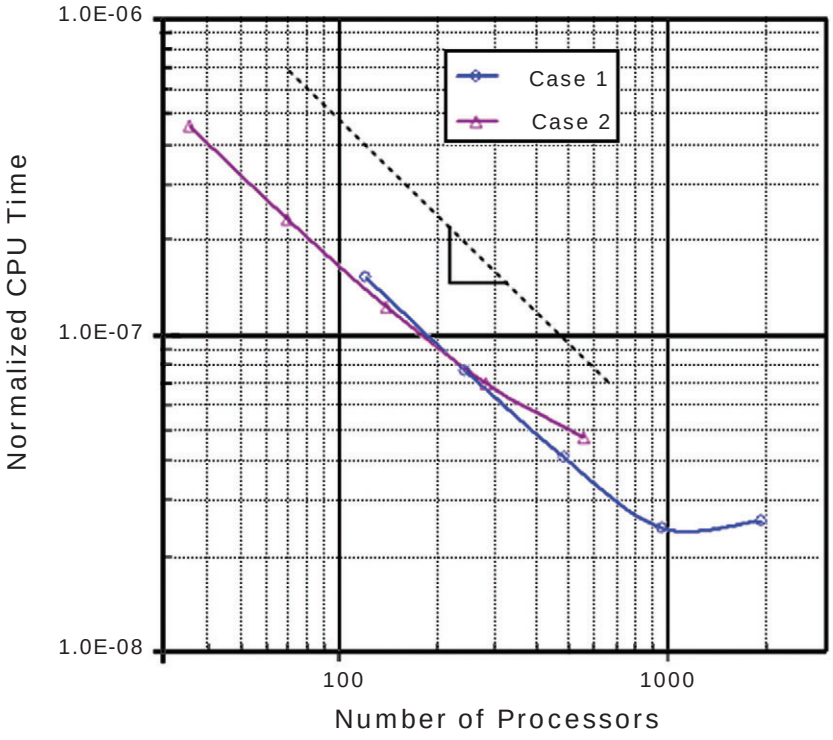

Figure 1. Scalability Results on BABBAGE (IBM P5+).

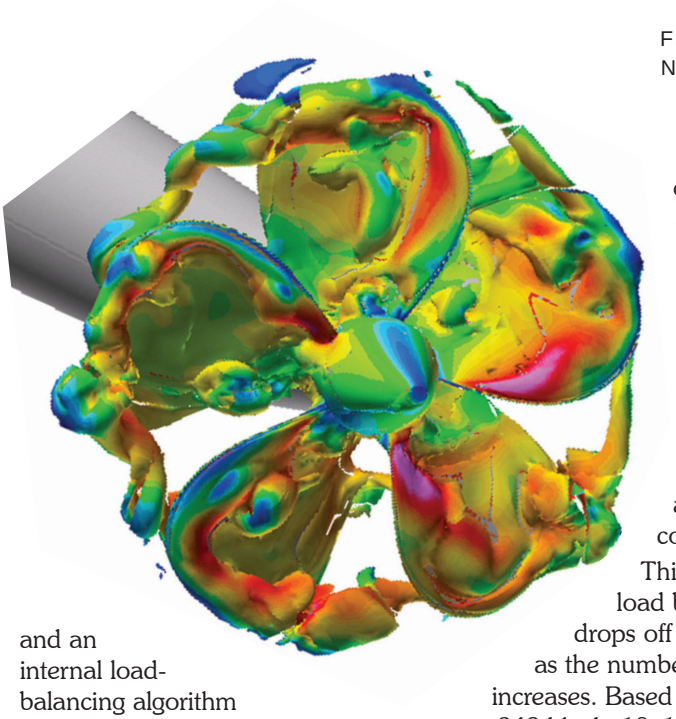

Figure 2. Isocontour of Vorticity Magnitude Colored by Normalized Velocity Magnitude.

and an internal load-balancing algorithm that distributes the grid blocks among the processors. Each processor sequentially works through the list of grid blocks assigned to it. Once the solution for a grid block on a processor has been advanced to the next time level, inter-block data required by other grid blocks on other processors are sent with MPI isend and irecv calls while the processor performs work on the next grid block.

Figure 1 shows Phase I scalability testing results performed for two different cases. Case 1 used a 30x106 node mesh with 1,936 equally sized grid blocks, and Case 2 used a 4.8x106 node mesh with 560 different sized grid blocks.

The results for Case 1 demonstrate good scalability until the last data point for which there is only one grid block per processor (resulting in serial communication and floating point operations).

For Case 2 the scalability is reasonable, but there is a noticeable degradation in parallel efficiency as the processor count increases. This is due to the load balance, which drops off monotonically as the number of processors increases. Based on these results, a 949 block, 10x106 node grid for Propeller 4381 was generated such that a load balance of 98 percent was achieved with 433 processors. For this processor count, two to three grid blocks are assigned to each processor.

Crashback Simulation

A crashback simulation was performed for Propeller 4381 at an advance coefficient (J) of -0.5. The predicted unsteady character of the flow is depicted by the non-uniformity of the flow around the propeller in Figures 2 and 3. Figure 2 highlights the separated shear layers off of each blade surface, while Figure 3 shows a complex ring vortex generated by the flow recirculating around the propeller.

The ring vortex is complex and does not remain centered on the propeller axis – it meanders both radially and axially, causing large chaotic-like variations in thrust and side loads. The crashback simulation was performed for over 1x105 time steps using 40 subiterations per time step.

This provided close to 7½ seconds of simulation time, allowing resolution of frequencies down to approximately ¼ Hertz (Hz). The predicted side force is compared with the Reference 1 experimental data in Figures 4 and 5. The average rms and spectral shape of the force history are in good agreement with the experiment. Most significantly, there is good agreement in the predicted spectral peak relative to the experimental measurements.

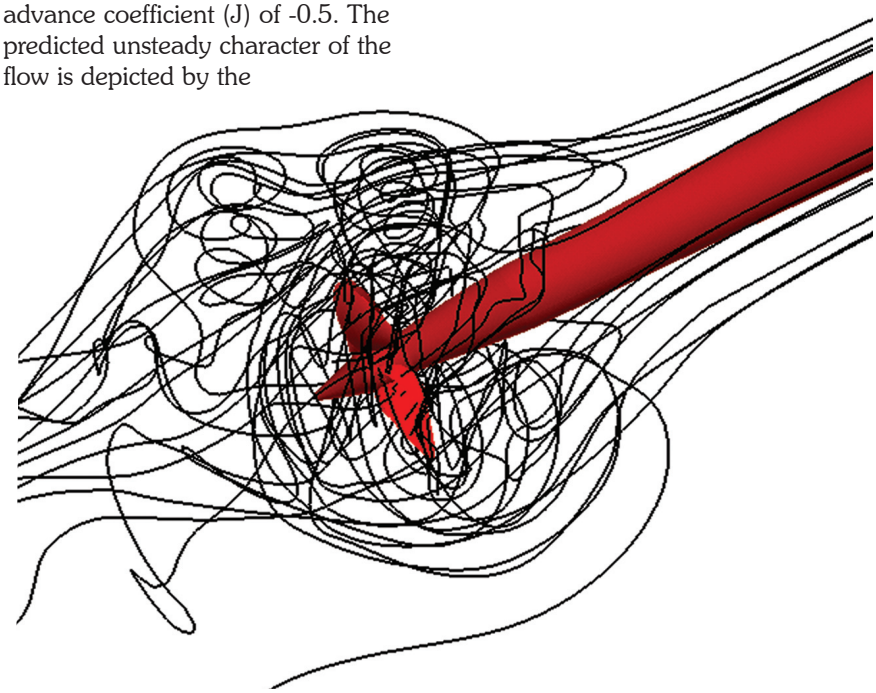

Figure 3. Streamlines Illustrating Radial Extent and Circumferential Variation of Ring Vortex.

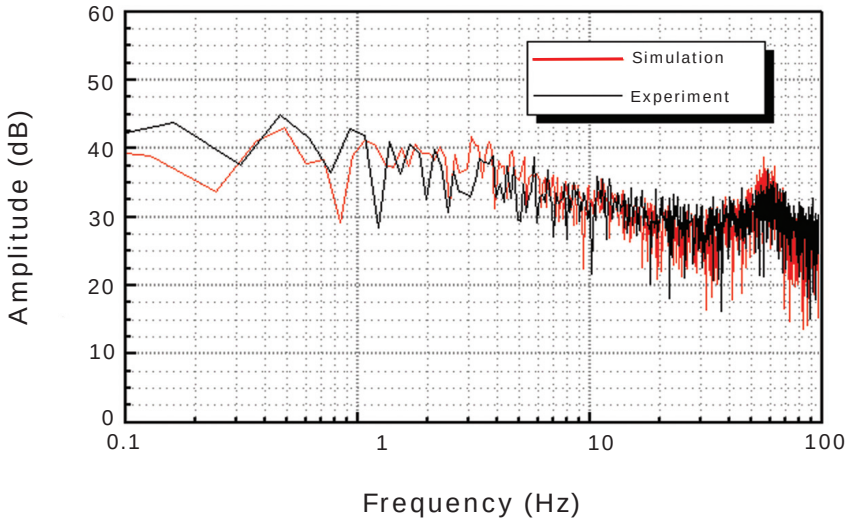

Figure 4. Radial Force Time History, Propeller 4381, J=-0.5.

Summary

Use of BABBAGE under the CAP program has been extremely valuable to this research. It allowed the researchers to enhance communication and load sharing protocols that appreciably improved the scalability of a production-based solver.

The simulations results from Phase II are significant in that the researchers were able to accurately resolve the ultra-low frequency propeller side force generated during crashbacks. This is a compelling result since this side force is not well understood and has significant implications in propeller design and marine vehicle maneuvering.

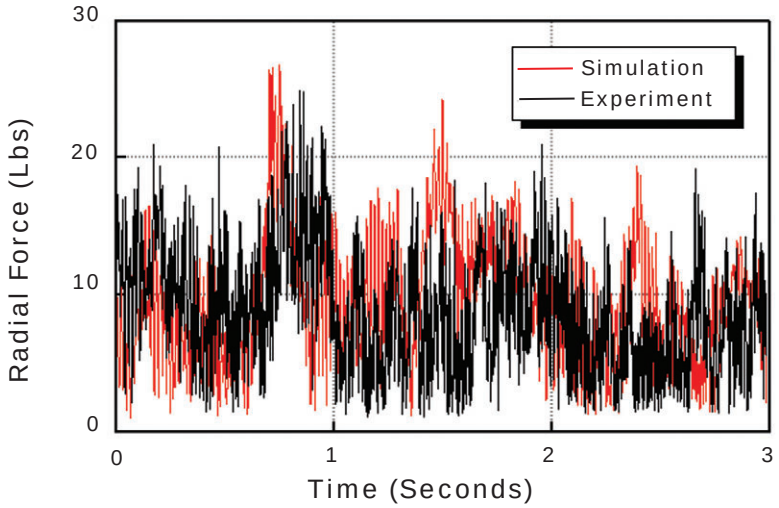

Figure 5. Radial Force Spectra, Propeller 4381, J=-0.5.

Acknowledgments

The authors would like to thank Troy Hollingsworth and Garth Dolphin for generating the computational grids used in this study. The support of John Skinner and the support group at NAVO MSRC was greatly appreciated during the CAP. The authors also acknowledge helpful discussions regarding the experimental data with Dr. Stuart Jessup and Martin Donnelly at NSWC-CD. This work was funded using an internal research and development project at Electric Boat Corporation.

References

1. Jessup, S. D., J. T. Park, C. J. Chesnakas, M. J. Donnelly, and D. J. Fry, "Performance of Propeller 4381 in Crashback," 2004, Storming Media.
2. Spalart, P. R., W. H. Jou, M. Stretlets, and S. R. Allmaras, "Comments on the Feasibility of LES for Wings and on the Hybrid RANS/LES Approach," 1997, Advances in DNS/LES, Proceedings of the First Air Force Office of Scientific Research (AFOSR) International Conference on Direct Numerical Simulation (DNS)/Large Eddy Simulation (LES).
3. Slimon, S. A., "Computation of Internal Separated Flows Using a Zonal Detached Eddy Simulation Approach," Proceedings of International Mechanical Engineering Congress and Exposition (IMECE), 2003 American Society of Mechanical Engineers (ASME) International Engineering Congress and Exposition, Washington, D.C.

A Practical Approach for Software License Access on Large High Performance Computing Clusters

NAVO MSRC System Administration and Information Assurance Staff

Introduction

Achieving the proper balance between security and usability is one of the classic Information Technology (IT) conundrums. While a near-perfect security can be achieved by eliminating all vestiges of useful features, capabilities, and user accounts, this is not a very meaningful accomplishment.

In contrast, an extremely usable environment can be created, but it is an environment vulnerable to even the simplest hacker exploits. Hence, IT developers almost always strive for equilibrium between security concerns and functionality. This classic conundrum is quite evident when enabling High Performance Computing (HPC) clusters to communicate with resources outside of their local area.

The Challenge

The High Performance Computing Modernization Program (HPCMP) has embarked on a cost savings initiative via software license consolidation. This allows the program to achieve economies of scale on high-use software by aggregating each of the Major Shared Resource Centers (MSRCs) and Distributed Centers (DCs) software demands.

The HPCMP creates a centralized license server that provides "licenses" each time a user accesses one of these consolidated software products. However, to leverage this capability, each computing resource must be able to communicate with a license server that may be external to a particular center.

This poses a challenge to systems with a cluster architecture. Typically, these systems are configured to allow outside access from only a small number of nodes, called front-end nodes, thus protecting the vast majority of the nodes, called compute nodes, from inadvertent or unauthorized access in order to protect user's programs running on these nodes.

Consequently, to realize this software savings, the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) staff needed to develop a secure method that enabled the compute nodes from their large cluster systems to communicate with the consolidated software license server.

Possible Solutions

The NAVO MSRC staff selected a proxy license server approach. (See Figure 1.) With this method, a license request is forwarded, or "proxied," from a local license server to the external license server. Correspondingly, the actual license key is forwarded from the external server to the license requestor via the local license server.

Since there is no direct communication between the compute nodes and the external license server, this is a more secure communication approach. Methods that allow a direct communication path create a potential vector that could be exploited by nefarious individuals. Furthermore, this method allows a compute node to obtain a license key from an external license server.

Product Selection

After examining and evaluating various proxy solutions, the NAVO MSRC staff chose a product simply called "proxy" (http://sourceforge.net/projects/proxy/). There were many reasons that this product stood out above other possible solutions, but three primary reasons are of note.

First of all, "proxy" is open source code. This allowed the NAVO MSRC staff to modify the source code as they see fit, be able to audit the code to make sure it was secure, and to have the option of using the code on multiple architectures if need be. This was especially important in order to proxy FlexLM licenses.

Most of the licenses provided by the HPCMP consolidated license servers are FlexLM-type licenses. In order to proxy these licenses, the proxy code had to be modified to use a 174 byte buffer. Other license software types, such as that used by EnSight, do not

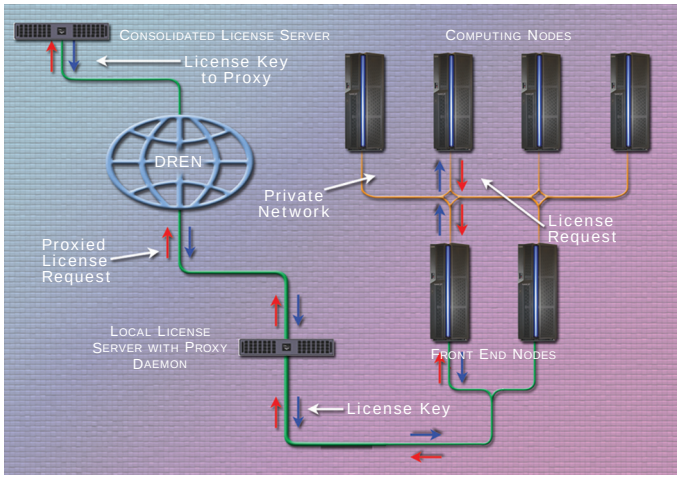

Figure 1. Proxy License Server Architecture at NAVO.

Continued Page 19

Atmospheric Decision Aids and the Need for Detailed Turbulence Simulations

Joe Werne (Senior Research Scientist, NorthWest Research Associates, CoRA Division)

The United States Air Force and the Missile Defense Agency are supporting the development of a flexible Atmospheric Decision Aid (ADA) framework for missile defense applications. These include mechanical turbulence forecasts for the High Altitude Airship (HAA), currently under development by Lockheed Martin, and optical turbulence forecasts for the Airborne Laser (ABL).

The major challenges for this work include (1) the small vertical extent of stratospheric turbulence events, typically O (10-100m) only, and (2) their anisotropic and episodic nature (i.e., they develop from locally unstable flow features in the atmosphere that give rise to vigorous mixing) but they also eventually regain stability as a result of the background vertical stratification.

Unfortunately, current Numerical Weather Prediction (NWP) codes lack the vertical resolution to capture these events, and their SubGrid-Scale (SGS) turbulence parameterizations are perforce applied to spatial scales that are far too large to be consistent with the theoretical assumptions on which the parameterizations are based.

In order to address this situation, a much more general and sophisticated SGS parameterization scheme, based on Bayesian Hierarchical Modeling (BHM), has been developed, which combines real-time NWP output with statistical catalogs developed from previously obtained atmospheric turbulence measurements and the compiled results from high-resolution turbulence simulations. This algorithm is a hybrid deterministic/probabilistic SGS modeling strategy that is self-consistent and which simultaneously predicts both the desired forecast variables and their uncertainties. Furthermore, because of the general nature of the approach and the ubiquity of the underlying atmospheric dynamics (e.g., wind-shear, gravity waves, their instability and interaction, and the ensuing turbulence), it is possible to incorporate the same input information and methods for significant components of both these and other ADAs. Hence, very different applications can significantly leverage identical resources, including: (1) improved NWP output, (2) atmospheric turbulence measurements (primarily balloon and aircraft data), and (3) rescaled and cataloged key high-resolution numerical solutions.

Continued Next Page...

Figure 1. Turbulent wind-shear billows viewed from the side for Ri=0.20 (top), Ri=0.15 (middle), and Ri=0.05 **(bottom) at t≈70. The higher Ri values are associated with flatter billows, which are unstable to turbulence** production, while the round Ri=0.05 billows are stabilized by solid-body rotation, leading to turbulence initiation at the billow edges and in the braid region between billows. The intact braid regions are clearly visible for Ri=0.15 and Ri=0.20, while the Ri=0.05 braids have been wiped out by turbulence. The Ri=0.05 case was computed during CAP 2004 on KRAKEN, while the Ri=0.15 and 0.20 cases were computed during CAP 2006 on BABBAGE.

This methodology is a welcomed development as the simulation and observational results are expensive to obtain, and reuse of the compiled data represents a significant cost savings for both the Department of Defense (DoD) and the American taxpayer.

To support this project, new wind-shear simulations were carried out as part of a 2006 DoD Capability Applications Project (CAP) with higher Reynolds numbers than have been computed previously. The simulations focused on high atmospheric density stratification (i.e., Richardson numbers of Ri=0.15 and Ri=0.20). As many as 3600x1800x1800 spectral modes were needed to obtain these solutions, which consumed 750,000 Central Processing Unit (CPU) hours and generated over 120 Terabytes (TBs) of archival data that will be analyzed in the coming year. In addition, a low-Ri case (with Ri=0.05) was previously computed on the Naval Oceanographic Major Shared Resource Center (NAVO MSRC) IBM P4+ (KRAKEN) as part of CAP 2004.

Figure 1 shows the different flow morphologies that result as Ri is varied for the solutions computed. At advective time t≈70 h/U0 (where 2h is the initial shear-layer depth and 2U0 is the velocity jump across the layer), the Kelvin-Helmholtz billows that spontaneously form in wind shear are transitioning to turbulence.

As the figure shows, between Ri=0.05 and Ri=0.15 the flow experiences a fundamental shift in this transition process. Whereas the Ri=0.05 case has round billows that benefit from the stabilizing influence of their near solid-body rotation, the higher-Ri billows are flattened by stronger background stratification, and their oblong cores become nearly immediately convectively unstable as soon as they overturn.

This shifts the location of turbulence initiation from the billow periphery and braid regions (between billows) at Ri=0.05 to the billow cores when Ri ≥ 0.15. As a result, the braids at high Ri are preserved intact to late times, while those at low Ri are obliterated by turbulence early in the flow evolution. Also, at low Ri a secondary transition to turbulence results when the convectively unstable billow cores eventually also become unstable. Animations of these solutions can be found at www.cora.nwra.com/~werne/cap2.

Recent aircraft wind-shear data clearly demonstrate numerous examples of intact braids when examining the edges of the jet stream,[1] implying that relevant upper tropospheric and stratospheric wind-shear events have Ri values of Ri=0.15 or greater and rarely as low as the more extensively numerically studied Ri=0.05 case. Since the dynamics, evolution, and turbulence statistics all depend critically on Ri (especially above and below the transitional value near Ri≈0.15), the high-Ri solutions computed during CAP-2006 are especially important for the HAA and ABL ADA projects because both of these systems will operate in the upper troposphere and/or the stratosphere near where the aircraft measurements were made.

Detailed comparisons between the aircraft data and the new high-Ri solutions shown in Figure 1 are extremely encouraging, with simulated event durations and extents, billow aspect ratios, braid tilt angles, turbulence spectral slopes, and structure-function anisotropy ratios all in agreement with the atmospheric measurements.

These new and previously computed CAP solutions are currently being used, along with aircraft and balloon measurements, to create a census of atmospheric wind-shear events. Both the census data and the computed turbulence statistics will be used in

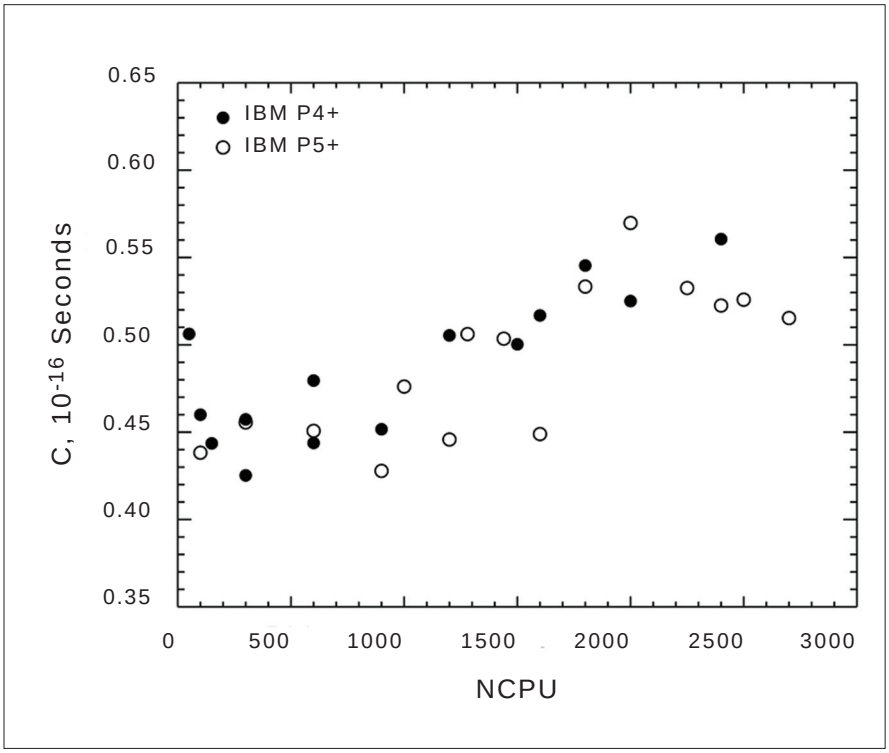

Figure 2. Parallel scaling of TRIPLE on IBM P4+ (KRAKEN) and P5+ (BABBAGE) for fixed per-processor problem size. C is proportional to the wall-clock time per compute cycle. Fluctuations in the data result from the cache-reuse characteristics of the different radix FFT routines used for different problem shapes.

the BHM-based ADAs currently under development.

All of the 2006 CAP work discussed here was carried out using the pseudo-spectral Direct-Numerical-Simulation (DNS) code TRIPLE on the new NAVO MSRC IBM P5+ (BABBAGE).[2] Before being able to compute the new solutions in Figure 1 however, acceptable performance for TRIPLE had to be demonstrated on the new system. Initially there was concern that TRIPLE would not perform as well on BABBAGE as it did on the older NAVO MSRC IBM P4+ (KRAKEN).

There were two primary reasons for this concern. First, TRIPLE makes extensive use of three-dimensional Fast Fourier Transforms (FFTs), meaning it has a very demanding all-to-all communication pattern. Since BABBAGE has twice as many processors per node as KRAKEN, while the two machines have essentially the same internal network,* the authors anticipated the inter-processor communication performance on BABBAGE would be roughly half that on KRAKEN. Second, past experience on IBM platforms indicated that intra-node housekeeping operations can seriously upset the load balancing among processors on a node.

As a result, it was discovered that in order to realize the greatest overall parallel performance, it was best to leave one processor per node idle to handle these intra-node tasks. For example, on KRAKEN, TRIPLE experiences a whopping 30 percent speed increase when only seven of the eight processors on each of KRAKEN's nodes are used. Since BABBAGE has twice as many processors per node, twice the number of intra-node operations is expected, and it was uncertain exactly how this would impact TRIPLE's performance.

Figure 2 shows the results of the Phase I tests on BABBAGE. It presents asymptotic parallel-efficiency data for up to 2700 processors. Earlier results for KRAKEN are also included. The tests were carried out using a fixed per-processor problem size so that as larger numbers of processors are employed, a larger total problem size is treated.

The variable C plotted is proportional to the wall-clock time per compute cycle. Lower C is better, and a flat (constant C) curve indicates perfect 100 percent parallel efficiency. Fitting the KRAKEN data for NCPU>500 indicates a 99.976 percent (95 percent) parallel efficiency according to Amdahl's law (the scaled grind time).

Clearly the concern that TRIPLE would not perform as well on BABBAGE is not the case. For every value of NCPU tested (except for NCPU=2000) BABBAGE outperformed KRAKEN. On average TRIPLE was 3.6 percent faster on BABBAGE, and for particular values of NCPU (e.g., 1200, 1600, and 2400), BABBAGE was significantly faster (e.g., 13, 15, and 7 percent, respectively).** BABBAGE also exhibited enhanced performance for the largest values of NCPU tested. The reasons for this enhanced performance for certain values of NCPU are not fully understood, but it was exploited as often as possible when performing CAP simulations. With regard to the specific concerns mentioned above, it appears that IBM has improved its inter-processor communication algorithm in the time between scalability tests performed on KRAKEN in December 2004 and on BABBAGE in December 2006. This is because the anticipated near factor-of-two increase in communication time that might be the case for twice as many processors sharing only marginally more bandwidth per node was not observed.

Also, when TRIPLE was timed while using only 15 of the available 16 processors per node, only a 2-to-6 percent speed increase was observed, making it much less compelling to leave one processor per node idle on BABBAGE, especially given that more nodes (and therefore longer queue wait times) would be required. Therefore, as Figure 2 makes clear, despite initial concerns, the IBM P5+ (BABBAGE) at the NAVO MSRC proved itself to be an excellent platform on which to run 3D-FFT-based codes like TRIPLE.

BABBAGE's inter-node bandwidth is about 10 percent larger than KRAKEN's.

*** To put these measured speed increases in perspective: BABBAGE (7.6 Giga-Floating Point Operations (Gflops) per processor) is theoretically 11.8 percent faster than KRAKEN (6.8 Gflops per processor).*

References

1. Wroblewski, D., O.R. Coté, J.M. Hacker, and R.J. Dobosy, "Cliff-Ramp Patterns and Kelvin-Helmholtz Billows in Stably Stratified Shear Flow in the Upper Troposphere: Analysis of Aircraft Measurements," Journal of Atmospheric Science, 2007 (submitted).
2. Werne, J., T. Lund, B. Pettersson-Reif, P. Sullivan, and D. Fritts, "CAP Phase II Simulations for the Air Force HEL-JTO Project: Atmospheric Turbulence Simulations on NAVO's 3000-Processor IBM P4+ and ARL's 2000-Processor Intel Xeon EM64T Cluster," 15th DoD HPC User Group Conference, Nashville, TN, June, 2005.

Lattice Model of

Jeffrey Yepez, Air Force Research Laboratory, USAF
George Vahala, Department of Physics, William & Mary
Linda Vahala, Department of Electrical & Computer Engineering, Old Dominion University
Min Soe, Department of Mathematics & Science, Rogers State University
Sean Ziegeler, High Performance Computing Modernization Program, Mississippi State University

The lattice Boltzmann model is an optimal way to model a vast many-particle system (such as a fluid) with complicated particle-particle interactions and irregular boundary conditions. With it the fluid dynamicist can achieve higher nonlinearilty (measured by Reynolds number), unconditional stability and accuracy, with the less memory and processor time than with other models of turbulence for situations with tortuous boundaries. For an engineer, it is simple to code, runs perfectly on parallel supercomputers, and is suited to a plethora of computational physics applications. As demonstrated here, it is a competitive alternative to large eddy simulations with Smagorinsky sub-grid closure. To theorists and experimentalist in quantum information science, its kinetic transport equation is a special case of the quantum lattice Boltzmann equation governing a parallel array of quantum processors, a quantum computer architecture available today. Presented are turbulent fluid simulations using the lattice Boltzmann model carried out on the supercomputer Babbage. As an illustration of the efficiency of the lattice model, presented is a discovery of a universal range in the morphological evolution of the laminar-to-turbulent flow transition: the breaking subrange.

INTRODUCTION

In the defense community there are a plethora of neutral-fluid flows problems, particularly flows by regions with non-trivial spatial boundaries, such around an aircraft fuselage or a ship hull, through a jet engine compartment, or over the Earth's surface. While analytical solutions to such problems remain elusive and generally intractable, today they are within the reach of the computational physicist. To tackle such problems, practitioners usually resort to direct numerical simulation methods; yet here the amount of computer memory and processing time grows faster that the number of desired computed field points. Even in cases with simple or periodic boundaries, where more efficient numerical representations are available (such as a psuedo spectral

Continued Page 14

The turbulent flow shown in the background was computed on a 1600^3 grid using the Q15 lattice model at Re=165,649.

Fluid Turbulence

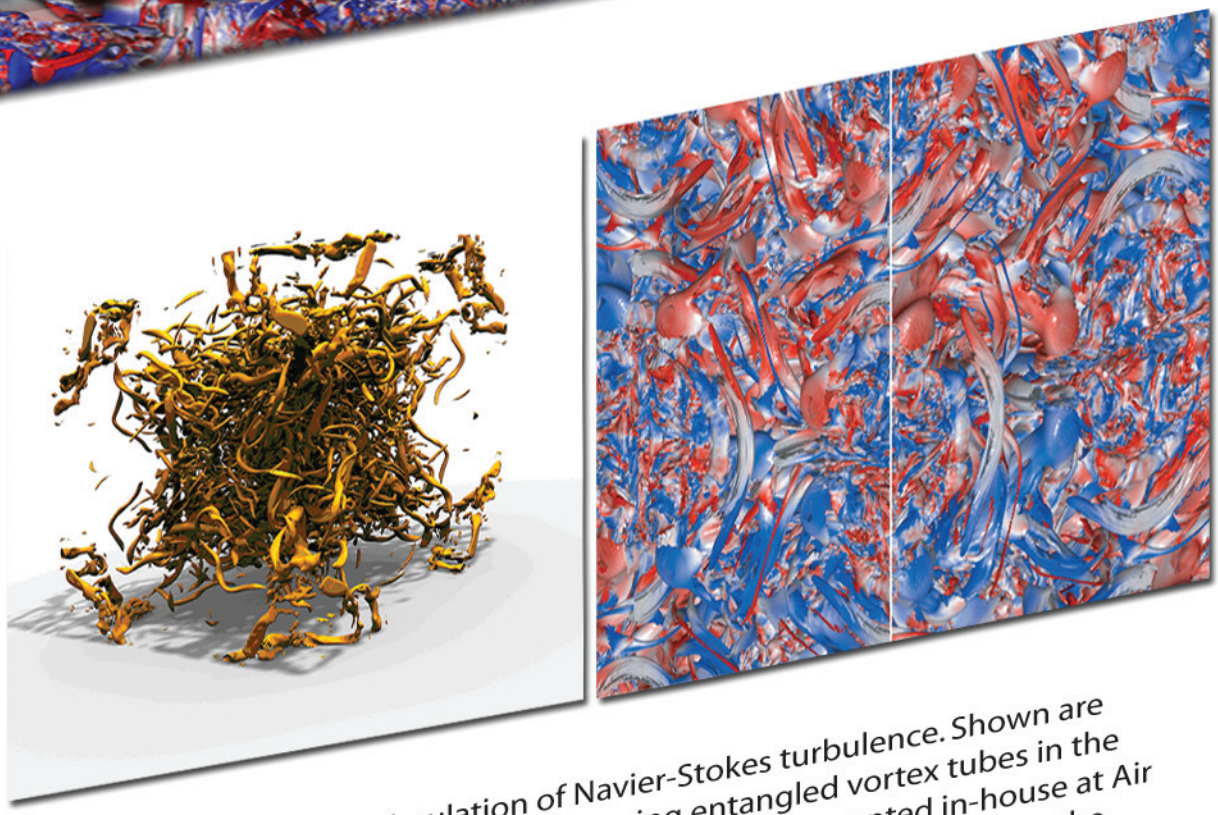

Figure 1. Supercomputer simulation of Navier-Stokes turbulence. Shown are streamlines and a vorticity isosurface, showing entangled vortex tubes in the breaking super-range. The computational method was invented in-house at Air Force Research Laboratory (AFRL) but required numerous groups across the country to bring to fruition. The authors have the fastest executing code for computational fluid dynamics simulations of turbulence to date as executed on the newest defense department supercomputer, BABBAGE, located at the NAVO MSRC at the Stennis Space Center in Mississippi, through a Capability Applications Project (CAP) II grant. They tested new sub-grid models of turbulence developed using lattice Boltzmann equation techniques, an entropic method, and a Smagorinsky closure method. These codes are close cousins to quantum algorithms for aerodynamics. The author's CAP II simulations used over one million hours (using a dedicated block of thousands of high performance processors over the period of four weeks), generating terabytes of data per day. These large-scale simulations provide a better understanding of the morphological evolution and structural development of turbulence in fluids, but they also give a preview of the kind of numerical output obtainable from future quantum computers.

approach pioneered by G.I. Taylor in the 1930's and S.A. Orszag et al. in the 1970's),[1] the scaling of computer resources with grid size is daunting.

This computational complexity has translated into significant annual costs to defense department high performance computing offices purchasing large numbers of processing elements (typically thousands or ten of thousands), configured in massive parallel computing arrays. But this cost, increasing year after year, has been borne so engineers can solve mission critical fluid problems, perpetually requiring higher resolution grids and more accurate and faithful simulations. For example, high resolution flow simulations are vital to aeronautical engineers designing the shape of advanced fighter jets or unmanned aerial vehicles or submarines to economize on fuel consumption and minimize maneuvering instabilities and wakes, to propulsion engineers designing nozzle and flow control orifices to maximize thrust, or to meteorologists trying to understand intermittent turbulence induced in the upper atmosphere under the jet stream to maximize laser propagation from airborne platforms. Yet to the theoretical physicist, the situation is even more dire: the prediction of any aspects of turbulence (beyond Kolmogorov's 1941 universality hypothesis), using advanced statistical methods and perturbation methods, borrowed from triumphant quantum field theory and statistical mechanics, remains the oldest and most prominent of classical grand challenge problems, open now for over 150 years.

This dire situation arises because, even in the macroscopic limit, strong correlations and feedback mechanisms between large scale and small scale flow structures, over many decades of spatial separation, dominate the overall flow evolution. The clearest high level picture capturing the essential physics of this problem, with restrictioned attention to divergence free and low Mach number flows, are the incompressible Navier-Stokes equations. The strong correlation between disparate scales is captured by the extremely simple non-local convective derivative (the second order nonlinearity in the velocity field).

Laminar to Turbulent Flow Transition

The lattice model now affords a deep insight into the origin and essential inner workings of free shear turbulence. This is a kinetic lattice gas model, the clearest low level picture

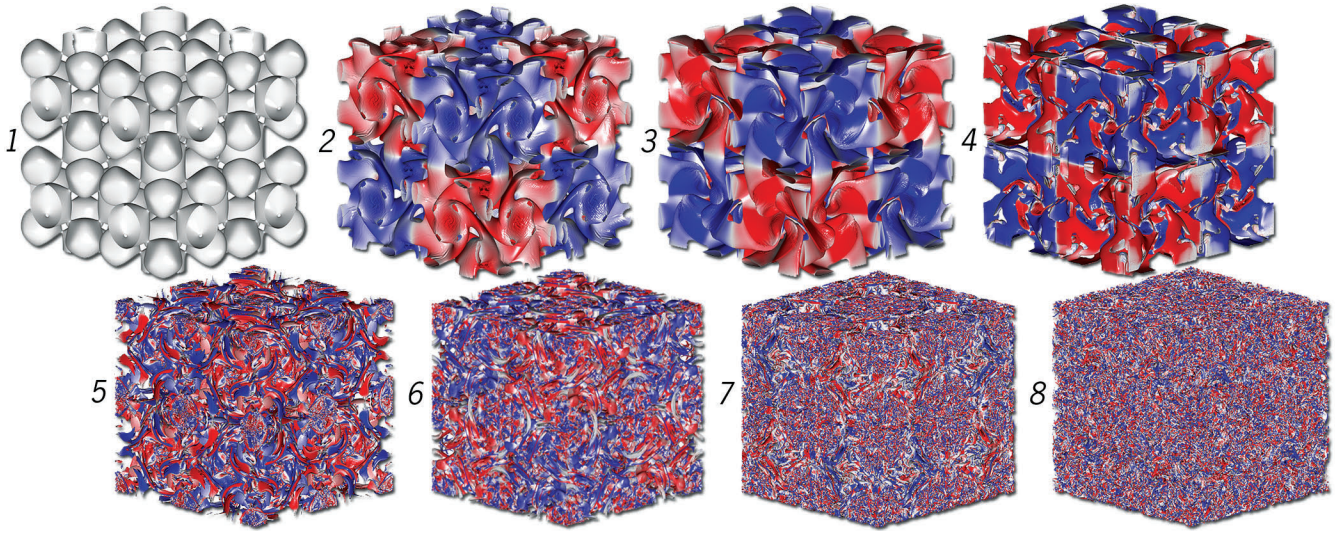

Figure 2. Surfaces of constant enstrophy ($\frac{1}{2}\int dV|\vec{\omega}|^2$ where $\vec{\omega} = \nabla \times \vec{u}$) illustrating an incompressible fluid's morphological evolution from t = 0 up to t = 7, 000Δt iterations, in time steps of 1, 000Δt on a cubical cartesian grid of size L = 512Δx. Surface coloring uses $\vec{v} \cdot \vec{\omega}$ (red equals -1 and blue 1). This 3+1 dimensional turbulent neutral fluid simulation run was on the newest defense department supercomputer, BABBAGE, using the entropic lattice Boltzmann equation with 15-body particle-particle collisions (ELB-Q15 model) computed at every lattice site at each time step. At each site, local relaxation of the single-particle probability distribution a desired equilibrium function, represented as a low Mach number polynomial expansion. The initial flow is a Kida and Murakami profile[2] with a super cell size set to $L_\circ = 512\Delta x$, the total grid size. So the flow configurations within all 8 octants of the large grid are initially identical. The characteristic flow speed is $u_0 = 0.07071\frac{\Delta x}{\Delta t}$. The collisional inversion parameter is set to $\beta = 0.99592$, corresponding to a kinematic viscosity of $\nu_\circ = 6.8 \times 10^{-4}\frac{\Delta x^2}{\Delta t}$, for $\alpha = 2$. The Reynolds number is $\mathrm{Re} = \frac{L_\circ u_\circ}{\nu_\circ} = 53,024$. The resulting turbulence is not fully resolved down to the dissipation scale, which in the model is the cell size Δx. To do this, set $L = \mathrm{Re}^{\frac{3}{4}} \sim 2,078\Delta x$. So the flow is under resolved by about a factor of 4.

correctly capturing the essential physics and hydrodynamics of the problem. As the well known Ising lattice gas model is fundamental to a statistical mechanics understanding of the essential physics of ferromagnetism and the order-disorder phase transition, the kinetic lattice gas model is fundamental to a dynamical mechanics understanding of fluids and the laminar-to-turbulent flow transition. See Appendix A for a brief mathematical overview of the model.

A new universal feature of the laminar-to-turbulent transition becomes clear in high Reynolds number simulations. Following an initial period of laminar flow with vortex sheet stretching, and preceding the final inertial subrange period of isotropic and homogeneous turbulent flow with self-similar vortex tubes, there exists a well defined interim period, here termed the *breaking subrange*. The hypothesis is that this subrange manifests self-organized criticality. The breaking subrange is dominated by anisotropic and non-homogeneous turbulent flow. Avalanches occur intermittently, where large coherent spatial structures grow, become unstable under maximal shear, and subsequently break into isotropic and homogeneous turbulence. These avalanches occur progressively in time across the entire space. See the bottom image of Figure 1 which shows the kind of anisotropic turbulent flow that occurs during an avalanche event, with young vortex tendrils. The time rate of change of the enstrophy has a generally negatively sloped avalanche cascade, with marked peaks, indicating the successive breaking points, as shown in Figure 3.

By carefully comparing the renderings in Figure 2 to the enstrophy plots in Figure 3, it is possible to see three distinct morphological stages of the flow: vortex stretching range, breaking subrange, and the inertial subrange. The morphological evolution transitions from (1) large vortex sheets to (2) convoluted vortex sheets with virgin vortex tendrils to (3) small entangled vortex tubes.

Q Models versus S Models: Closure of Sub-Grid Effects

The continuum hydrodynamical equations are projections of the Entropic Lattice Boltzmann (ELB) equation, a projection down from the $Q \times L^D$-dimensional kinetic phase space on to a $4 \times L^D$-dimensional hydrodynamics null space.* This projection recovers the Navier-Stokes equations in the Chapman-

Continued Next Page...

* A subset of all the eigenvectors in the Q-dimensional kinetic space has zero eigenvalues. These eigenvectors span the hydrodynamic null space. It is 4 dimensional because the authors' kinetic lattice gas model conserves probability (mass) and probability flux along the spatially orthogonal directions (three components of the momentum vector).

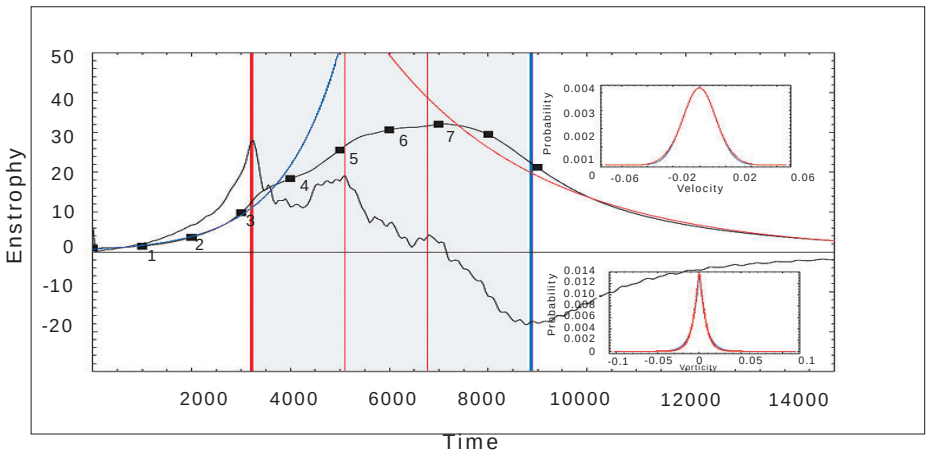

Figure 3. Plot of enstrophy versus time (smooth black curve) showing three stages in the morphological evolution: (1) vortex stretching range ($t < 3,200\Delta t$), (2) breaking subrange ($3,200 < t < 9,000\Delta t$), and (3) inertial subrange ($t > 9,000\Delta t$). The enstrophy is normalized so that at t = 0 it is unity. The isovalues used to visualize the 8 images in Figure 2 are shown (black squares). Stage 1: The initial exponential increase (blue curve) of enstrophy designates the initial *vortex stretching range* with characteristic laminar flow. There is excellent agreement between the analytical fit (blue curve) and the enstrophy data (black curve). Stage 2: The time derivative of the enstrophy curve (jagged black curve) is also plotted. The time period of generally negative slope of the enstrophy derivative (gray shaded region) is here termed the *breaking subrange*, where large scale anisotropic ordering of turbulence occurs and intermittently breaks down over time. The first major breaking point occurs at about $t = 3,200\Delta t$ (vertical red line) and subsequent intermediate avalanches occur at about $t = 5,100\Delta t$ and $t = 6,750\Delta t$ (thin vertical red lines), respectively. Stage 3: The final exponential decrease of enstrophy (red curve) designates the inertial subrange where the homogeneous and isotropic "small scale" turbulent flow morphology, with characteristic entangled vortex tubes, is organized in a spatially self-similar way. Here the energy spectral density obeys the Kolmogorov universality hypothesis, the famous $k^{-\frac{5}{3}}$ power law for energy cascade downward to smaller scales. The onset of the inertial subrange occurs close to $t > 9,000\Delta t$ (vertical blue line). Here the velocity probability distribution functions, for each component, are Gaussian (top inset) and the vorticity probability distribution function approaches an exponential (bottom inset). There is excellent agreement between the analytical fit (blue curve) and the enstrophy data (black curve). There exists a fourth stage of the morphology of turbulence at late times ($t > 14,000\Delta t$), not shown here, called the *viscous subrange*.

Enskog limit. This has an important consequence: there exist many "qubit" models (or Q models for short) with a different local stencils (i.e., lattice vectors sets with different finite point group symmetries and coordination number), which will also recover the Navier-Stokes equations asymptotically (continuous rotational symmetry).

ELB is ideal for large eddy simulation (LES) closures since in LES one typically deals with mean strain rates for modeling the eddy viscosity. These nonlocal fluid calculations are immediately recovered from simple local moments in ELB.

At this stage, the ELB runs are a validation of the method, and this is important because ELB is a crucial precursor to a viable quantum lattice Boltzmann method.[3] The good numerical agreement between the LES-LB (lattice Boltzmann equation with sub-grid Smagorinsky closure here referred to as an S model) and the basic ELB Q models is encouraging. Now fully convinced of the validity of ELB, work to speed it up for present day supercomputer implementations is underway. But the most promising opportunity is to simply build a type-II quantum computer,[4,5] which could far outstrip any classical supercomputers, for turbulence fluid simulations. In the fullness of time, the ELB will outpace the LES-LB, even without quantum computers.

There are a few reasons for this view. First, the kinetic lattice gases obey detail balance while sub-grid closure methods, such as the LES-LB, do not. And another advantage of the ELB over LES-LB is that ELB obeys the second law of thermodynamics while LES-LB and other LES methods do not necessarily obey the second law. Third, in the LES-LB the strain tensor must be computed at every site. With no-slip boundaries, computing the strain tensor becomes problematic. In contrast, ELB is purely local, so grid sites near boundaries are handled as easily as sites far away from boundaries. All these are important differences when the model is used for practical engineering grade applications.

Timings and Scaling

Turbulent dynamics are easily solved since the underlying kinetic equations, see (1) and (2) in the appendix, have only local algebraic nonlinearities in the macroscopic variables and simple linear advection. At this mesoscopic level there are various kinetic lattices (Q=15, 19, or 27) with different lattice vectors on a cubic lattice, which model the Navier-Stokes equation to leading order in the Chapman-Enskog perturbative asymptotics.

With the CAP data obtained on BABBAGE, the effects of the underlying lattice symmetry on the fluid turbulence statistics (through autocorrelation tensors of velocity, vorticity, pdfs of vorticity, and the like) can be determined, but there is not sufficient space to present details here. The Q15 model seems to be the most efficient model. An example output of this model is shown in Figure 2. Even on a relatively modest size 512^3 grid, researchers can achieve such a high degree of resolved nonlinearity (Re=26,512) that the consequent isotropic turbulence at the onset of the inertial subrange (at about $t\sim 9{,}000\Delta t$) outstrips the ability to visualize the myriad vortex tubes. Figure 2 is just a test case. The sustained floating point per second (MFLOPS/PE) achieved are the best of all scientific codes run, for example on the Earth Simulator, attaining over 67% of peak performance on 4,800 processors. Achieving the world's highest Reynolds number in the field of computational fluid dynamics should also occur in the near future. The current tested code on BABBAGE, on $1{,}600^3$ grid with 2,048 processing elements, can achieve a Reynold's number of Re=565,667. Modeling atmospheric scale turbulence, in the range of Re $\sim 10^6$, is possible today for the first time in the half century long history of numerical digital computers applied to aerodynamics.

Some runs with these models–in particular the Q15, Q19, Q27 and S27 models–have been completed. The advantages of these lower Q models are reduced wallclock times with less memory demands. The nonlinear convective derivatives of in the Navier-Stokes equation are recovered from purely local moments of kinetic space distribution function. This is the basic reason why ELB scales so well with PEs: the algorithm consists of simple local computations and streaming of information only to near neighbor grid sites.

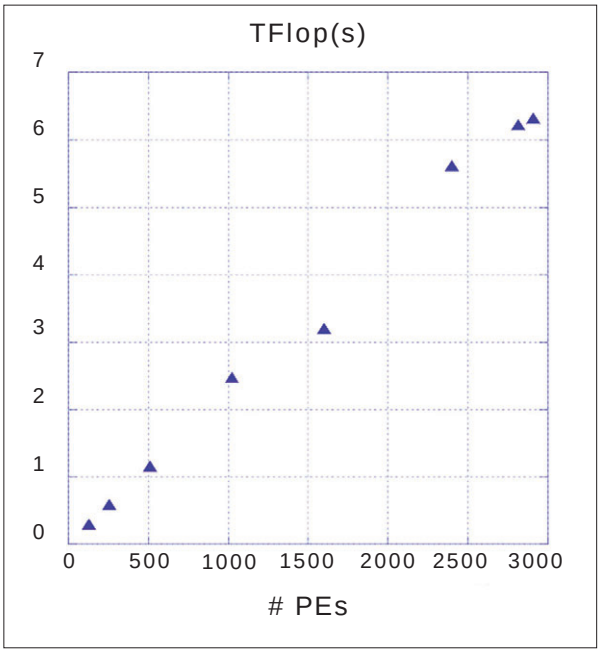

Figure 4. Terrabyte FLoating point Operations Per Second (TFLOPS) scaling of ELB-Q27 code on BABBAGE with number of Central Processing Units (CPUs).

During Phase I, investigating the scaling properties of the Q27 code, over 6.3 tera flops per second on the full 2912 processor elements available on BABBAGE was achieved. (See Figure 4.)

The LES-LB-S27 code, which no longer needs the solution of an entropy constraint equation, has also been tested in Phase I. It is less computationally intensive (due to the avoidance of log-calls and the need for a Newton-Raphson root finder at each spatial node and time iteration) and shorter wallclock time than the ELB-Q27 code. (See Table 1.)

No. PEs	GRID	MODEL	WALLCLOCK(s)	GFlop(s) per PE
2912	1950^3	ELB-Q27	7,554.7	2.17
2912	1950^3	ELB-Q19	5,602.7	2.24
2912	1950^3	ELB-Q15	4,798.4	2.05
2912	1950^3	LES-LB-S27	4,451.2	1.05

TABLE 1. The gigaflops per second per processor element for 2912 CPU runs on a 1952 × 1946 × 1950 grid for four lattice Boltzmann codes variants. The wallclock time is for 2,000Δt (lattice time steps). A full turbulence simulation takes about 54,000 time steps.

Quantum Information Processing

Quantum Information Processing (QIP) and quantum communications will be integral to the 21st century. For many years, QIP has been included in the Developing Science and Technologies List, in the section of critical information systems technology. It appears that superconductive quantum information processing may soon be elevated to the status of a militarily critical technology. There now exists the rather imminent possibility of the development of large quantum computer arrays, potentially far outstripping any supercomputers now used for defense department computations.

On Babbage Q27, Q19, Q15 lattice Boltzmann codes were tested, like the one shown in Figure 1. Using 2,048 processors it took two days to complete a single job. A 1,0243 grid takes about 44 hours for Q27 on 512 processors, while the corresponding run for Q15 takes about 28 hours.

The cost of the largest supercomputer parallel arrays annually adds to a significant fraction of a billion dollars for new government-owned systems the United States (e.g., 19,000 processor Franklin at DOE/NERSC cost about $50 million and occupies the space of a gymnasium).

Remarkably, exploiting quantum mechanical complexity new quantum device technology available today can be used to efficiently compute the collision operator of the Q15 lattice Boltzmann code, the basic engine of the code.

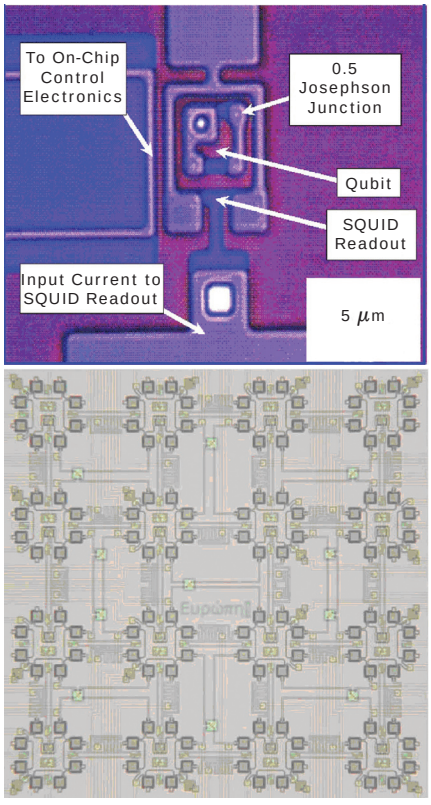

Shown is a newsworthy achievement of a first quantum information processor circuit to embody a qubit (top), a basic device for storing quantum information, designed by Terry Orlando of the Massachusetts Institute of Technology (MIT) and build at MIT Lincoln Laboratory at Hanscom AFB in 2000.[6,7] The circuit must be placed in a dilution refrigerator. Air Force Office of Scientific Research (AFOSR) supported novel quantum computing technology based on superconductive electronics under the Quantum Computation for Physical Modeling (QCPM) theme, and this research has recently found follow-on use. A superconducting wire loop with multiple Josephson junctions forms a qubit and such qubits are coupled together to make quantum logic gates. AFOSR funded this new solid-state technology, fabricated at the superconductive electronics foundry at MIT Lincoln laboratory. This helped establish the basic fabrication techniques to build scalable quantum computers and mapped the quantum control methodology, proven with Nuclear Magnetic Resonance (NMR) spectroscopy,[4,5] onto the field of superconductive electronics for quantum information processing, mapping pulse protocols to allow qubit-qubit logical operations, constituting a basic 2-qubit quantum processor. Going from a 2-qubit processor to a 16-qubit processor necessarily entails a significant applied research effort recently announced by D-Wave, a Canadian start-up company. D-Wave's 16-qubit processor, fabricated by the National Aeronautics and Space Administration (NASA), is shown (bottom).

APPENDIX A: KINETIC LATTICE GAS MODEL

In the lattice model, dynamics is projected into a discrete kinetic phase space. The logical "1" state (the excited state) of a qubit $|q\rangle$ associated with the spacetime point $(\vec{x}, t)$ encodes the probability f_q of the existence of a particle at that point moving with velocity $\vec{c}_q = \frac{\Delta \vec{x}_q}{\Delta t}$, where $\Delta \vec{x}_q$ are lattice vectors, for $q = 1, 2, \ldots, Q$. A fundamental property of the lattice model is that particle motions in momentum space and position space occur independently [8]. Particle momentum and position space motions are generated by the combination of an engineered qubit-qubit interaction Hamiltonian q and a free Hamiltonian, $-i\hbar \sum_q \vec{c}_q \cdot \nabla$ respectively.

All the particle-particle interactions, the 2-body up to and including all the Q-body interactions, generated by H' are mapped to a local collision function $H' \mapsto \Omega_q$ that depends on all the f_q's at the lattice site.[3]

In the type-II quantum computing case, quantum entanglement is localized among qubits associated with the same $(\vec{x}, t)$,[9] so:

$$f'_q(\vec{x}, t) = f_q(\vec{x}, t) + \Omega_q(f_1, f_2, \ldots f_Q) \quad (1)$$

$$f'_q(\vec{x}, t) = f_q(\vec{x}, t) + \Omega_q(f_1, f_2, \ldots f_Q) \quad (2)$$

where f_q and f'_q are called the incoming and outgoing probabilities, respectively. In the classical limit, there exists a fundamental entropy function

$$\mathscr{H}(f_1, \ldots, f_Q) = \sum_{q=1}^{Q} f_q \ln(\gamma_q f_q), \quad (3)$$

where the γ_q are self-consistently determined positive weights $(\sum_q \gamma_q = 1)$. Ω_q in (1) is determined by the constant entropy condition

$$\mathscr{H}(f'_1, \ldots, f'_Q) = \mathscr{H}(f_1, \ldots, f_Q). \quad (4)$$

It is (1) through (4) that constitutes the lattice model of fluid turbulence suited for parallel supercomputers.

In the Q-dimensional kinetic space, the equilibrium distribution f_q^{eq}, taken as a vector, is the bisector of the difference of the incoming and outgoing kinetic vectors:

$$f_q = f_q^{\mathrm{eq}} - \frac{1}{\alpha\beta}\Omega_q \quad (5)$$

$$f'_q = f_q^{\mathrm{eq}} + \left(1 - \frac{1}{\alpha\beta}\right)\Omega_q, \quad (6)$$

when $\lim_{\eta \to 0} \alpha\beta = 2$ and where f_q^{eq} is analytically determined by extremizing (4) subject to the two constraints of conversation of probability and probability flux. Eliminating Ω_q from (5) and (6) yields an operative collision equation simpler than (1)

$$f'_q = f_q + \alpha\beta\left(f_q^{\mathrm{eq}} - f_q\right) \quad (7)$$

Finally, eliminating f'_q in (2) and (7) yields the lattice Boltzmann equation

$$f_q(\vec{x} + \Delta\vec{x}_q, t + \Delta t) = f_q(\vec{x}, t) + \alpha\beta\left[f_q^{\mathrm{eq}}(\vec{x}, t) - f_q(\vec{x}, t)\right] \quad (8)$$

in the BGK collisional limit.[10] Since $\alpha\beta$ and f_q^{eq} are determined by (3), (8) is called the *entropic lattice Boltzmann equation*.[11]

The kinetic lattice gas model becomes equivalent to the Navier Stokes equations:

$$\partial_t \vec{u} + \vec{u} \cdot \nabla \vec{u} = -\nabla P + \eta \nabla^2 \vec{u} \quad \text{and} \quad \nabla \cdot \vec{u} = 0 \quad (9)$$

where $\vec{u} = \vec{u}(\vec{x}, t)$ is the vectorial flow field, where the pressure is proportional to the field density ($P = \rho c_s^2$, with spatial dimension D and sound speed $c_s = \frac{1}{\sqrt{D}}\frac{\Delta r}{\Delta t}$), and where the shear viscosity is the measure of dissipation (a renormalized transport coefficient for momentum diffusion). In the limit when $\Delta\vec{r} \to 0$ and $\Delta t \to 0$ and the hydrodynamic variables are cast as moments of the probability distribution, $(\rho c, \vec{u}) = \sum_q (c, \vec{c}_q) f_q$.[12] The hydrodynamic variables are independently evaluated at each spacetime point $(\vec{x}, t)$. The shear viscosity is analytically determined, and the result is

$$\eta = \rho c_s^2 \Delta t \left(\frac{1}{\alpha\beta} - \frac{1}{2}\right), \quad (10)$$

so the model approaches the inviscid limit where $\eta \to 0$ as $\alpha\beta \to 2$.[13]

Acknowledgments

This work was supported in part by a grant of computer time from the Department of Defense High Performance Computing Modernization Program, with computational support from the Major Shared Resource Centers: the Naval Oceanographic Office and the U.S. Army Engineer Research and Development Center. Long-term support was provided through the Computational Mathematics Program of the Air Force Office of Scientific Research for algorithm development research beginning in 1992 under the Novel Strategies for Parallel Computing Initiative at the Air Force Research Laboratory.

References

1. Orszag, S. A., and G.S. Patterson, "Statistical Models and Turbulence," Springer-Verlag, New York, 1972. M. Rosenblatt and C. Van Atta, editors.
2. Kida, S. and Y. Murakami, "Kolmogorov Similarity in Freely Decaying Turbulence," Physics of Fluids, 30(7): 2030-2039, July 1987.
3. Yepez, J., "Quantum Lattice-Gas Model for Computational Fluid Dynamics," Physical Review E, 63(4):046702, March 2001.
4. Pravia, M., Z. Chen, J. Yepez, and D.G. Cory, "Towards an NMR Implementation of a Quantum Lattice-Gas Algorithm," Computer Physics Communications, 146(3):339–344, 2002.
5. Chen, Z., J. Yepez, and D.G. Cory, "Simulation of the Burgers Equation by NMR Quantum-Information Processing," Physical Review A, 74(4):042321, October 2006.
6. Mooij, J.E., T.P. Orlando, L. Levitov, L. Tian, C.H. van der Wal, and S. Lloyd, "Josephson Persistent-Current Qubit," Science, 285:1036–1039, 1999.
7. Orlando, T.P., J.E. Mooij, L. Tian, C. H. van der Wal, L.S. Levitov, S. Lloyd, and J.J. Mazo. "Super-Conducting Persistent-Current Qubit," Physical Review B, 60(22):15398–15413, 1999.
8. Yepez, J., "Relativistic Path Integral as a Lattice-based Quantum Algorithm," Quantum Information Processing, 4(6):471-509, December 2005.
9. Yepez, J., "Open Quantum System Model of the One-Dimensional Burgers Equation with Tunable Shear Viscosity," Physical Review A, 74(4):042322, October 2006.
10. Bhatnagar, P. L., E. P. Gross, and M. Krook, "A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems," Physical Review, 94(3): 511-525, May 1954.
11. Boghosian, B. M., J. Yepez, P. Coveney, and A.Wagner, 2001, "Entropic Lattice Boltzmann Methods," Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 457(2007):717-766, March 2001.
12. Boghosian, B. M., P. J. Love, P. V. Coveney, I. V. Karlin, S. Succi, and J. Yepez, "Galilean-Invariant Lattice-Boltzmann Models with H Theorem," Physical Review E, 68(2):025103, August 2003.
13. Keating, B., G. Vahala, J. Yepez, M. Soe, and L. Vahala, "Entropic Lattice Boltzmann Representations Required to Recover Navier-Stokes Flows," Physical Review E75, (3):036712, 2007.

Proxy License...Continued

seem to be sensitive to buffer size like FlexLM.

The second advantage to using proxy is that it uses a simple, lightweight, design that requires a minimal amount of computational resources but still managed to scale easily. More about the scalability and adding new license types is discussed later in this article.

Lastly, and maybe most importantly, it was available for free.

Implementation

The NAVO MSRC staff deployed three separate proxy servers. Each proxy server provides a one-to-one relationship with one of the three HPCMP consolidated license servers. To provide redundancy and increase fault tolerance, each server is hosted in a different physical location with independent power and local network connections.

Each proxy server runs multiple instances of the proxy daemon. The proxy software is designed so that one instance of the daemon handles only one particular type of software license. This is because each of the software licenses provided by the HPCMP consolidated license servers are associated with a different network port, and each instance of the proxy daemon can listen on ony one port at a time. This configuration makes it easy to add proxying capabilities for new licenses: Simply start another proxy daemon with the correct port options and the licenses become available without affecting any of the other software licenses already being offered.

Compute nodes in our computational clusters reach these local proxy servers via each cluster's front-end nodes. The /etc/hosts file on the compute nodes is configured so that the IP addresses for the local proxy servers are returned when querying for the IP addresses of the HPCMP consolidated license servers. This "tricks" the compute nodes into thinking that they are talking to the HPCMP consolidated license servers when they are really using the proxy servers. This allows back-end compute nodes to obtain licenses without then having to have access outside of the local area network.

Conclusion

A proxy license server provides the balance between security and capability. It allows compute nodes on large, HPC cluster to obtain license keys without directly communicating with an external license server.

Furthermore, through the use of Open Source software, this solution is very inexpensive. The proxy approach to software license serving provides an unusual optimization of a three critical factors: security, functionality, and cost. "Real" world solutions are not always this balanced.

The Porthole

Chaplain McClellan looks on as RMDL McGee congratulates the outgoing Commanding Officer of NAVOCEANO CAPT Brown... and welcomes aboard CAPT Cousins as the new Commanding Officer.

Brigadier General Capasso and his wife tour the NAVO MSRC.

Members of Senator Trent Lott's staff visit the NAVO MSRC.

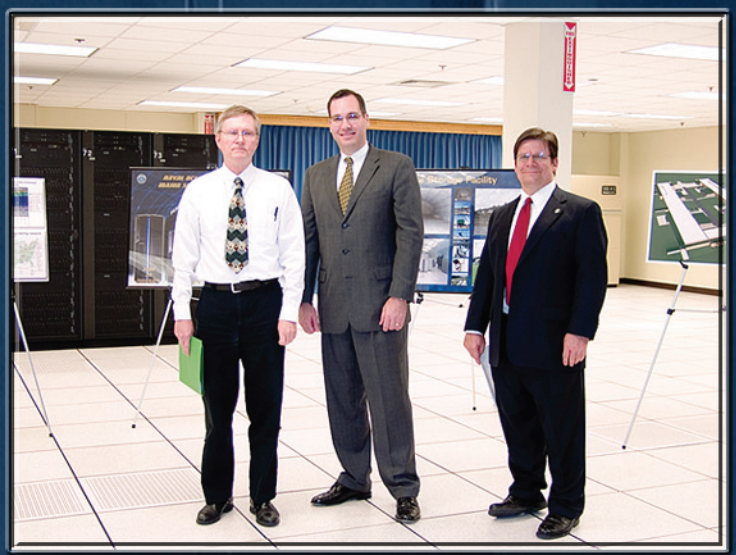

Tom Dunn, NAVO MSRC Director; Mark Honecker, Executive Director and Chief of Staff, U.S. Fleet Forces Command; and Ed Gough, Deputy/Technical Director COMNAVMETOCCOM visit the NAVO MSRC.

CAPT Jorge Ibarra, Director of the Chilean Navy Hydrographic and Oceanographic Service (SHOA), and students from the U.S., Chile, Mexico, and Ecuador tour the NAVO MSRC.

Bobby Knesel, NAVO MSRC Deputy Director; Maxine Brown, National Centers for Environmental Prediction (NCEP) Central Operations; Tom Dunn, NAVO MSRC Director; Ben Kyger, NCEP Acting Director, Central Operations tour the NAVO MSRC.

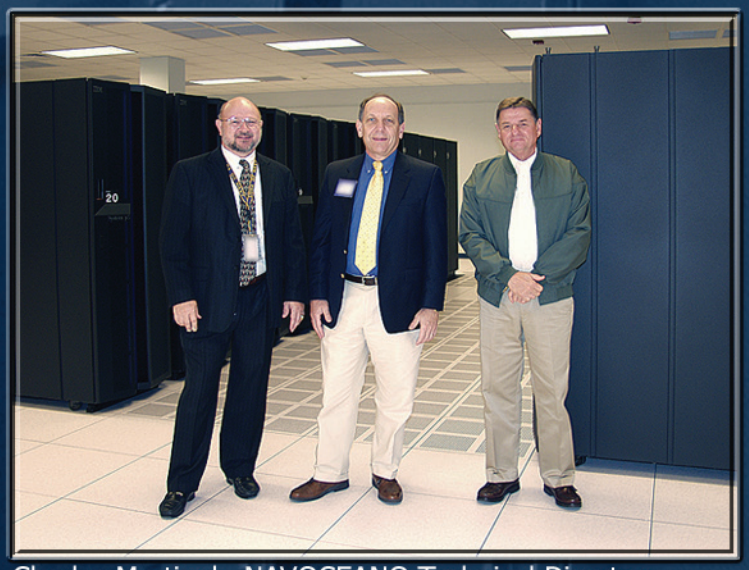

Charles Martinek, NAVOCEANO Technical Director; Dr. Louis Uccellini, National Centers for Environmental Prediction (NCEP); and Bobby Knesel, NAVO MSRC Deputy Director tour the NAVO MSRC.

Representatives from the Aeronautical Systems Center MSRC visit the NAVO MSRC to tour the Data Storage Facilities.

Dr. Robert Woolsey, Director, University of Mississippi Mineral Resources Institute (MMRI) and staff tour the NAVO MSRC.

The Basic Oceanography Accession Training (BOAT) Class students visit the NAVO MSRC.

Representatives of the High Performance Computing Modernization Office and ERDC MSRC visit the NAVO MSRC and tour the Data Storage Facilities.

Navigator Tools and Tips

Using DoD-Kerberized MPSCP on BABBAGE, KRAKEN, and ROMULUS

Sheila Carbonette, NAVO MSRC User Support

The DoD-kerberized MPSCP (Multi-Path Secure Copy), a high-performance remote file copy utility, is now available on the unclassified IBM P5+ system, BABBAGE, and the IBM P4+ systems, KRAKEN and ROMULUS.

This utility allows users to transfer files at higher transfer rates because of the ability to spread the data over multiple Transmission Control Protocol (TCP) streams to segment a file into pieces and send the pieces in parallel, reassembling them into a single file on the remote system.

MPSCP should be available in your default search path, but if for some reason it is not, it can be found in the */node/bin* directory. It is currently configured to transfer files to/from/between the local IBM systems and other MSRC sites. Note that MPSCP is limited to four data streams per transfer to prevent negative impact on the system, and it only uses the configuration file located at */node/etc/mpscp.conf*.

Please note that the utility is not currently available on the archive systems, JULES and VINCENT.

MPSCP syntax is similar to the commands *krcp* and *scp*. However, the command is limited in that you must have a kerberos ticket on the system from which you initiate the command.

What follows are several examples of using MPSCP:

1. Transfer a 2 Gigabyte (GB) file from */scr* work directory on KRAKEN to */scr* work directory on BABBAGE:

```
kraken% mpscp -w 4 /scr/shecar/2GBfile babbage:/scr/shecar
```

2. Transfer a 2 GB file over DREN from */scr* work directory on KRAKEN to home directory on an Aeronautical System Center (ASC) MSRC platform:

```
kraken% mpscp -w 4 /scr/shecar/2GBfile falcon.asc.hpc.mil:
```

3. Transfer the */scr* work directory, modeldir, on BABBAGE to home directory on an ASC platform:

```
babbage% mpscp -w 4 -r /scr/shecar/modeldir falcon.asc.hpc.mil:
```

4. Transfer all fortran source code files in the */scr* work directory, modeldir, on KRAKEN to home directory on BABBAGE:

```
kraken% mpscp -w 4 /scr/shecar/modeldir/*.F babbage:
```

The utility also has a "*-v*" option that can be used for debugging and a "*-p*" that can be used to preserve modification times. More information on these options can be found in the on-line man page or by contacting the NAVO MSRC User Support at msrchelp@navo.hpc.mil, (800) 993-7677, or (228) 688-7677.

Director's Corner...Continued

(ARSC), and the Maui High Performance Computing Center (MHPCC).

The rate is anticipated to increase further as larger HPC systems become operational at these sites in the next few years. Like the NAVO MSRC upgrades, we introduced STK T10000 tape drive and cartridge technology into the RSF storage environment.

As a result, the STK PowderHorn silos that currently house HPCMP RSF data have been upgraded, and six T10000 drives have been installed. The amount of data stored at the NAVO MSRC and RSF facility is projected to increase from 5.2 petabytes (10**15) in January 2007 to approximately 12.5 petabytes by January, 2009.

We have been working with National Aeronautics and Space Administration (NASA) to establish a shared resilient, 10 gigabit, northbound communications service for the Stennis Space Center. This initiative will improve the operational availability of the NAVO MSRC for the nationwide DOD HPCMP user community. A contract award has been made, and the northbound service should be a reality in the June/July timeframe.

As always, we invite you to contact us and let us know how we can better serve you.

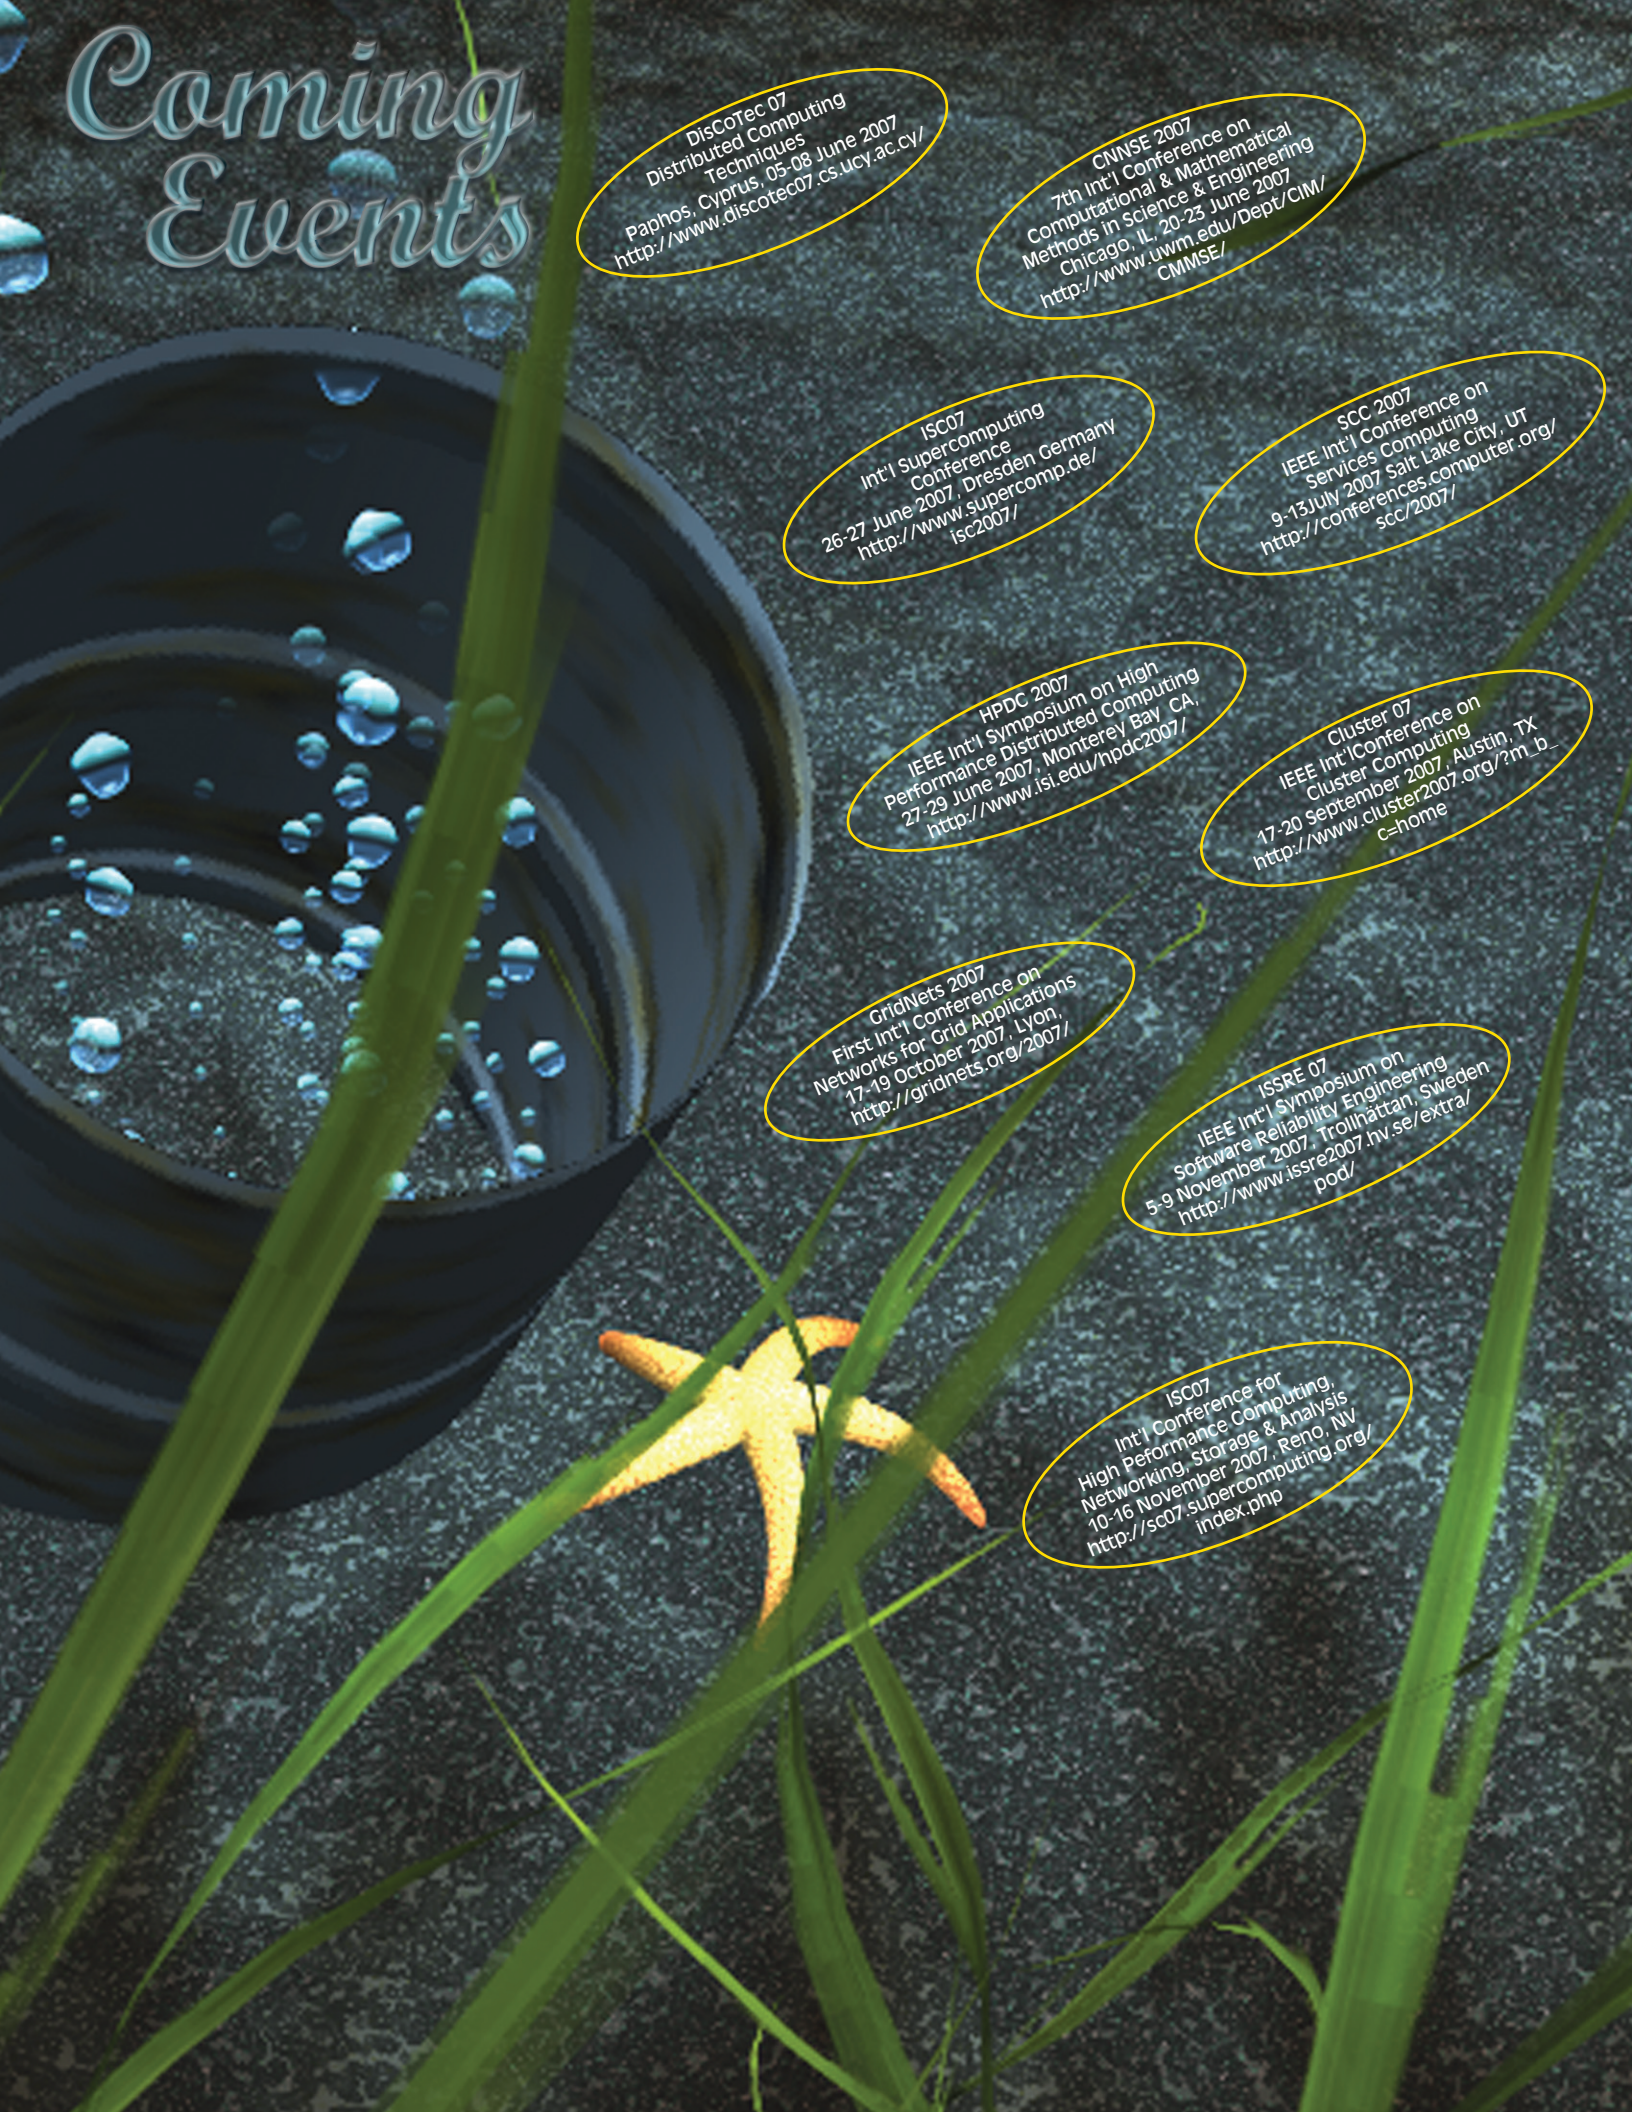
Coming Events
DisCoTec 07
Distributed Computing Techniques
Paphos, Cyprus, 05-08 June 2007
http://www.discotec07.cs.ucy.ac.cy/
CNNSE 2007
7th Int'l Conference on Computational & Mathematical Methods in Science & Engineering
Chicago, IL, 20-23 June 2007
http://www.uwm.edu/Dept/CIM/CMMSE/
ISC07
Int'l Supercomputing Conference
26-27 June 2007, Dresden Germany
http://www.supercomp.de/isc2007/
SCC 2007
IEEE Int'l Conference on Services Computing
9-13July 2007 Salt Lake City, UT
http://conferences.computer.org/scc/2007/
HPDC 2007
IEEE Int'l Symposium on High Performance Distributed Computing
27-29 June 2007, Monterey Bay CA,
http://www.isi.edu/hpdc2007/
Cluster 07
IEEE Int'lConference on Cluster Computing
17-20 September 2007, Austin, TX
http://www.cluster2007.org/?m_b_c=home
GridNets 2007
First Int'l Conference on Networks for Grid Applications
17-19 October 2007, Lyon,
http://gridnets.org/2007/
ISSRE 07
IEEE Int'l Symposium on Software Reliability Engineering
5-9 November 2007, Trollhättan, Sweden
http://www.issre2007.hv.se/extra/pod/
ISC07
Int'l Conference for High Performance Computing, Networking, Storage & Analysis
10-16 November 2007, Reno, NV
http://sc07.supercomputing.org/index.php

Naval Oceanographic Office * Major Shared Resource Center
1002 Balch Boulevard, Stennis Space Center, Mississippi 39522