MAUI HIGH PERFORMANCE COMPUTING CENTER

APPLICATION BRIEFS 2007

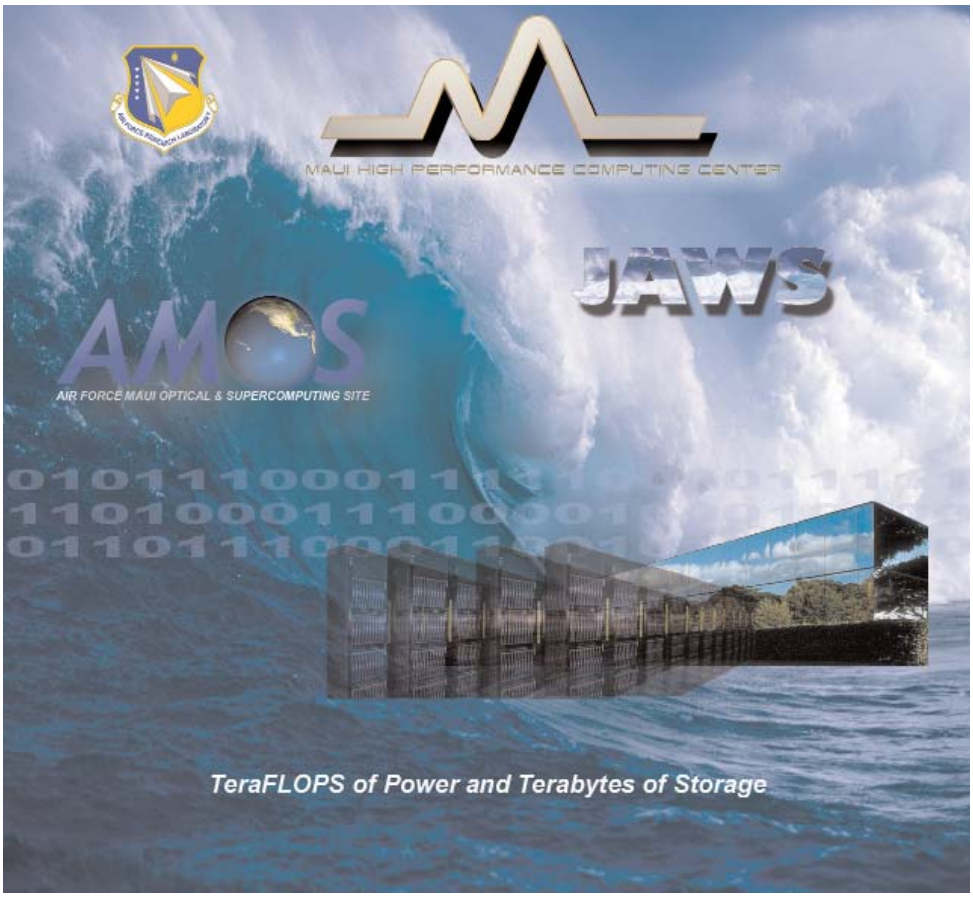

MAUI HIGH PERFORMANCE COMPUTING CENTER

550 Lipoa Parkway, Kihei-Maui, HI 96753
(808) 879-5077 • Fax: (808) 879-5018
E-mail: info@mhpcc.hpc.mil
URL: www.mhpcc.hpc.mil

An Air Force Research Laboratory Center Managed by the University of Hawaii

WELCOME

This is the thirteenth annual edition of Maui High Performance Computing Center's (MHPCC) *Application Briefs* which highlights some of the successes our staff and clients have achieved this year.

MHPCC, established in September 1993, is an Air Force Research Laboratory (AFRL) Center managed by the University of Hawaii. A leader in scalable parallel computing technologies, MHPCC is chartered primarily to support the Department of Defense (DoD) and other federal government organizations.

MHPCC offers an innovative environment for High Performance Computing (HPC) applications. This includes:

- Computational Resources: Stable and secure parallel computing platforms for prototyping, benchmarking, and testing applications. MHPCC is ranked as one of the premier HPC centers in the Department of Defense in terms of computational capabilities.

- High-Speed Communications Infrastructure: OC12 connections, offering 622 megabit per second (Mbps) capacity, provide direct access to MHPCC resources — over the Defense Research and Engineering Network (DREN) and the Hawaii Intranet Consortium (HIC).

- Support Services: An expert staff provides MHPCC users with systems, network, and applications support in addition to assistance with code porting, optimization, and application development.

MHPCC is a well-established member of the High Performance Computing community, participating in collaborations and partnerships that extend its capabilities. MHPCC is a direct contributor to the Department of Defense as a:

- Allocated Distributed Center within the DoD High Performance Computing Modernization Program (HPCMP). MHPCC provides resources to the DoD research community, as well as Pacific Region DoD organizations, including the Air Force's Maui Space Surveillance Complex.

- Center within the Air Force Research Laboratory. MHPCC works closely with DoD and other government researchers to support Research, Development, Testing, and Evaluation (RDT&E) efforts.

- Member of Hawaii's growing science and technology community.

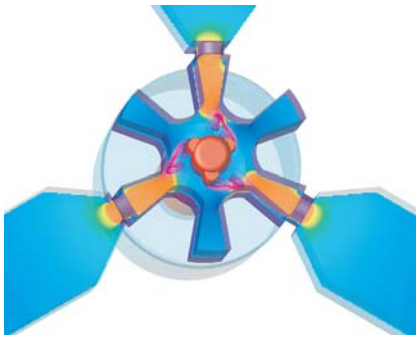

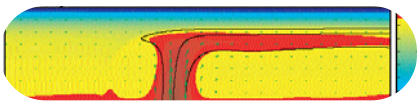

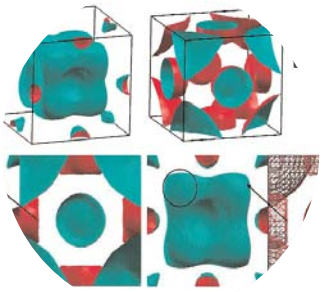

APPLICATION BRIEFS

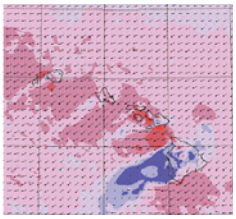

The user application briefs in this publication represent selected research efforts that have taken place at MHPCC during 2007. Each Application Brief was written by an individual researcher or research team, and reflects their first-hand experiences using MHPCC resources. These articles reflect the diverse nature of our clients and projects.

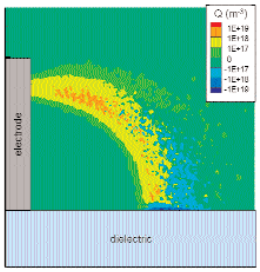

The Application Briefs in this document are the result of the efforts of more than 50 authors. We acknowledge the contributions of each of these individuals and are grateful for their work. We welcome back those authors who have become regular and frequent contributors. We also welcome those making their MHPCC Application Briefs debut this year.

The shaded box at the top of each brief's first page is a short summary of the article. Author and/or organizational contact information can be found in the shaded box at the end of each brief. The notation at the bottom of each page indicates each project's primary functional area (DoD, Government, or Academic).

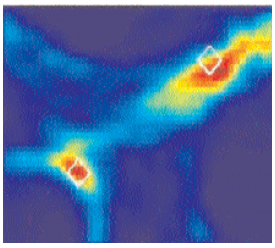

All the efforts described in this document were performed using resources provided by the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP). Additional sponsorship has come from a variety of Research, Development, Test and Evaluation sources, including Research Laboratories in the Defense Services.

Thank you for your support.

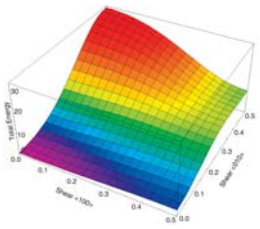

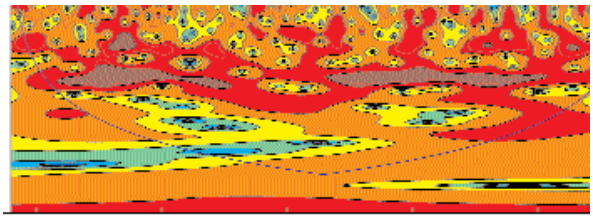

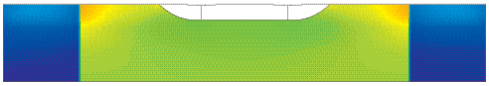

TABLE OF CONTENTS

The Kernal Energy Method: Application to a Collagen Triple Helix **1**
Lulu Huang, Lou Massa , and Jerome Karle

High Performance Computing Software Applications for Space Situational Awareness **2**
Francis Chun, Concetto Giuliano, Paul Schumacher, Charles Matson,
Kathy Borelli, Kevin Roe, and Bruce Duncan

Tracking Humpback Whales with Passive Acoustics **4**
Eva-Marie Nosal and L. Neil Frazer

High Resolution Forecasts to Support AMOS Using WRF **6**
Kevin P. Roe

Plasma Actuator Modelling **8**
Gabriel I. Font-Rodriguez, Connor Caples, Shawn Hackett, and Cordell Hachinsky

Mantle Convection, Melting, and Hotspot Geochemistry **10**
Todd Anthony Bianco, Garrett Ito, Jeroen van Hunen, Maxim Ballmer, and John Mahoney

The Electronic Mechanism of Nickel Titanium Shape Memory Behavior **12**
Nicholas B. Hatcher, Oleg Yu. Kontsevoi, and Arthur J. Freeman

Parallel Smart Structural Design Optimization System (PASOS) **14**
Rong C. Shieh

The Unmanned Systems Test Bed **16**
Matt Burnham and Shannon Wigent

Partitioning and Archiving the Freight Desk Database at ONI **18**
Partha Mandayam, Dave Norman, Carl Holmberg, and Robert Dant

Asynchronous Parallel Optimization of Relative Magnetron Cathode for Counter Electronics **20**
P. J. Mardahl, M. T. Bettencourt, and K. L. Cartwright

Trade Wind Inversion Variability in Hawai'i **22**
Guangxia Cao and Duane Stevens

Modeling Heterogenous Patient Populations Using Gene Expression for Colorectal Cancer Prognosis **24**
Vidya Kamath, Lawrence O. Hall, Dmitry Goldgof, Rangachar Kasturi,
Timothy J. Yeatman, and Steven Eschrich

Seismic Soil-Structure Interaction Analysis of the Kealakaha Stream Bridge on Parallel Computers **26**
Seung Ha Lee and Si-Hwan Park

TABLE OF CONTENTS CONTINUED

UHF Track-Reactive Simulation **28**
Donald J. Fabozzi II

PCID and Aspire 2.0- The Next Generation AMOS Image Processing Environment **30**
Charles L. Matson, Charles C. Beckner, Jr., Kathy Borelli, Tom Soo Hoo, Shiho You, Brandoch Calef, Maria Murphy, and Ron Viloria

A Comparative Study on Boron and Olefin Metathesis **33**
Eluvathingal D. Jemmis, Susmita De, and Pattiyil Parameswaran

Theater UnderSea Warfare (TUSW) **34**
Carl Holmberg and Robert Dant

Index of Authors **37**

The Kernel Energy Method: Application to a Collagen Triple Helix

Lulu Huang, Lou Massa, and Jerome Karle

There is a rapid growth in computational difficulty with the number of atoms, when quantum mechanics is applied to the study of biological molecules. Two things alleviate this difficulty, viz., the advance of parallel supercomputers, and the use of a quantum crystallographic formalism based upon quantum kernels. The kernel methodology is well suited for parallel computation. Recently published articles have applied these advances to calculate the quantum mechanical *ab-initio* molecular energy of peptides, protein, DNA, and RNA. The results were found to have high accuracy. It is possible to use the full power of *ab-initio* quantum mechanics to calculate the interaction of long chain molecules of biological and medicinal interest, containing thousands of atoms. The calculations are simplified by representing a full molecule by smaller "kernels" of atoms. The general case is illustrated by a specific example using a triple helix collagen molecule of known molecular structure (Figure 1). The results (Table 1) show that such helix chain interactions are well represented by application of the KEM to this triple helix.[1]

Table 1. Interaction Energy Calculations* of Collagen Triple Helix (945 atoms and 9 kernels) by HF/STO-3G.

$E_{HF(abc)}$ [au]	$E_{HF(a+b+c)}$ [au]	$E_{KEM(abc)}$ [au]	$E_{KEM(a+b+c)}$ [au]	I_{HF} [kcal/mol]	I_{KEM} [kcal/mol]	I_{HF} -I_{KEM} [kcal/mol]
-22146.9171	-22146.8510	-22146.9112	-22146.8508	-41.4778	-37.9010	-3.5768

* Interaction Energies, $I_{abc}=E_{abc} - E_{a+b+c}$, $E_{a+b+c} = E_a+E_b+E_c$

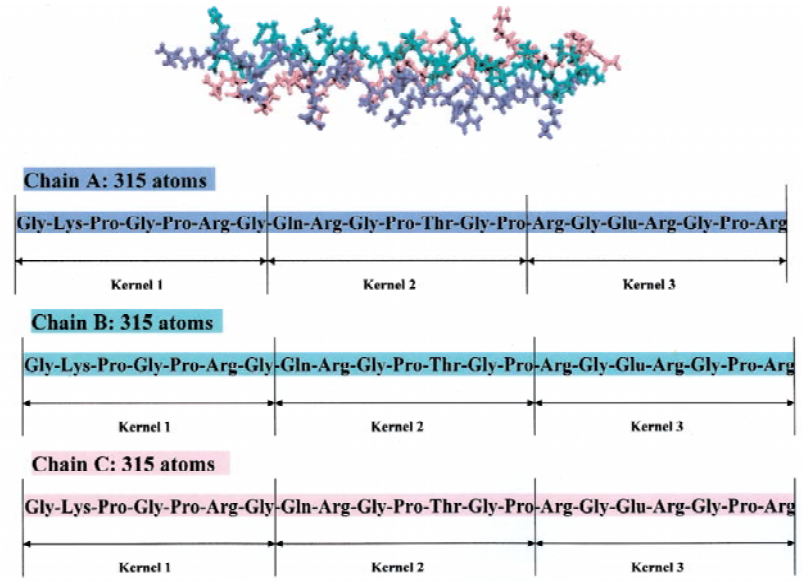

Figure 1. A picture of the collagen triple helix, 1A89, and the primary structure of each of its individual protein chains broken into kernels.

References:
1) Kernel Energy Method: The Interaction Energy of the Collagen Triple Helix, Huang, L.; Massa, L.; and Karle, J. Journal of Chemical Theory and Computation, Vol. 3, No. 4, 1337-1341 (2007).

Author and Contact: Lulu Huang
Organization: Scientist at the Laboratory for the Structure of Matter, Code 6030, Naval Research Laboratory, 4555 Overlook Ave., SW. Washington, DC, 20375-5341
Author: Lou Massa
Organization: Professor in Chemistry and Physics Departments at Hunter College and the Graduate School, City University of New York, New York, NY, 10021.
Author: Jerome Karle
Organization: Nobel Laureate and Chief Scientist at Laboratory for the Structure of Matter, Code 6030, Naval Research Laboratory, 4555 Overlook Ave., SW., Washington, DC, 20375-5341
Resources: IBM P3/P4 *Tempest* at MHPCC, and SGI Origin 3900, IRIX64 release 6.5, hpc11. asc.hpc.mil (ASC), and IBM p670, AIX version 5.2 (Hunter College)
Sponsorship: This project was supported by the Office of Naval Research. One of us (L. M.) wishes to thank the U. S. Navy Summer Faculty Research Program administered by the American Society of Engineering Education for the opportunity to spend summers at NRL. L. M. thanks NIH for grants (NIGMS MBRS SCORE5 S06GM606654, and RR-03037 the National Center For Research Resources) and NSF for CREST grant support.

High Performance Computing Software Applications for Space Situational Awareness

Francis Chun, Concetto Giuliano, Paul Schumacher, Charles Matson,
Kathy Borelli, Kevin Roe, and Bruce Duncan

The High Performance Computing Software Applications Institute for Space Situational Awareness (HSAI-SSA) has completed another full year of applications development. The emphasis of our work during this year was to continue improving space surveillance sensor model and image enhancement software. These applications are the Space Surveillance Network Analysis Model (SSNAM) and the Physically Constrained Iterative Deconvolution (PCID) image enhancement software tool. Specifically, we have demonstrated further speed-up in those codes running on the latest Cray XD-1 Linux supercomputer (*Hoku)* at the Maui High Performance Computing Center. The software applications improvements that the HSAI-SSA has made, has had significant impact to the warfighter and has fundamentally changed the role of high performance computing in SSA.

Introduction: The mission of the High Performance Computing Software Applications Institute for Space Situational Awareness (HSAI-SSA) is to support SSA needs of stakeholders by developing High Performance Computing (HPC) software applications for SSA. The Institute collaborates directly with experts and end users in the Joint Space Control mission area in order to identify those "leverage points" at which HPC can be applied to make a strategic difference in how the mission is performed. The HSAI-SSA has had a successful period of performance during the past year, working with multiple software applications, as briefly highlighted below.

Project Updates: HSAI-SSA work is currently primarily focused on two software applications, with one responding to Strategic Goal 1 (Astrodynamics) objectives, and one responding to Strategic Goal 2 (Image Enhancement) objectives. All applications now run on the Maui High Performance Computing Center (MHPCC) newest Cray XD-1 Linux supercluster named *Hoku* (or star in Hawaiian). The Space Surveillance Network Analysis Model (SSNAM) application is used by Air Force Space Command to analyze the performance of the global space tracking network, or SSN, and is being re-engineered to run as a scalable parallel process. The Physically Constrained Iterative Deconvolution (PCID) image enhancement software tool for space surveillance has also been further optimized for faster, more precise operations.

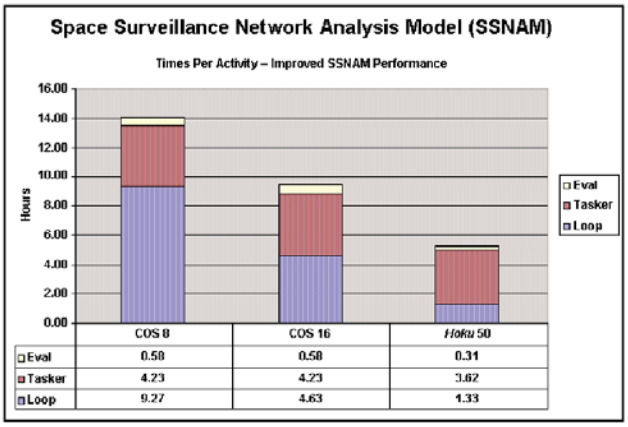

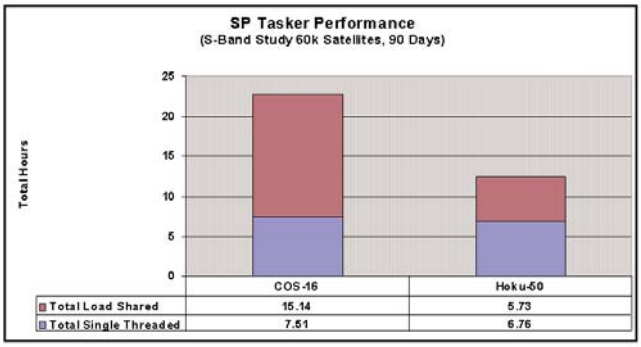

Figure 1. The panel on the left shows the six-fold speedup in the Loop module. The panel on the right shows that *Hoku*-50 runs of the SP Tasker take approximately 12 hours, which is 24x faster than what a non-load shared SP Tasker would take on a single PC.

During the past year, the original version of the Space Surveillance Network Analysis Model (SSNAM), which used the less accurate General Perturbations (GP) orbit propagation, was placed into operation at MHPCC. Also the SSNAM software architecture was updated and the speedup and parallelization of the Special Perturbations (SP) Tasker portion of this software for higher precision SP processing was completed. This code now executes twenty-four times faster compared to the serial SP Tasker (see the right panel in Figure 1). The HSAI-SSA conducted test runs of SSNAM with the parallelized SP Tasker at MHPCC using the *Hoku* Cluster (50 CPUs) and achieved more than a five-fold speed-up compared to runs using the Colorado Springs PC Cluster (16 CPUs) at the Air Force Space Command SSNAM laboratory.

PCID image processing software engineering work and successes during the past year have resulted in additional optimization of this code. An optimized version of the Numerical Recipes' conjugate gradient routine has been developed, implemented, and tested. The HSAI-SSA team has implemented the fastest 2-D distributed-memory Fast Fourier Transform (FFT) routine, optimal cost function-based regularization, and automatic parameter calculation - which have helped us achieve an additional speedup of more than 2.5-times versus the previously improved version of this code (see Figure 2).

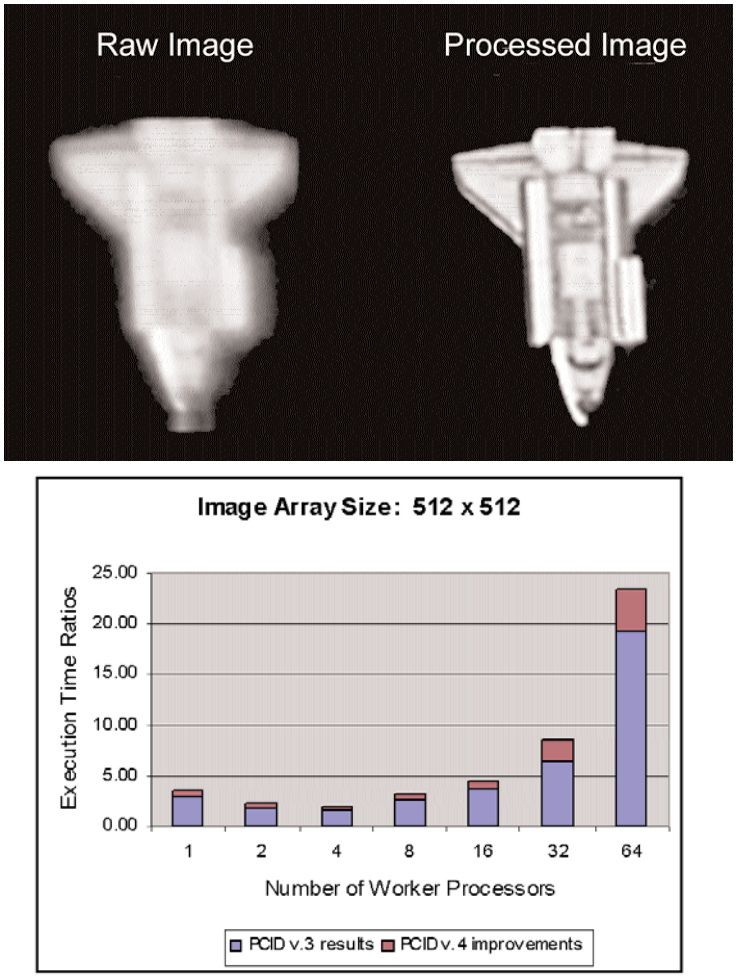

Figure 2. Hardware improvements by migrating from *Tempest* to *Hoku* achieved a > 10-fold speedup (lower panel). An additional speedup of at least 2.5-fold was obtained during the past year with software optimization modifications.

Conclusion: With sustained modest funding, the HSAI-SSA has continued making significant progress in developing HPC SSA applications for the warfighter. We've been able to build on the order of magnitude speed-up in space surveillance models and image enhancement software that was achieved during the past year. The HSAI-SSA team has employed proven software engineering practices to demonstrate scalability and performance improvements. HSAI-SSA is well focused and bringing the appropriate technical expertise together with the right computing resources to develop, test, and transition Department of Defense HPC software applications in the critical SSA arena.

Figure 3. HPC Software Applications Institute for Space Situational Awareness.

Author and Contact: Francis Chun
Authors: Concetto Giuliano and Paul Schumacher
Organization: Air Force Research Laboratory/DESM, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Author: Charles Matson
Organization: Air Force Research Laboratory/DESA, Kirtland AFB, NM, 87117-5776
Author: Kathy Borelli
Organization: KJS Consulting, 71 Awalau Road, Haiku, HI
Authors: Kevin Roe and Bruce Duncan
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Resources: Cray XD-1 Linux Supercluster *Hoku* and IBM P3 Servers *Tempest* at MHPCC
Sponsorship: This work is sponsored by the Department of Defense High Performance Computing Modernization Program and the Air Force Research Laboratory's Directed Energy Directorate.

Tracking Humpback Whales with Passive Acoustics

(2004 Engagement Grant Recipient)

Eva-Marie Nosal and L. Neil Frazer

Since many marine mammals are very vocal, passive acoustic techniques present a highly promising approach to marine mammal studies, monitoring, and human impact mitigation. They can be used as an addition/alternative to more traditional visual and tagging techniques, which may be costly, invasive, and limited to favorable environmental considerations (such as sunlight and calm seas). Advantages of passive acoustic methods include relative cost-efficiency, non-invasiveness, and potential for long-term monitoring.

To track marine mammals using passive acoustics, vocalizations are recorded on several hydrophones (an array) towed behind a boat, mounted on the ocean-bottom, or suspended from buoys near the surface. These recordings are processed to find the position of the animal at incremental time steps. Our work involves the development and testing of the processing methods for tracking humpback whales.

Methods: The processing methods used depends on several factors, including: (1) the type of call being produced (including the duration of the call, its frequency range and characteristics - Figure1 shows a humpback whale signal); (2) environmental parameters (including bottom depth and sound-speed profile); and (3) the type of hydrophone array. Several techniques exist for localizing underwater sound sources, such as the "time difference of arrival" (TOAD) method and "matched-field" (MF) methods. There are several reasons that these methods are not always successful for tracking humpback whales. In TOAD methods, the longer-duration calls of humpbacks, the fact that they favor shallow waters, and the presence of multiple singers makes it difficult to identify and associate direct reflections. MF methods rely on accurate models of sound propagation (only possible for low frequencies), long line-arrays (costly), and knowledge of the source signature (not available for humpbacks).

To track humpbacks using arrays with a minimal number of hydrophones, we extended conventional MF techniques to get our "pair-wise spectrogram" (PWS) method. In this method, recordings are processed along pairs of receivers. The signal (data) at the first hydrophone in the pair is propagated (via convolution with a modeled Green's function) from the second hydrophone position to a candidate source location. This gives a resulting signal, H_{12}. We get H_{21} similarly. If the candidate source is at the correct source location, these two resulting signal waveforms should be identical. However, because we cannot model the environment perfectly and because there is noise in the data, they may differ significantly. Since spectrograms are less sensitive (than full waveforms) to noise and modeling uncertainties, we compare the spectrograms (rather than the waveforms) of H_{12} and H_{21} (by taking normalized inner products over frequency, time, and receiver pairs). This procedure is repeated for a grid of candidate source locations, and the one that gives the best agreement between spectrograms is the estimated source location.

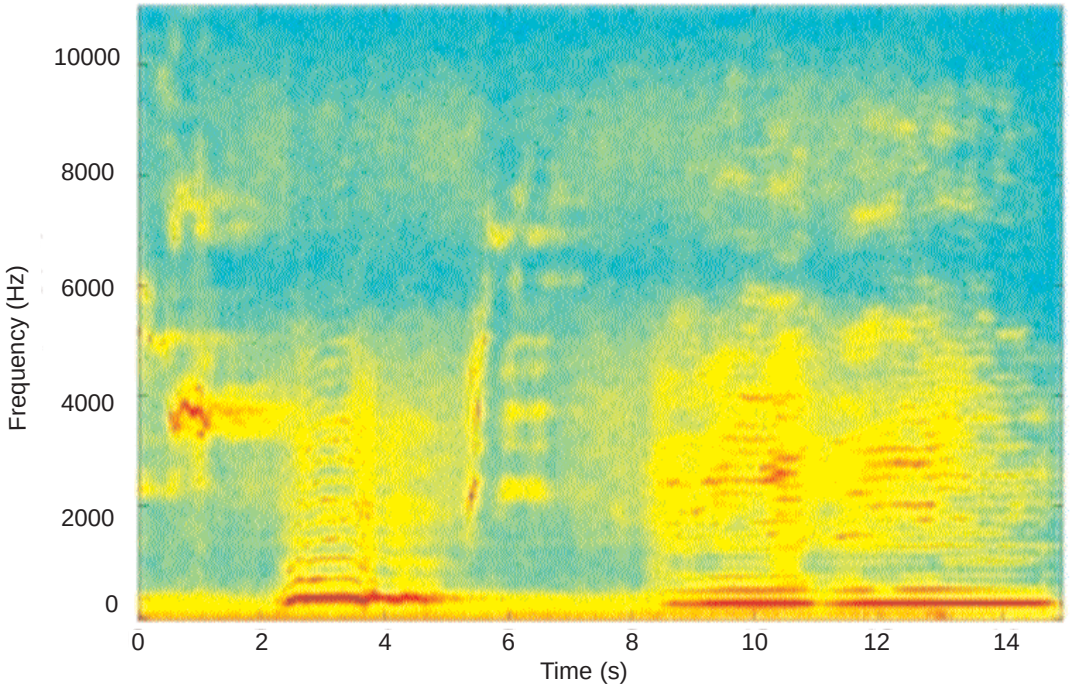

Figure 1. Spectrogram of a humpback whale signal.

Use of HPC Resources: Computational demands associated with the prediction of Green's functions for each receiver/candidate source location pair increase with the product of frequency and range. Since our localization uses frequencies up to several kHz and we wish to localize at ranges of up to tens of kilometers, high-performance computing (HPC) resources are required to use the PWS method. Fortunately, the computations of Green's functions at different locations are independent, which means they parallelize well. Further computational demands are associated with spectrogram computation and inner products. Once again, since computations for difference candidate source locations are independent, this step parallelizes well. Once our algorithms were developed and implemented on a partial scale (fewer frequencies, shorter ranges) using desktop PCs, they were parallelized and ported to the *Squall* system at MHPCC for full scale testing and simulations.

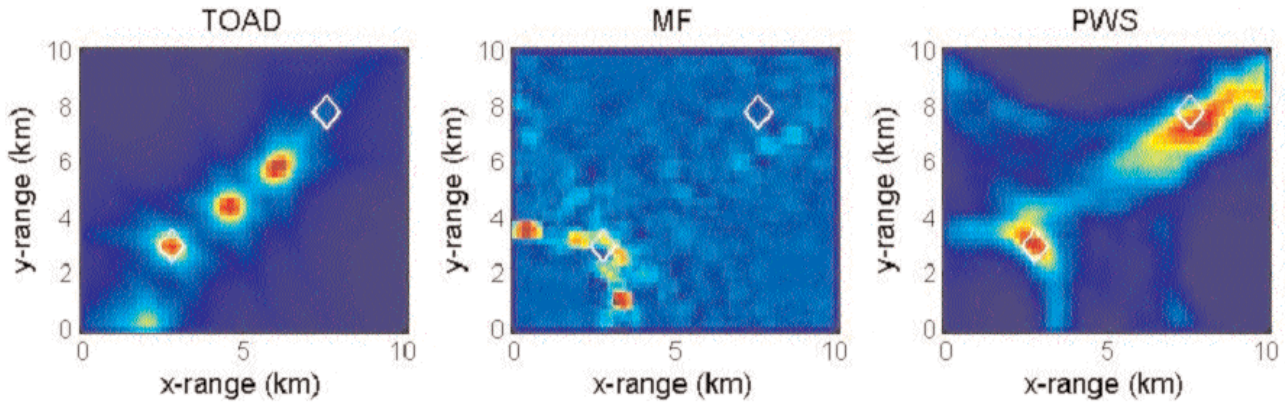

Figure 2. Ambiguity surfaces (probabilistic indicators of source location with red indicating highest probability) in plan-view using the TOAD, MF, and PWS processing methods. Three receivers were used to localize two sources (indicated with white triangles) in 0 dB signal-to-noise ratio. Environmental uncertainty is in the form of incorrect bottom depth: 200 m for the forward model, 204 m for the inversion. Only PWS finds both sources correctly.

Results and Future Work: We developed and tested PWS processing on simulated data.[1] One of our more recent simulations is shown in Figure 1. The PWS method clearly outperforms both the TOAD and MF method in localizing both sources. We are currently collecting data of singing humpback in Hawaiian waters that will be used to test PWS in a real life situation. Processing of this data will require HPC resources for which we plan to continue to use the *Squall* system, and possibly the *Tempest* system, at MHPCC. We are also working on a variety of different processing techniques for localizing other species, such as sperm whales,[2] blue whales, and spinner dolphins.

References:

1) Nosal, E.-M., and Frazer, L. N., "Pair-wise processing of spectrograms for localization of multiple broadband CW sources," Proceeding of the IEEE Oceans '05 Europe meeting, Brest, France, June 20-23, 2005. Special publication in the Newsletter of the IEEE OES, Winter 2006.
2) Nosal, E.-M., and Frazer, L. N., "Delays between direct and reflected arrivals used to track a single sperm whale," Applied Acoustics, 87 (11-12), 1187-1201, 2006.

Author and Contact: Eva-Marie Nosal
Author: L. Neil Frazer
Organization: Department of Geology and Geophysics, School of Ocean and Earth Science and Technology, University of Hawaii at Manoa, 1680 East-West Road, POST 813, Honolulu, HI, 96822
Resources: IBM P3 Server *Squall* at MHPCC
Acknowledgement: Funding was provided by the University of Hawaii through a MHPCC Student Engagement Grant. HPC resources were provided by MHPCC.

High-Resolution Forecasts to Support AMOS Using WRF

Kevin P. Roe

The Hawaiian Islands contain a variety of microclimates in a very small region. Some islands have rainforests within a few miles of deserts; some have 10,000 feet summits only a few miles away from the coastline. Because of this, weather models must be run at a much finer resolution to accurately predict in these regions. The Weather Research and Forecasting (WRF)[1,2,3] modeling system (version 2.1.2) is run from a coarse 54 km horizontal resolution (surrounding an area of approximately 7000 by 7000 km) nested down to a 2 km horizontal resolution (and 55 vertical levels) daily. Since the computational requirements are high to accomplish this in a reasonable time frame (as to still be a forecast) WRF is run in parallel on MHPCC's Cray Opteron cluster. Utilizing 32 nodes (2 processors/node) the WRF model is run daily over the above conditions in under 4 hours for a 48 hours simulation. Although these forecasts are relatively new, over seven years of numerical weather simulation experience with MM5 and RSM have contributed to its effective use. Operators at the telescope on Haleakala, Maui, utilize this WRF simulation in their daily planning.

Research Objective: The telescope operations on Haleakala are highly dependent on weather conditions on the Hawaiian Island of Maui. If the wind speed is too high then the telescope cannot be utilized. Problems also exist if there are clouds overhead. Rainfall and relative humidity are also factors in determining the capabilities of the telescopes. Lastly, optical turbulence (or "seeing") is predicted over Haleakala using the output conditions from WRF[4].

In order to effectively schedule telescope operations, an accurate weather prediction is extremely valuable. Current forecasts that are available from the National Weather Service (NWS) give good indications of approaching storm fronts but only at a coarser level (10-12 km resolution). Because of this and the location of the telescope on Maui, this can be insufficient for their needs. The additional benefit of the telescope operators having access to an accurate forecast (even for only a day in advance) is that they can still perform some scheduling. If a storm is predicted they can plan maintenance for this time period. This allows them to function more effectively by giving them the capability to schedule downtime. This in turn saves time, improves operating efficiency, and potentially saves money.

Daily Operations: Every night at Midnight Hawaiian Standard Time (HST), a PERL script is run to handle all the operations necessary to produce a forecast on MHPCC supercomputers, prepare images of weather fields (wind, temperature, relative humidity, rainfall, cn_2, etc.), and post it to the MHPCC web page (http://weather.mhpcc.edu).

Figure 1. Haleakala Weather Center Homepage.

Web Output: Now that the above processes have created images, they must be made available for the telescope operators. This is accomplished by posting the images to the MHPCC web page; specifically, http://weather.mhpcc.edu. This title page gives the user the option of what area and resolution they would like to examine. From the title page, the user can select the all island area at a 54, 18, or 6 km resolution or 1 of the 4 counties (Hawaii, Maui, Oahu, and Kauai) at a 2 km resolution. Once one of the above has been selected, the user is transported to a web page that initially includes an image of the wind in the selected area. On this regional web page, the viewer can select to see the previous or next image through the use of a small JavaScript. If the viewer prefers, an animation of the images (in 1 hour increments) can be started and stopped. Finally, the user can select any of the other images from a pull down menu. If the viewer would like to change the field being examined, a pull down menu on the left side of the page will transport the user back to the main menu, a different county, or allow them to choose a different weather field.

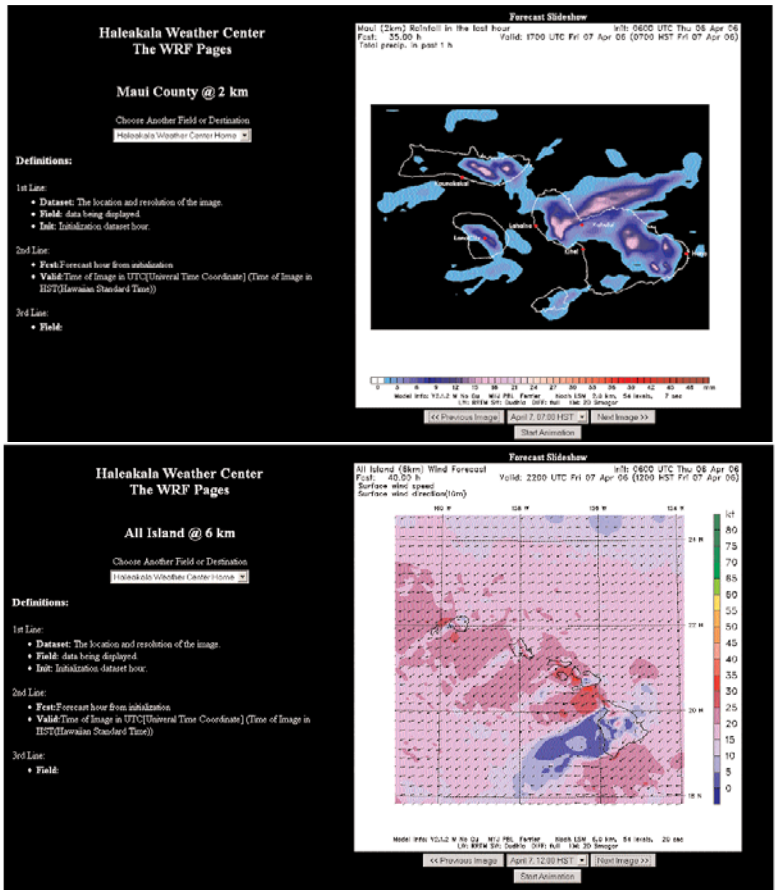

Figure 2. Examples of weather conditions such as rainfall and surface wind speed and direction.

References:

1) Michalakes, J., J. Dudhia, D. Gill, T. Henderson, J. Klemp, W. Skamarock, and W. Wang: The Weather Research and Forecast Model: Software Architecture and Performance. Proceedings of the Eleventh ECMWF Workshop on the Use of High Performance Computing in Meteorology. Eds. Walter Zwieflhofer and George Mozdzynski. World Scientific, 2005, pp. 156-168.
2) Michalakes, J., S. Chen, J. Dudhia, L. Hart, J. Klemp, J. Middlecoff, and W. Skamarock: Development of a Next Generation Regional Weather Research and Forecast Model. Developments in Teracomputing: Proceedings of the Ninth ECMWF Workshop on the Use of High Performance Computing in Meteorology. Eds. Walter Zwieflhofer and Norbert Kreitz. World Scientific, 2001, pp. 269-276.
3) Skamarock, W. C., J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, and J. G. Powers, 2005: A Description of the Advanced Research WRF Version 2. NCAR Tech Notes-468+STR.
4) Ruggiero, F., Roe, K. P., DeBenedicts, D. A. "Comparison of WRF versus MM5 for Optical Turbulence Prediction." DoD High Performance Computing Modernization Program User Group Conference. Williamsburg, VA, 2004.

Author and Contact: Kevin P. Roe
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Hi, 96753

Resources: Cray XD-1 at MHPCC
Sponsorship: Air Force Research Laboratory
Acknowledgement: MHPCC would like to thank the Maui Space Surveillance System for their cooperation and feedback on this project.

Plasma Actuator Modelling

Gabriel I. Font-Rodriguez, Connor Caples, Shawn Hackett, and Cordell Hachinsky

Plasma actuators are electrical devices that produce a plasma (ionized gas) over an airfoil for the purpose of minimizing separated flow which results in increased lift and reduced drag. These actuators hold the promise of improving the efficiency and performance of aircraft through the use of an inexpensive electrical device. The present study utilizes Particle-In-Cell and Direct-Simulation-Monte-Carlo methods in an effort to simulate the plasma formation and learn about the interaction of the plasma with the air. Current findings address the applicability of the plasma actuator for high altitude flight as well as design strategies for improving the actuator performance. The study is carried out by faculty of the U.S. Air Force Academy and cadets of the U.S. Military and Air Force academies.

Research Objectives: The goal of this research is to increase the present understanding of the physics behind the formation of the plasma and its interaction with the air in order to improve actuator performance to a degree that it becomes applicable to current flight vehicles.

Background: Aerodynamic flow control is typically achieved through the use of pneumatic means, such as blowing or suction, or through the use of physical means, such as vortex generators. In recent years, another method has become available which utilizes an atmospheric pressure plasma discharge. Plasma actuators have been successfully employed to promote boundary layer attachment on airfoils at high angle-of-attack. Experimental studies have shown that the plasma actuator supplies momentum to the boundary layer. The manner in which this momentum was imparted has been the subject of much study. This research attempts to improve the physical understanding by simulating the plasma actuator using particle methods (Particle-In-Cell and Direct-Simulation-Monte-Carlo).

The plasma actuator is a simple device consisting of two electrodes separated by a dielectric and staggered in the flow direction, as shown in Figure 1. The upper electrode is exposed to the free stream and typically subjected to a bias of several thousand volts at a frequency of several to tens of kilo-Hertz. The buried electrode is usually electrically grounded. When the actuator is turned on, a plasma is observed to form downstream of the exposed electrode and impart a force to the air in the direction of the buried electrode. The power requirements of this device are measured in only tens of Watts, while the potential for improving the efficiency of aircraft and saving significant amounts on their operating cost is worthy of consideration.

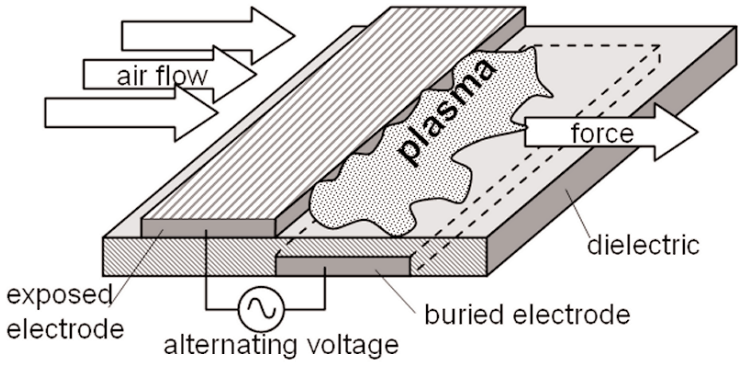

Figure 1. Schematic of plasma actuator.

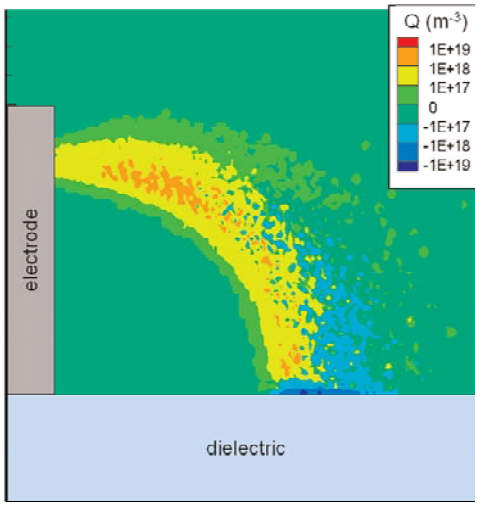

Figure 2. Plasma density contours for the actuator during the beginning of the discharge.

Results: An illustration of the plasma actuator discharge is shown in Figure 2, where the charge density is shown. An electron avalanche from the exposed electrode to the buried electrode has populated the region with electrons, positive ions, and negative ions. Simulations carried out by faculty and cadets reveal several new facets of the physics that governs the plasma actuator behavior. The first one deals with the plasma chemistry. The plasma discharge is governed by several hundred reactions involving ionization, excitation, dissociation, momentum transfer, and attachment. The set has been reduced down to less than twenty reactions in order to minimize computational cost while still trying to include the most relevant and important physics.

The results of the two chemistry sets are compared to experiments in Figure 3. The second chemistry set attempts to capture the energy loss mechanisms of highly accelerated electrons. While the force-bias behavior is better for the second set, clearly more work is still needed.

Inspection of the momentum transfer mechanism revealed another aspect of the plasma interaction with the air. After comparing momentum transfer totals for electrons, positive ions, and negative ions, it was found that the bulk of the force is imparted to the air through ion-neutral collisions by the positive ions. This is illustrated in Figure 4. This data was collected from summing the momentum transfer of millions of particles during different simulations at many plasma conditions.

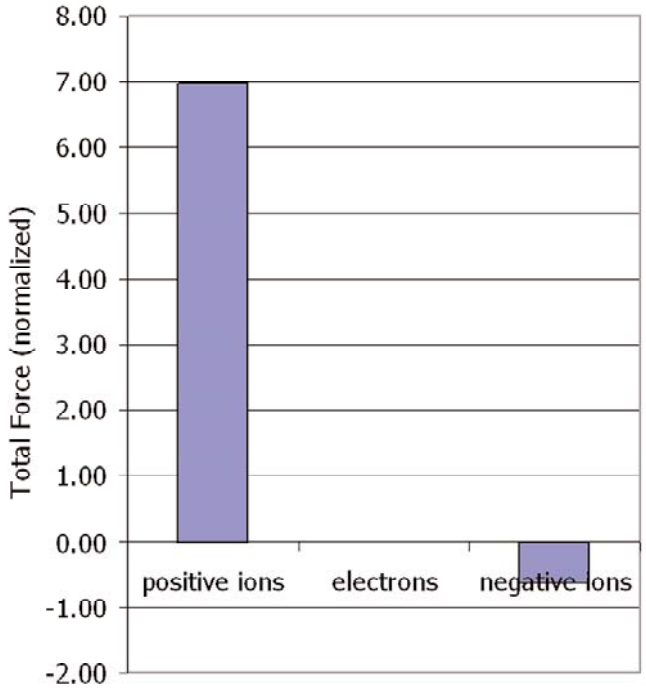

Figure 4. Total force imparted by the different charged species in the plasma during the entire bias cycle.

Significance: This study explored the physics behind plasma actuators, an emerging flow control technology. These actuators hold the promise of increasing the efficiency and performance of aircraft. The current study helps to increase the understanding of how the plasma is formed and interacts with the air flow. This will help in the design of more effective actuators and speed their eventual installment on aircraft.

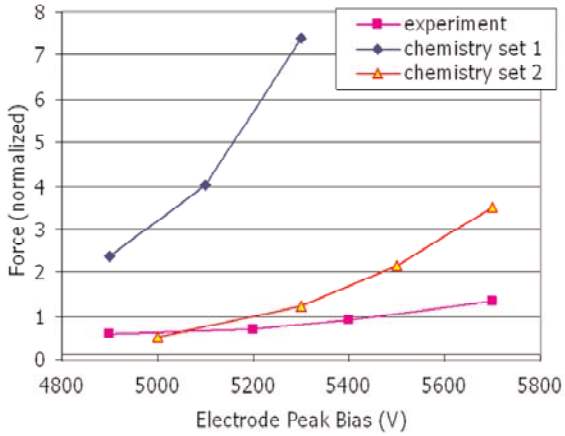

Figure 3. Comparison of normalized force imparted by the actuator and experiments for two sample chemical reaction sets.

The behavior of the plasma actuator at high altitude is another topic of interest. Computations were carried out with diminishing background air densities in order to simulate the effect of increasing altitude. The results reveal that the actuator force increases with altitude due to higher plasma densities caused by high energy electrons, as shown in Figure 5. The power required, however, also increases implying a loss in efficiency.

Finally, geometric aspects of the actuator design were varied in order to determine their effects on actuator efficiency. Changes in several parameters produced force gains of more than an order of magnitude. This continues to be investigated and will now be attempted experimentally.

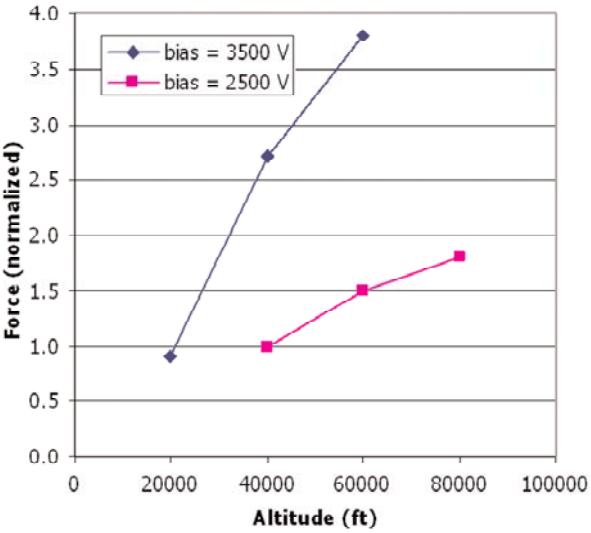

Figure 5. Effect of altitude on force imparted by plasma actuator.

Author and Contact: Gabriel I. Font-Rodriguez
Authors: Cadet Connor Caples and Cadet Shawn Hackett
Organization: USAF Academy, Department of Physics, 2354 Fairchild Dr., Suite 2A169, CO, 80840
Author: Cadet Cordell Hachinsky
Organization: U.S. Military Academy, West Point, NY, 10996
Resources: Initial Development was at USAFA. Computational parallelization development and computational experiments were executed on *Jaws* at MHPCC.
Sponsorship: This research was sponsored by the DOD High Performance Computing Modernization Program (HPCMP).

Mantle Convection, Melting, and Hotspot Geochemistry

Todd Anthony Bianco, Garrett Ito, Jeroen van Hunen,
Maxim Ballmer, and John Mahoney

This work studies the cause of observed geographic variations in lava compositions at hotspot volcanoes. For example, at Hawaii, key geochemical tracers of mantle composition (e.g., Sr, Nd, and Pb isotope ratios) vary with age and between different volcanoes. Such observations indicate that the underlying mantle is geochemically heterogeneous, but geoscientists struggle to understand how and over what spatial scales this heterogeneity occurs in the mantle. This study simulates the 3-D dynamics of a hot plume of mantle rising beneath the tectonic plate to generate hotspot volcanoes. The key advancement is to simulate the transport, heat transfer, melting, and mixing of different geochemical components in an internally consistent numerical model. We are finding that the spatial variability in melting alone can contribute to some of the key observations of lava composition in Hawaii. These results suggest that mantle heterogeneity on the small scale (10^1-10^3 m) is as, or more significant than at the large scale (10^5-10^6 m). This finding represents a benchmark for understanding how mantle heterogeneity is created and evolves through geologic time.

Description: Hotspot volcanoes are observed throughout the earth's surface, and are marked by globally as well as locally varying composition and eruption rate. A popular way to explain the cause of hotspot volcanism is with the mantle plume hypothesis, by which a hot and buoyant plume of rock in the mantle rises below the tectonic plate. If the excess temperature of the rising rock plume is high enough, the rock melts to form magma that percolates through the plate and erupts on the surface. The plume hypothesis explains many observations at hotspots, such as the well-known linear age progression of Hawaiian volcanoes. A challenge in invoking the mantle plume hypothesis is to explain the observed geographic variations of lava composition at hotspots. For example, Hawaiian volcanoes erupt lavas with Sr, Nd, and Pb isotope ratios (geochemical tracers of mantle composition) that vary systematically with volcano age and position. Such observations are clues to the chemical structure and evolution of the mantle. Our work examines how upper mantle flow and melting, of chemically heterogeneous mantle contributes to geographic variations in volcano compositions at hotspots. We used the MHPCC facilities to numerically simulate the full 3-D dynamics of a mantle plume as it melts and interacts with the overlying, rigid lithospheric plate. Such calculations are needed to relate the 3-D mantle dynamics to surface magma composition.

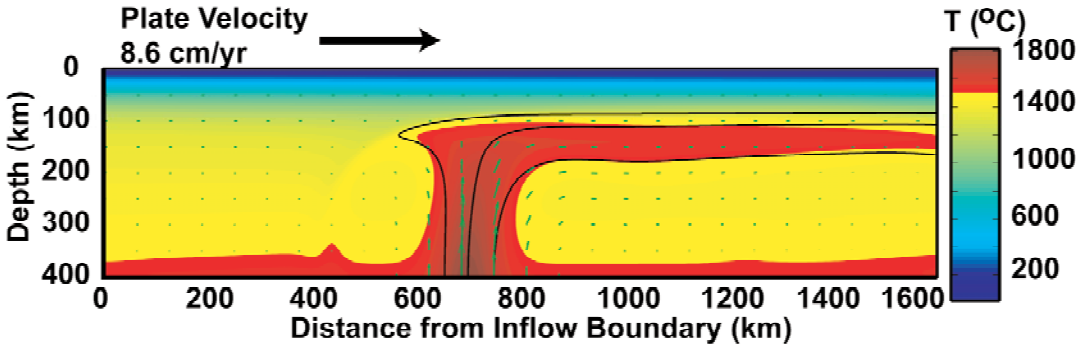

Figure 1. A profile of the temperature and flow fields at the plume axis (y=0). Vertical axis is depth and horizontal axis is distance from the inflow boundary. Colors show temperature in °C, arrows are velocity vectors, and black lines are streamlines (tracer trajectories). Full box dimensions are 400 km deep, 1600 km long, and 800 km wide (not shown) with 5 km resolution in depth and 6.25 km resolution in horizontal dimensions. A 65 km radius temperature anomaly is prescribed at (x,z,y) = (700,400,0) km. A no-slip velocity condition is set on the top boundary to simulate 8.6 cm/yr plate velocity in the positive x-direction. Upwelling rate and temperature are highest at the center of the plume, and these conditions control where and at what rates different geochemical components are sampled by melting.

The project requires performing two sets of coupled calculations: one modeling upper mantle convection and one modeling composition and partial melting. To simulate upper mantle convection, we employ CITCOM, a finite element code that numerically solves equations of conservation of mass, momentum and energy for an incompressible fluid, with zero Reynold's number, high Rayleigh number (10^5-10^6) and with large viscosity variations (10^{17}-10^{22} Pa s), all of which are conditions expected for the ductile creep of Earth's upper mantle *[Moresi and Solomatov, 1995[1]; Moresi and Gurnis, 1996[2]; Zhong et al.[3], 2000; Zhong and Watts, 2002[4]].* Figure 1 shows an example calculation (2-D cross-section) of the temperature and velocity field. To simulate plate motion, a horizontal velocity is imposed on the top of the model mantle, seen here as moving to the right. Also, to initiate a mantle plume, we impose a thermal anomaly at the base of the model. Melting is simulated using experimentally constrained parameterizations in which the fraction of partial melt depends on temperature, pressure, and composition. The amount of melt produced at a particular position in the mantle is controlled by the material flux (therefore velocity) through that position. The extent and rate of partial melting are next used to predict magma fluxes and concentrations of the key geochemical tracers of mantle heterogeneity.

A key assumption we make is that heterogeneity in mantle composition is present on scales (10^1-10^3 m) that are much smaller than the zone of partial melting. To predict compositions at the surface, we assume magma rises vertically, and thus we average compositions in each column of model elements. To average the column properly, the compositions are weighted by the amount of liquid actually produced at a given location, in other words, the rate of melting at the location. Figure 2 shows the predicted compositions from the same calculation shown in Figure 1. In this figure, the viewer is looking down on model seafloor, and colored contours indicate concentration of "EC", which is our name for one geochemical (Sr and Nd isotope) component in the mantle. EC can begin melting at much lower temperatures than "DC" (the other geochemical component with distinct Sr and Nd composition), and therefore contribution from DC is only seen near the center of the hotspot, where temperature is highest.

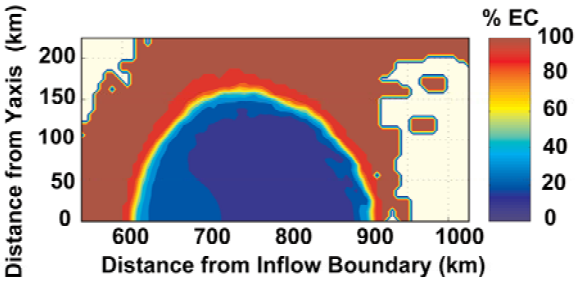

Figure 2. Fraction of EC contribution to magma composition, in mapview. Vertical axis is distance from the plume axis (the y-dimension; normal to plate motion); horizontal axis is distance from the inflow boundary; thus the viewer is looking down onto the top of the model box. Color contours are percentages of the "EC" geochemical component in the predicted surface magma. Compositions form a nearly radial pattern, with the strongest contributions from the other component, "DC", occurring near the center of the pattern. This figure shows that the model predicts significant and systematic variations in magma composition at length scales much larger than that assumed for the starting mantle.

References:

1) Moresi, L.-N. and V. S. Solomatov, Numerical investigation of 2D convection with extremely large viscosity variations. Phys. Fluids 7 (9): 2154-2162. 1995.
2) Moresi, L. N. and M. Gurnis, Constraints on the lateral strength of slabs from three dimensional dynamic flow models. Earth and Planetary Science Letters 138. 15-28. 1996.
3) Zhong, S., M. T. Zuber, L. N. Moresi and M. Gurnis, The role of temperature dependent viscosity and surface plates in spherical shell models of mantle convection. J. Geophys. Res. 105 11,063-11,082. 2000.
4) Zhong, S. J. and A. B. Watts, Constraints on the dynamics of mantle plumes from uplift of the Hawaiian Islands. Earth Planet. Sci. Lett. 203 (1): 105-116. 2002.

The pattern of varying composition in Figure 2 means that as a volcano moves over the hotspot (from right to left) it will sample different compositions. Figure 3 shows the average composition a volcano would sample as it moves over the hotspot. The solid black line is the volcano composition, and the dashed black line is volcano volume. As the volume increases the composition begins as being mostly EC and evolves to lower EC content, and as the volume increase slows, the composition switches back to EC-like. Figures 2 and 3 illustrate an important result: the dynamics of upper mantle convection and melting can transfer heterogeneity on small scales (10^1-10^3 m) in the mantle into to larger scales variations in magma at the surface. Previous studies have suggested that heterogeneity in the mantle is strongest over scales of (10^5-10^6 m), but our results reveal that the heterogeneity at the small scale can be as or more important.

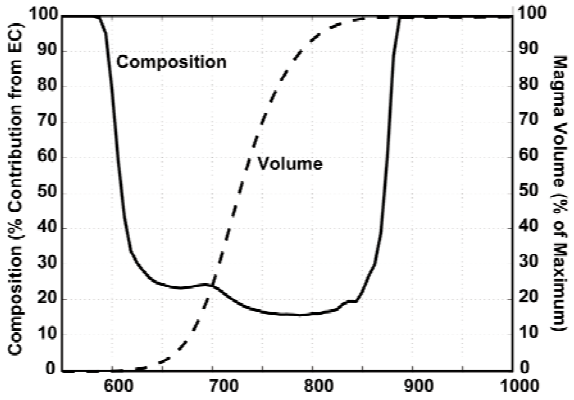

Figure 3. Fraction of EC contribution to magma composition and volcano growth. Predicted properties of a single volcano as it passes over the melting zone in Figure 2 from right to left. The horizontal axis is again distance from inflow boundary as in Figures 1 and 2. The solid black line is magma composition, shown as percentage of the EC geochemical component on the left vertical axis. The dashed black line is the volume of magma accumulated, shown as percent of total magma volume accumulated on the right vertical axis. This figure shows that as a volcano grows with time, its composition will vary systematically and in a manner that resembles some of the key variations at hotspots such as Hawaii.

Author and Contact: Todd Anthony Bianco
Authors: Garrett Ito and John Mahoney
Organization: Department of Geology and Geophysics, SOEST, University of Hawaii at Manoa, POST 701,1680 East-West Road, Honolulu, HI, 96822
Author: Jeroen van Hunen
Organization: Department of Earth Sciences, Durham University, Durham DH1 3LE, Durham, UK
Author: Maxim Ballmer
Organization Institut f. Geophysik, ETH Zurich, Zurich, Switzerland

Resources: *Akuatest* Cluster (Intel Xeon at SOEST, UH Manoa) and the *Hurricane System* (IBM SP P4) at MHPCC
Acknowledgements: The work completed on MHPCC resources was greatly expedited by the help of Mary Ann Ciuffini and Michele Hershey of MHPCC.
Sponsor: 2007 MHPCC Student Engagement Grant

The Electronic Mechanism of Nickel Titanium Shape Memory Behavior

Nicholas B. Hatcher, Oleg Yu. Kontsevoi, and Arthur J. Freeman

NiTi exhibits the unique property of undergoing a diffusionless structural transformation near room temperature. While this material has been widely studied and utilized in a large range of engineering applications including the development of "smart" wings for aircraft and prosthetic tendons, the mechanism causing this atomic reordering from a cubic to monoclinic structure is still unknown. We demonstrate that the mechanism is related to instabilities resulting from the so-called "nesting" of the electronic Fermi surface causing the transition from the high temperature CsCl structure to the low temperature monoclinic structure. This detailed explanation of the electronic mechanism furthers researchers' ability to predict, categorize, and understand martensitic transformations.

Introduction: Martensites, materials that exhibit a structural transformation in response to heat and pressure variations, have been studied for decades. NiTi has generated considerable interest for its reversible martensitic transformation near room-temperature. This alloy transforms from a high temperature cubic CsCl type (B2) structure into a low temperature monoclinic (B19') structure, thus changing its shape with changes in temperature and pressure. Due to its shape memory effect, nitinol may be used in future flapless "smart" wings, prosthetic tendons, or more flexible surgical tools all of which can be activated to change form with slight changes of heat or pressure. While experimental and theoretical work on this material is robust, the explanation of the electronic mechanism causing the martensitic transformation is still not complete. We seek to explain this phenomenon and further the understanding of the mechanism of the NiTi shape memory behavior

Results and Significance: Calculations were performed using density functional theory with our highly precise full-potential linearized augmented plane wave (FLAPW)[1] method. We determined electronic structures, total energies, and stacking fault and elastic properties for the intermediate, high temperature, and low temperature structures of NiTi. The elastic constants were calculated for the B2 structure and compared to other NiTi structures. The C' and C_{44} values of the B2 structure were found to be much lower than other intermediate and low temperature structures, which suggests an instability of the B2 to crystal deformation. Additionally, gamma surfaces,[2] i.e., calculations of differences in the total energy with respect to generalized stacking faults, in the {100} and {110} planes were analyzed. In our search for structural instabilities of the B2 structure, we find a low energy barrier path in the <100> direction of the {110} plane. The energy surface of this stacking fault calculation is shown in Figure 1.

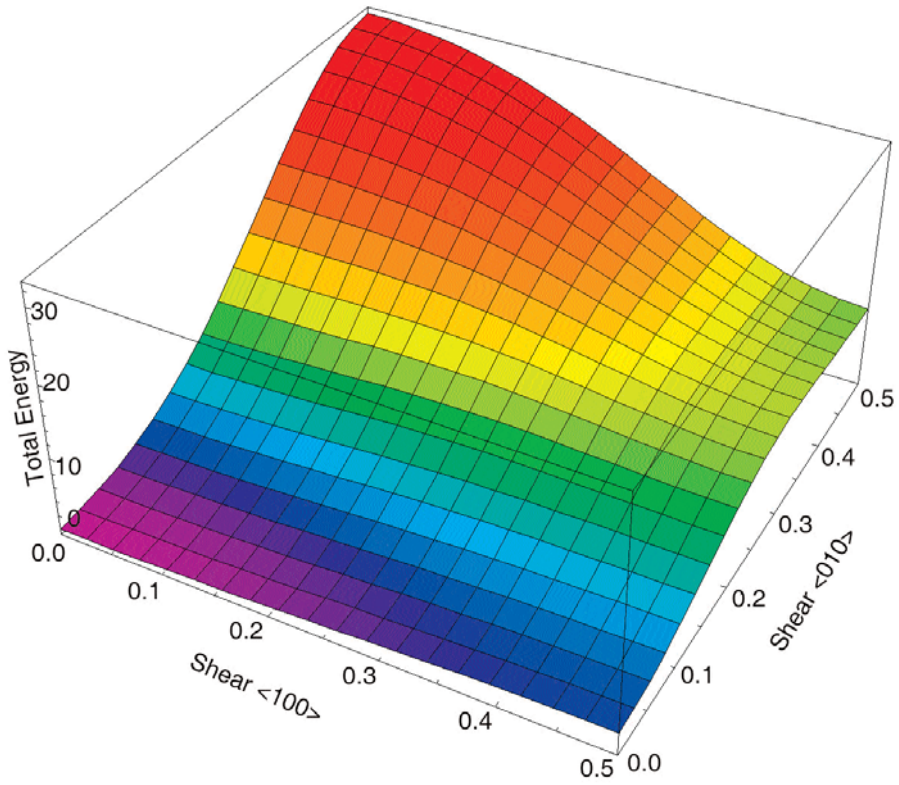

Figure 1. Above is a figure of the total energy of the structure verses stacking fault for the {110} plane. The vertical axis represents the energy difference with respect to the B2 structure in J/m^2. The horizontal axes show stacking faults amounts in the <100> and <010> directions.

We searched for instabilities in the electronic structure to determine the mechanism of this transformation. We focused on the Fermi surface and investigated plane parallel surfaces that may show so called Fermi surface "nesting". Nesting typically denotes instabilities in the electronic structure, and can predict charge density wave formation (i.e., a sinusoidal type distortion of the charge density with respect to the lattice). Charge density waves are often a precursor to structural deformations, such as martensitic transitions. Therefore, Fermi surface nesting can be indicative of structural instabilities in a material.

The Fermi surface of B2 NiTi is shown in Figure 2. Note how the distorted hemisphere of band 8 can be translated to match the surface of band 8 in the <110> and <111> directions. These Fermi nesting regions were compared with experimental data. In addition, neutron diffraction experiments have shown absorption phenomena that suggest charge density wave formation.[3,4] The neutron absorption patterns match our calculated nesting vectors of both the <110> and <111> directions and magnitudes and signify instabilities to the monoclinic and tetragonal distortions.

In summary, in this study of NiTi we have calculated structural and electronic properties from first principles. Our calculations of structural properties show agreement with experiment. The calculated elastic constants and shear energetics exhibit precursory behavior to the martensitic transformation. Most interestingly, the exhibition of an electronic instability is shown as Fermi surface nesting in B2 NiTi. It is this electronic instability that suggests charge density wave formation and leads to the structural transformation of NiTi. This explanation of the electronic mechanism furthers the fundamental understanding of how shape memory NiTi undergoes this transformation.

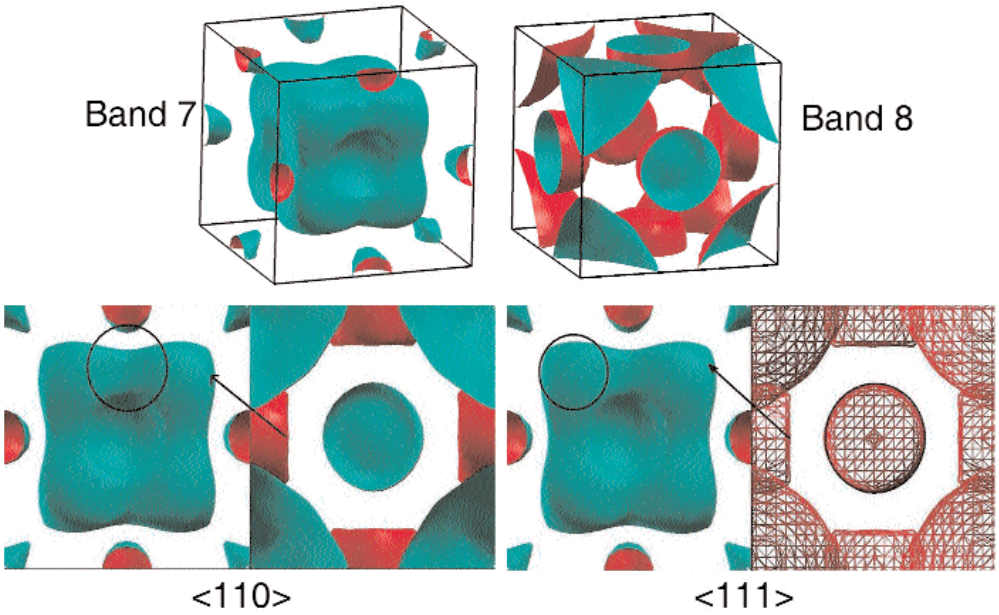

Figure 2. This shows the Fermi surface of B2 NiTi. Vectors represent the <110> and <111> nesting vectors that map band 7 to nest with band 8.

References:

1) Full-potential self-consistent linearized-augmented-plane-wave method for calculating the electronic structure of molecules and surfaces: O_2 molecule; E. Wimmer, H. Krakauer, M. Weinert, and A. J. Freeman; Phys. Rev. B; 1981.
2) First-principles total-energy calculations for planar shear and cleavage decohesion processes in B2-ordered NiAl and FeAl; N. I. Medvedeva, O. N. Mryasov, Yu. N. Gornostyrev, D. L. Novikov, and A. J. Freeman; Phys. Rev. B; 1996; p. 13506.
3) Observation of a soft-phonon mode and a pre-martensitic phas in the intermetallic compound $Ti_{50}Ni_{47}Fe_3$ studied by inelastic neutron scattering; P. Moine, J. Allain, and B. Renker; J. Phys. F.: Met. Phys.; 1984; pp. 2517-2523.
4) Dynamical properties of premartensitic NiTi; H. Tietze, M. Müllner, and B. Renker; J. Phys. C: Solid State Phys; 1984; pp.L529-L532.

Author and Contact: Nicholas B. Hatcher
Authors: Oleg Yu. Kontsevoi and Arthur J. Freeman
Organization: Department of Physics and Astronomy and the Materials Research Center, Northwestern University, 2145 N Sheridan Rd, Rm F275, Evanston, IL, 60208-3112
Resources: *Jaws* (Dell Woodcrest Cluster) and *Tempest* (IBM SP4) at MHPCC
Sponsorship: The Air Force Office of Scientific Research #15573FR1 and the Office of Naval Research # 07319050 sponsored this work. We thank Bobbi Ruf (AFOSR) and Annette May (ONR) for their invaluable assistance with the assigned resources.

Parallel Smart Structural Design Optmization System (PASOS)

Rong C. Shieh

This computational structural mechanics project was to develop parallel computational methodologies and an associated general purpose, prototype computer code system, PASOS, for design optimization of both plain and adaptive smart structures. The successfully developed PASOS code, which is based on the uniquely efficient Integral Global/Local Optimization (IGLO) algorithms and parallel finite element structural analysis method, is applicable to both uni- and multi-objective design and active element placement optimization of piezoelectric-type smart structures with continuous and/or discrete design variables. Such parallel design optimization software tool is particularly important and useful in significantly improving the global minimum solution quality and computational efficiency and, thus, design cycle time of large structures and hardware systems. The IGLO-based PASOS code capability was found to be capable of outperforming that based on the Genetic Algorithms by one to two orders of magnitude in the minimum weight design of a 72-bar truss with 72 design variables (Figure 1).

PASOS Methodologies: The PASOS code is based on the following innovative methodologies:

- A set of generalized parallel stochastic-type Integral Global/Local Optimization (IGLO) algorithms and some efficient generalized or improved optimization constraint treatment (CT) schemes, such as the Maximum Constraint Function-based Exterior Penalty Function (MCFEPF), Quasi-Linear Design Scaling (QLDS), Generalized Nonlinear Design Scaling (GNDS), Corrector Augmented External Penalty Function (CAEPF), etc.
- Parallel or parallelized (i) multi-objective optimization methods based on some existing efficient methods, (ii) finite element (FE) analysis procedures of passive and active structures, (iii) a generalized structural design optimization concept that includes both traditional sizing design variables and comprehensive active element design, placement, and selection variables (such as control gain values, locations, and number of active elements), and (iv) uni- and multi-objective structural optimization methodologies based on the above algorithms/schemes/procedures/concepts.
- A parallel computational framework and implementation procedure of the above IGLO algorithms/smart adaptive structural optimization methodologies

Optim. algor/ cnstrn type[1]	W (lbs.)	Wt.-ratio $R_1 = W/W_{GA}$	N_F = No. of fns. evaluated	~Speedup factor (N_F-ratio) relative to GA case
Best GA/EPF [2]	$485.20=W_{GA}$	**1**	100K	1
IGLO/FDSO	480.95	1	2,620	38
IGLO/CAEPF	485.20	1	2,898	35
IGLO/MCEPF	473.43	0.975	2,217	45
IGLO/QLDS	463.64	0.955	381	262
IGLO/QLDS	**$290.50=W_{IM}$**	0.598=1/1.672	100K	1

Table 1a. The GA and various constrained IGLO algorithm-based optimal weight design results of asymmetric 72-bar truss with 72 discrete design variables ({h}={x_L}={0.01} in^2; P3 processor case.

No. of processors	Weight (lbs.)	Elapsed time (s)	Speedup factor
1	291.89**	4050	1.0
2	293.80	2078	1.95
3	297.53	1389	2.92
6	297.45	764	5.30
12	296.65	369	10.98

$N_{GP}/N_{GL}/N_{LM}$ = 40/100/20; {x_L} = {0.01} in^2; stop criterion: $S_{C3} \leq 0.01\%$ and $e_r \leq 0.0001$ (cf. Eq. 12).
* Approximately 35,000 function evaluations were needed to obtain these optimal weight results.

Table 1b. Optimal design performance timing results as functions of no. of parallel processors for an asymmetric 72-bar truss with 72 continuous design variables.

Numerical Results: A numerical evaluation study of the PASOS code algorithms and methodologies was performed on the IBM P3 computer platform at the Maui High Performance Computer Center (MHPCC) in Maui, Hawaii via various example problems with number of design variables ranging from n_x = 16 to 1008 and the number of parallel processors ranging from 2 to 128. Some of the numerical evaluation results (Figure 1 and Table 1a) show that the PASOS code algorithms can outperform those based on the "Best" version of the widely used Genetic Algorithms (GAs)[2] by a factor of up to two orders of magnitude in computational speed in obtaining a semi-optimization weight result and by a factor of nearly 2 in near-global minim weight value after 100K function evaluations. It was also found via the 72-bar/72-design variable (DV) and 1008-bar/1008-DV truss weight minimization problems that the PASOS code timing results are nearly linearly scalable with respect to the number of processors used, i.e., 1 to 16 for the former and 16 to 128 for the latter, with the scalability efficiency factor ranging from 92 % to 100 % (Tables 1b and 2).

Case	Np	Event	W	W/W_0	WC s/h:m:s	Speedup factor
1a: 1st run		@ $W=\sim 0.1W_0$	24,205	0.1	3,715/1:01:55	
1b: Ditto		@ W=14,931	14,931	0.0617	14,910/4:08:30	
1c: Ditto	32	@ $S_{C3}<0.1\%$	14,780	0.0611	14,986/4:09:46	1.0
1d: 2nd run*		@ $S_{C3}<0.01\%$	13,414	0.0554	14,503/4:01:43	P3 proc case
1e: 3c+3d runs		@ $S_{C3}<0.01\%$	13,414	0.0554	29,489/8:11:29	
2a: 1st run		@ $W=\sim 0.1W_0$	24,153	0.0998	2,038/0:33:58	
2b: Ditto		@ W=14,931	14,924	0.0616	7,298/2:01:58	
2c	64	@W=14,780	14,780	0.0611	7,560/2:06:00	1.982
2d: Ditto		@ $S_{C3}<0.1\%$	14,576	0.0602	8,477/2:21:17	P3 proc case
2e: 2nd run*		@ $S_{C3}<0.01\%$	13,474	0.0557	9,745/2:42:25	
2f: 2d+2e runs		@ $S_{C3}<0.01\%$	13,474	0.0557	18,222/5:03:42	
p4 proc case	64	@ $S_{C3}<0.1\%$	14,662		2,290	p4 proc case
		@ $S_{C3}<0.01\%$	13,370		4508s	

Notes: Wo = 242,050# = the minimum structural weight among the initially randomly selected 50 design variables.

Table 2. Performance Timing Results of PASOS (Algor. A6) vs. No. of Processors (Np) Used in Weight (W) Minimization of Asymmetric 1,008-Bar/1,008 Design Variable Truss (min. allow. area=0.1 in^2; Δx=0.01 in^2; allow. disp = 24" & stress = 25ksi).

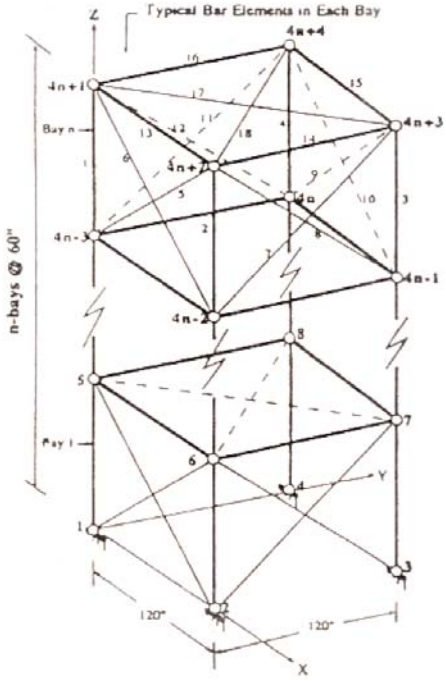

Figure 1a. Minimum weight design of 72-bar truss with 72-design variables (n = 4 bays).

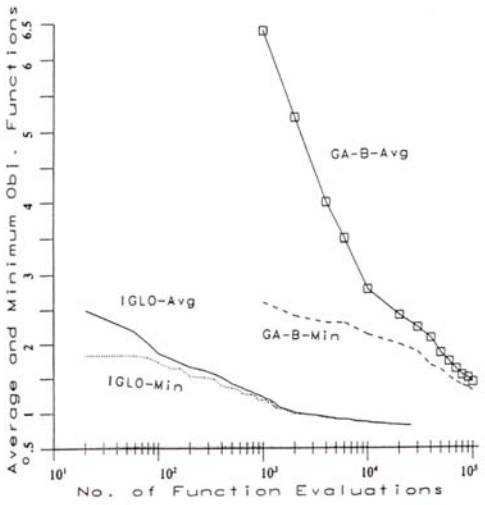

Figure 1b. IGLO and "Best" Genetic Algorithm (GA-B) based weight design histories.

References:

1) Shieh, R. C., "Parallel Design Methodologies for Optimization of Smart Structures," AMT TR-07-01, Final Technical Report for DOD/BMDO SBIR Phase II Contract No. DASG60-03-C-0004, Feb. 2007.
2) Lin, C.-Y. and Hajela, P., "Genetic Search. Strategies in Large Scale Optimization," Proc. 34th AIAA SDM Conf., La Jolla, CA, April 19-22, 1993, pp. 2437-2447.

Author and Contact: Rong C. Shieh
Organization: Principal Investigator/President of Advanced Mechanics Technology, Ltd., 6508 Anna Maria Ct., McLean, VA, 22101

Resources: The high performance computer resource IBM P3-based *Tempest* system was used in developing the PASOS code capability under the earlier Phase I and present Phase II SBIR projects. All HPC resources were provided by Maui High Performance Computer Center (MHPCC) in Maui, HI.
Acknowledgements: This work was supported by DOD Missile Defense Agency under SBIR Phase II Contract No.DASG60-03-C-0004 and managed by the U.S. Army Missile and Space Command. Technical Monitors were Richard Rodger and Alex Gilmore.

The Unmanned Systems Test Bed

Matt Burnham and Shannon Wigent

The primary objective of the Unmanned Systems Test Bed (USTB) program is to develop new training, test, and evaluation concepts and capabilities for unmanned systems and to conduct a series of demonstrations.

The Central Test & Evaluation Investment Program (CTEIP) has teamed up with the U.S. Army Corps of Engineers Topographic Engineering Center (TEC) to examine how to augment DOD Test and Training Range resources with new capabilities in support of unmanned systems. CTEIP is particularly interested in finding effective uses for advanced visualization software and high-performance computing assets to support the needs of unmanned systems. Related goals include leveraging the Test and Training Enabling Architecture (TENA) for cross-range and cross-facility data collaboration. Accordingly, the initial stages of the USTB program focuses on applying advanced visualization and high-performance computing technologies to support the needs of unmanned systems users.

Research Objective: The U.S. military services are evolving their war fighting systems with an increasing emphasis on unmanned vehicles. Although the trend began with intelligence, surveillance, and reconnaissance (ISR) assets (e.g., Predator and Global Hawk), it is evolving to include combat vehicles (e.g., Armed Predator and Air Force and Navy Unmanned Combat Air Vehicle) and utility platforms (Quick Delivery ACTD). The evolving network-centric warfare concepts of all services rely increasingly on unmanned platforms. The associated command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) systems are adopting interoperability standards that include unmanned vehicle mission planning, control, and data dissemination. Recent DOD R&D efforts reflect emphasis toward all types of unmanned systems: air, ground, sea-surface, and underwater platforms. Transformational programs such as the Army's Future Combat System and Navy's SSGN point to future capabilities dependent on unmanned systems. Consequently, future operations may involve the employment of unmanned vehicles at all echelons of command and interoperability and integration of mixed unmanned vehicle types. It is increasingly important for the individual services and the Joint Commands to develop concepts and facilities for training exercises, CONOPS development, and system T&E which include single and combined unmanned vehicle operations. Although much of this will be accomplished in distributed exercises involving computer-generated virtual environments and simulation, it must also include live exercise components on a test range that provides for air, ground, sea-surface, and underwater unmanned vehicle operations.

The USTB Concept

The key concepts behind this test bed are as follows:

1. Gather three-dimensional terrain, imagery and feature data of the test range.
2. Utilize the terrain, imagery, and feature data to render a synthetic environment as a realistic view of the environment the unmanned system is operating in.
3. Integrate, in real time, the movements of the unmanned system into this synthetic environment.
4. Integrate any sensor data (e.g., video) from the unmanned system with the synthetic environment. Project the sensor data like a spotlight on the rendered synthetic environment providing a window to the real world displaying any movement, change, or action in the observed area.
5. Integrate Live, Virtual and Constructive (LVC) entities (i.e., troops, vehicles, etc.) into the synthetic environment to create an augmented environment to support various test, evaluation, and training scenarios.

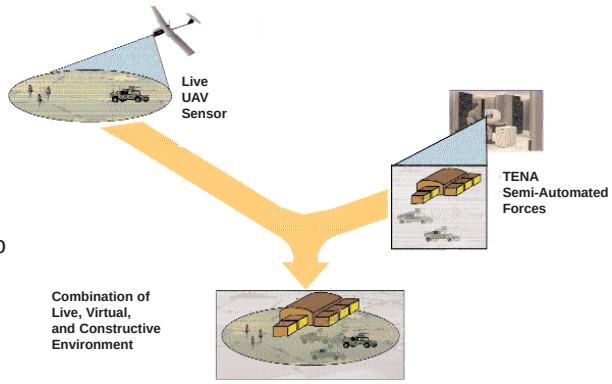

Figure 1. USTB General Concepts.

Figure 1 illustrates the USTB concept showing how the rendered synthetic scenery, live video stream captured by the unmanned system sensor, and LVC entities are combined to create an aggregate battlespace environment.

The USTB System Components

Based on the concepts mentioned, the USTB development efforts fall into three areas of focus:

1. Creation of correlated synthetic terrain datasets that can be used by application in a distributed system, allowing each application to correctly position LVC entities.
2. Integrate these applications into a distributed system using the Test and Training Enabling Architecture (TENA) middleware to create an aggregate synthetic environment which includes the synthetic terrain plus LVC entities.
3. Ingest sensor video streams from unmanned vehicles and projecting them onto the synthetic terrain, thereby augmenting the aggregate synthetic environment.

The following applications have been integrated into the USTB System for use in performing test, evaluation, and training scenarios.

Joint Semi-Automated Forces (JSAF) – Battlespace management software which is used to simulate constructive entities, and control the interaction of these entities with other live and virtual entities. The battlespace is rendered as a two-dimensional map-like view that depicts the entities (troops, buildings, ground vehicles, aircraft, ships, etc.) as icons.

Joint Automated Deep Operations Coordination System (JADOCS) –Joint mission management application that provides the warfighter with a battlespace view for planning, coordination, and execution of targets. The primary use of JADOCS in the USTB scenarios is to issue Air Task Order (ATO) strike orders to allow JSAF entities to engage targets. Like JSAF, the battlespace is rendered as a two-dimensional map-like view that depicts the entities as icons. The primary communication mechanism for JADOCS is though Cursor on Target (CoT) events via a JADOCS server. To integrate JADOCS into a TENA enabled distributed system, a gateway is used to translate CoT events to and from TENA events.

Scene Generator – Software that renders the synthetic terrain data, simulated entities and sensor video streams, to produce the aggregate synthetic environment. Currently Terrex SOFViz is utilized as the rendering system with the added functionality being done via a plug-in. This application has the following functionality:

- Positions all entities in the synthetic environment based on the position in the TENA objects and updates entities locations as updates are received.
- Allows the observer position to be manipulated manually to view the synthetic environment from any desired location.
- Allows the observer location to be moved to the current position of one of the entities.
- Allows the observer to be attached to any of the entities and thereby view the synthetic environment from the point of view of the entity.
- Project and render each frame of the sensor video on the synthetic terrain.

Video Projection and Alignment

Even though this is part of the Scene Generator process, it is a significant software component that receives individual frames of the sensor video stream on the unmanned system, processes the video, and prepares it to be integrated with the synthetic terrain. The primary goal is to align and register sensor video frames in real time to rendered views of the synthetic terrain dataset. A secondary goal is to obtain accurate coordinate information for the permanent and transient objects, features, and entities observed in the sensor video by utilizing the synthetic terrain dataset that has been precisely aligned to geocoordinates. Currently the projection part of the process has been incorporated into the real-time process and some analysis and offline alignment has been done. The results of this work will be used to determine the scope and expected success of incorporating alignment as a real-time process.

Video Projection

The term video projection refers to utilizing an algorithm to map the video onto the three-dimensional synthetic terrain. The simplest projection may involve matching the four corners of the video frame with locations in the synthetic scene by utilizing telemetry data (location, orientation, etc.) of the sensor and texture mapping the video frame onto this area.

Video Alignment

The terms video alignment, video image registration, and video registration are often used interchangeably in the industry, and they refer to detecting and matching (i.e., aligning or registering) the image information from a video with another image. This usually involves image processing to detect shapes, edges, and features on the video frame and the previously captured image, matching the features from both, and aligning them. This can yield the following benefits:

1. Better alignment of the video and the synthetic terrain can be achieved by image processing the video frame and the synthetic terrain and finding matching features. If, for example, edges, features, or objects in the video can be detected and matched with the edges, features, or objects in the synthetic image sequence, a much better positional alignment between the two can be achieved.
2. Since the positional accuracy of the synthetic terrain data is significantly better than the position information of the GPS on the unmanned system, better geolocation information for the objects observed in the video can be calculated by utilizing the geocoordinate information from the synthetic environment. This can be used to determine the absolute location of the stationary objects and features (buildings, roads, etc.) in the video and to determine the relative location of the transient and moving objects (troops, vehicles, etc.) in the video relative to the permanent and stationary objects.

Author and Contact: Matt Burnham
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Author: Shannon Wigent
Organization: Science Applications International Corporation (SAIC), 9555 Kaumualii Hwy, Suite 104, Waimea, HI, 96796
Resources: Dell PCs running Windows XP and Fedora Core 5 Linux
Sponsorship: The Test Resource Management Center's (TRMC) Central Test and Evaluation Investment Program (CTEIP) sponsored this research.

Partitioning and Archiving the Freight Desk Database at ONI

Partha Mandayam, Dave Norman, Carl Holmberg, and Robert Dant

With increased usage over several years, the Freight Desk (FD) database at the Office of Naval Intelligence (ONI) has become huge. This raised both space and performance concerns at ONI:

Space: Some tables already have 500 M+ rows, raising the concern that the disks housing the Oracle data files may run out of space in the not so distant future.

Performance: Also, as the tables keep getting bigger, it takes longer each day to select date range blocks of the growing tables. This has a negative impact on the performance of applications that access the data in these tables, including the ETL processes.

To resolve the above situation, ONI tasked MHPCC to architect a partitioning and archiving strategy that would be easy to implement and maintain. MHPCC engineers utilized the latest Oracle 10g database features on its *Tempest* (IBM Power4 node with 32 processors) computing resource to design, develop, and test a robust partitioning and archiving strategy.

Research Objective: Data in large tables will be partitioned by date range on a new column called OBJECTCREATETIMESTAMP that will be created and populated by Freight Desk Technologies (the vendor of the Freight Desk software). MHPCC will architect the processes and scripts to partition the high-volume tables in the FD schema, both initially and on an on-going basis.

Methodology: There are several methods available to convert a non-partitioned or flat table containing data into a partitioned table containing the same data in a manner that is seamless to all the applications accessing that table and its associated database objects. The option we pick depends on several criteria, as follows:

- **Availability:** Does the table have to be available to other applications while being converted? If so, is it okay if the table is made read-only to these apps while being converted?
- **Quiet Time:** Quiet time means the database is not being inserted into or updated. Does the National Cargo Tracking Program (NCTP) database have a window of quiet time? For how long and how often?
- **Space:** Do you have extra space in the database to accommodate at least twice the size of the flat table that is converted to a partitioned table? If so, is the space available in the current datafile directory or a new datafile directory? If not (worst case scenario), we would at least need enough space on any filesystem to store the flat table's export dump temporarily.
- **Redo Generation:** Some operations are designed to minimize redo. Large redo information generation means a larger number of archived log files, which may overwhelm your tape backup system. Of course, if your database runs in NOARCHIVELOG mode, this lessens the problem, but most production systems tend to be in ARCHIVELOG mode.
- **Rollback Segment:** How big are rollback segments in the databases? This is the second most important factor in the selection of a conversion method. Some techniques rely on reading the source tables. Due to the read consistency requirement of the Oracle database engine, the database must provide the image of the block that was present when the query started. This past image is constructed from the data stored in the rollback segment. If the query is long, the rollback segment may eventually run out of extents to store the past image data. When the long-running query tries to get the data from the segment, it faces an error called "ORA-1555 Snapshot Too Old". Even if the rollback segments are huge, there is a good chance that the query will face the ORA-1555 error if the run time is too long.
- **Retention of flat table:** Is it necessary to retain the original flat table with its data and other aspects (such as constraints, etc.) for a long time just for safe-keeping? This is really unnecessary if ONI deems the partitioning has been accomplished successfully.
- **Data types:** Some of these methods would not work if the table to be partitioned had certain data types such as LONG, LONG RAW, BLOB, etc.

Results: MHPCC reviewed eight well-know partitioning methods that are described in a fair amount of detail in an online white paper on the subject at http://www.dbazine.com/oracle/or-articles/nanda6 and http://www.dbazine.com/oracle/or-articles/nanda7. Among those, methods 4 (SQL*Loader) and 7 (Split-Split) seemed like good options for ONI, with its large tables. Note that this white paper was written in the Oracle 9i days.

With Oracle 10g (the current NCTP Data Warehouse platform), the new Data Pump utilities are available - expdp and impdp provide significant space and speed advantages over all the prior methods. ONI has experienced higher speeds of data transfer using the Data Pump utilities.

Given all this and after reviewing the data pump capabilities and developing and testing a prototype set of partitioning scripts for the complex LEGCONTAINER table (a large table with several associated database objects such as indexes, constraints, stored procedures and packages, a trigger, views, several foreign keys in this and dependent tables), MHPCC came up with the following process for converting a flat table to a partitioned table.

MHPCC's Partitioning Process for ONI

One-time conversion

- Obtain a complete dump of the FD schema from the production data warehouse at ONI, along with other requested partitioning parameters identified in the "Solution" section above.
- Create a development and unit test clone at MHPCC from the above information.
- Develop and unit-test all the necessary partitioning scripts.
- Create a system-test clone with real data (to be provided by ONI) and again thoroughly test all the partitioning scripts on that instance.
- Share the results of the system test with ONI to get their approval to proceed to deployment.
- Repeat the process followed above in system test with a fresh dump of the FD schema from the NCTP production data ware house, to create its partitioned clone.

On-going maintenance

- Follow a periodic, disciplined process to create one or more future monthly partitions to all partitioned tables. As a prerequisite, a named tablespace with a corresponding named datafile also should be created. Here are two sample commands to create a May 2007 partition for the LEGCONTAINER table, for which, let us say we have already created named partitions up until April 2007 and then a default, catch-all partition called LEGCONTAINERFUTURE:

```
CREATE TABLESPACE LEGCONTAINER200705 NOLOGGING
    DATAFILE '/oni/oradata/ONIDEV/LEGCONTAINER200705.dbf'
    SIZE 698K REUSE AUTOEXTEND ON NEXT 698K MAXSIZE UNLIMITED
    EXTENT MANAGEMENT LOCAL
    SEGMENT SPACE MANAGEMENT AUTO;

ALTER TABLE LEGCONTAINER
    SPLIT PARTITION LEGCONTAINERFUTURE AT
          (TO_DATE ('01-jun-2007','DD-MON-YYYY'))
       INTO (PARTITION LEGCONTAINER200705 TABLESPACE LEGCONTAINER200705,
             PARTITION LEGCONTAINERFUTURE TABLESPACE LEGCONTAINERFUTURE);
```

- Validate all database code objects such as triggers, etc. on that table subsequently.
- This process should be followed once a month, once a quarter, etc. as is convenient to ONI.
- This would prevent any newly inserted row from getting into the default, catch-all, anonymously named partition, LEGCONTAINERFUTURE. If rows ever do get inserted accidentally into the LEGCONTAINERFUTURE partition, it would take longer to get them out later to a named partition using the above "partition split" command.

Conclusions/Significance: Having laid a firm foundation with the above partitioning process, with the data being partitioned by month in separate tablespaces, the archival process becomes much easier. The OBJECTCREATETIMESTAMP column values would allow ONI to archive old partitions by specifying an appropriate date range.

Depending on ONI's specific archiving requirements, MHPCC would be able to provide scripts and or processes to enable any one of the following archiving methods:

- Backup the data in old partitions onto tape or other slow media, truncate those partitions and disable them. This would facilitate an easy import back into those partitions if the need ever arose in the future to query the old data.
- Use transportable tablespaces and the Oracle partition exchange facility to move the data in old partitions to dedicated historical datamarts and then drop those partitions from the NCTP data warehouse.
- Use transportable tablespaces and the Oracle partition exchange facility to move the datafiles corresponding to the old partition tablespaces to slower and cheaper read-only disks and thus continue to keep the old data available in the NCTP data warehouse for querying only.

Author and Contact: Partha Mandayam
Organization: Sr. Database Architect, SAIC - Imagery Technology & Sys. Div., 1305 North Holopono Street, Suite 6, Kihei, HI, 96753
Authors: Dave Norman, Carl Holmberg, and Robert Dant
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Resources: *Tempest* (IBM P4 Nodes) at MHPCC
Sponsorship: Office of Naval Intelligence (ONI)

Asynchronous Parallel Optimization of Relativistic Magnetron Cathode for Counter Electronics

P. J. Mardahl, M. T. Bettencourt, and K. L. Cartwright

An automated high power microwave (HPM) source design method using an integration of ICEPIC (Improved Concurrent Electromagnetic Particle-In-Cell[1]: used for simulating HPM sources) and APPSPack[2] (Asynchronous Parallel Pattern Search–used for nonlinear optimization) was developed. This design method was applied to the problem of producing an optimally shaped cathode for the AFRL A66 relativistic magnetron, improving power output by 1.4x, energy and peak efficiencies by 1.5x, and improving mode purity 1.2x. These gains were obtained with approximately 1/10th the time and manpower normally required. The new method continues to be applied and continues to yield improved designs, at a much faster pace than previous design methodology. The result is faster and cheaper development from concept to weapon.

Introduction: We replaced a human-driven optimization process with APPSPack, a nonlinear optimization program from Sandia National Laboratory. With the new design methodology, humans would still decide upon which parameters to vary and reasonable ranges for the parameters. However, the process of actually finding the optimum operation within the parameter space is now automated.

For this effort, we applied the optimizer to four parameters which described the shape of the magnetron cathode, a component from which electrons are emitted into the magnetron.

Problem and Methodology:
The general problem is to produce the most suitable HPM source for weaponry applications possible. Suitability includes many physical and performance characteristics, such as size, weight, form factor, output power and frequency, energy and power efficiency, operator safety, reliability, and maintainability.

For this particular effort we were examining the benefits which could be obtained by inserting a shaped cathode in place of the standard circular-cross-section cathode: Figure 1 shows a shaped cathode in the A6-3 magnetron.

The particular improvements we expected were 1) increased mode purity, 2) increased power and energy efficiency, 3) faster device start-up, 4) increased power output, and 5) better reliability of operation when non-ideal voltage and magnetic field were applied. We did not expect changes in device size or frequency of operation from changing the cathode: in fact, none resulted.

To summarize the cathode shape parameters we varied:
- Cathode core size
- Cathode bump group 1 size
- Cathode bump group 2 size
- Cathode clocking

These parameters formed a 4-dimensional parameter space which had to be explored. A straightforward exploration with 8 points in each dimension would have required 4096 designs be evaluated to produce an optimum.

We applied APPSPack to aid in exploring this parameter space for the optimum. APPSPack only required 84 designs to be evaluated before discovering an optimum.

During design, a maximum of 3,456 CPUs were employed, to evaluate 6 separate designs simultaneously, with each design evaluation requiring 9 runs of 64 CPUs each. For this number of CPUs and for this ICEPIC input, parallel efficiency was approximately 80%.

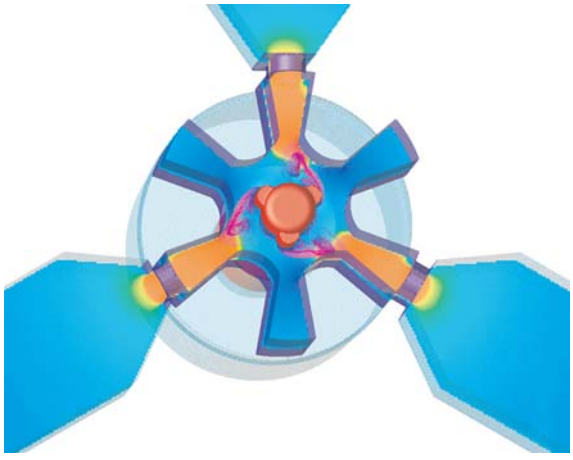

Figure 1. Shaped cathode in the A6-3 relativistic magnetron.

Results:
There were two very significant results of this effort. One was the application of this new design technique to the design of an optimal cathode, which led to a great improvement in the simulated performance of the magnetron. The other was the creation and successful testing of an automated technique for converting design ideas into optimized HPM designs: i.e., the successful integration of APPSPack and ICEPIC.

The new optimal cathode improved the power output 1.4x, averaged over a B-V scan, with each weighted equally. Peak power and energy efficiency were both increased 1.5x for a similar average. Likewise, impedance remained within acceptable ranges, while mode purity increased 1.2x.

Of more significance than the actual improvements to the magnetron design is the new design method. A similar cathode optimization for a previous magnetron design (the A6-3) took a year or more to develop-not a man-year, the problem was not worked full-time, but a calendar year. The new design method produced an optimal cathode in less than six calendar days, at a cost of 300,000 CPU-hours-roughly the same as a human-driven optimization. Further, this design brought improvements in several objectives: not just peak power, but energy and power efficiency, mode purity, and impedance: only peak power was considered previously in human-driven optimization.

References:
1) Peterkin, Jr., Robert E. and Luginsland, John W. March-April 2002. "Theory and design of a virtual prototyping environment for directed energy concepts," Computing in Science and Engineering 4, no. 2, 42.
2) Tamara G. Kolda, Revisiting Asynchronous Parallel Pattern Search for Nonlinear Optimization, SIAM J. Optimization, 16(2):563-586, Dec. 2005. (doi:10.1137/040603589) (http://epubs.siam.org/sam-bin/dbq/article/60358).

Author and Contact: Peter Mardahl
Authors: Matthew Bettencourt and Keith Cartwright
Organization: Air Force Research Laboratory, Directed Energy Directorate, 3550 Aberdeen Dr. SE, Kirtland AFB, NM, 87117
Resources: The CAP utilized the Dell PowerEdge cluster at MHPCC (*Jaws*). Follow-on work to the CAP was performed using the Linux Evolocity II computer at ARL, the Cray XT4 at ERDC, and on "mjm" at ARL. CTA: CEA
Sponsorship: This work was supported by the DoD High Performance Computing Modernization Program and by the Air Force.

Trade Wind Inversion Variability in Hawai'i

Guangxia Cao and Duane Stevens

In the lower troposphere, trade wind inversion (TWI) occurs primarily in fair weather, and is a layer of 200-300 m thickness, characterized by an increasing temperature with height. In contrast, the temperature of the free troposphere above TWI decreases with height. The trade winds and TWI are important controls for the weather and climate in Hawai'i. From sounding data, the TWI base height is about 2,225 and 2,076 m at Hilo and Lihu'e, respectively. Wavelet analysis shows that the TWI has a distinct annual cycle, with two maxima in April and September. Linear regression reveals that there is an increasing trend in the frequency of occurrence for the TWI during the period of 1973-2003. The numerical simulation from Weather Forecast and Research (WRF) model suggests that the island topography contributes to the lifting up of TWI base height at Hilo. Our findings bridge previous research gaps in the understanding of the TWI long-term variability and its regional distribution, and set a basis for the study of its future trend, which is fundamental to predicting the Hawaiian ecosystem and hydrology. Technically, the numerical experiences from this research underlie our current endeavor of severe storm study funded by MHPCC.

Introduction: The trade wind inversion (TWI) over the tropics and subtropics of the Atlantic and Pacific Oceans is one of the most important features of the Hadley circulation. The inversion in the trade-wind regime of the tropics and subtropics is the result of the interaction between large scale subsiding air from the upper troposphere and convection-driven rising air from lower levels (Riehl 1979). The inversion can occur with or without the prevailing trade winds. When the subtropical high lies over Hawaii, or the high moves east or west of Hawaii, the northeasterly trade-winds disappear, but an inversion may still be observed.

There are many gaps in our knowledge of the TWI. First, we have little information on TWI climatology. For example, we do not have convincing information on the annual cycle of the TWI. Schubert *et al.* (1995) suggest that before developing a complete trade wind theory, better understanding of long term TWI characteristics is necessary. It is also important to understand the long-term dynamics of the inversion in order to develop more realistic climate scenarios for Hawaii. We used 1979 - 2003 data from two Hawaiian sounding stations on Hilo and Lihu'e to study the long-term inversion dynamics.

According to Neiburger *et al.* (1961), von Ficker (1936)--the first to present an inversion base height map for the Atlantic Ocean, showed that in the subtropics the inversion base slopes upward from about 300 m near the coast of Africa to about 1,500 m at a distance of 1,500 km east of Africa, then remains fairly constant westward. It also slopes southward from 500 m near 15° N to about 2,000 m near 5° N. Neiburger *et al.* (1961) reported a similar pattern in the eastern North Pacific off the California coast. However, we have no information on the spatial distribution of TWI at a resolution below 10 km for a relative small region such as Hawai'i. The effect of the Hawaiian Islands on the ocean circulation has been documented by Xie *et al.* (2003), but it is not clear how the islands' topography affect the TWI. We use the Weather Research and Forecast (WRF) model to simulate the TWI, and thus deduce its spatial pattern.

Methodology: The study to address climatological inversion variability includes three main components: inversion identification, descriptive statistical analysis, and spectral analysis. The inversion identification procedure uses atmospheric sounding data as input, and produces a data set of inversion occurrence and characteristics. In the descriptive statistical analysis, basic statistical parameters and distribution functions, such as annual cycle and diurnal range, along with inversion occurrence frequency maps, are produced from the derived inversion data set. Spectral analysis, using the wavelet approach, is applied to identify the inversion variability pattern over a range of time scales. Trend analysis is used to identify tendencies in inversion frequency and height over a 25-year period.

.

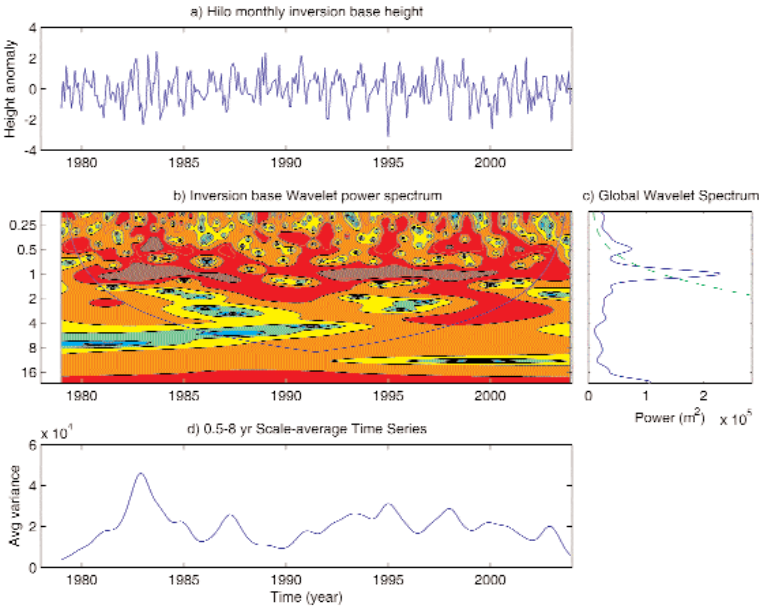

Figure 1. Wavelet analysis of the monthly inversion base height from the data at Hilo.

The overall methodology for studying inversion regional distribution through WRF includes three interrelated components: conducting numerical experiments using WRF with selected physics packages for identifying optimal model physics configuration, simulating two months of the TWI to study its temporal and spatial patterns, and statistical analysis of model output. The WRF modeling is conducted on the MHPCC *Tempest*. We use the NCEP's global final analysis data and U.S. Geographical Survey (USGS) topography data to initialize the WRF model.

Results and Significance: Figure 1 shows the result of wavelet analysis of the monthly inversion base height from the data at Hilo, clearly displaying a significant power peak at annual period. Overall, using 1979-2003 radiosonde data at Hilo and Lihu'e, Hawai'i, the trade-wind inversion (TWI) is found to occur approximately 82% of the time at each station, with average base heights of 2225 m (781.9 hPa) for Hilo and 2076 m (798.8 hPa) for Lihu'e. A diurnal pattern in base height of nighttime high and afternoon low is consistently found during summer at Hilo. Inversion base height has a September maximum and a secondary maximum in April. Frequency of inversion occurrence was found to be higher during winters and lower during summers of El Niño years than non- El Niño years. Significant upward trends were found for inversion frequency at Hilo for March, April, and May (MAM); June, July, and August (JJA); and September, October, and November (SON) seasons; and at Lihu'e for all seasons and for annual values

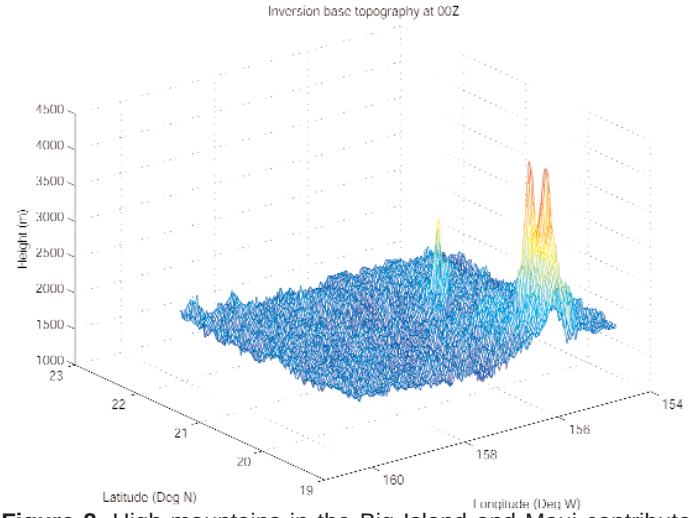

Figure 2. High mountains in the Big Island and Maui contribute to the lift up of inversion base height.

This project also simulates the month-long inversion base height for January and July 2003. The spatial simulation results show that high mountains in the Big Island and Maui contribute to the lift of inversion base height (Figure 2). However the TWI peaks in Figure 2 may be subject to model artifacts. Using the WRF model demonstrated that the microphysics and PBL packages influence simulation of the TWI. The model fails to show a consistent north-south inversion base height slope in the two simulation months, possibly caused by the climatological high in July and low in January north of the Hawaiian Islands.

In summary, this project concludes that the annual cycle for the Hawaii TWI reveals an increasing trend for the TWI frequency of occurrence and suggests its regional distribution. In the future we will look at some boundary layer problems for extreme weather in Hawai'i. Currently, we are conducting WRF experiments on Jaws at MHPCC, supported by a MHPCC Engagement Grant, to simulate one of the flooding events in Oahu - the 2004 Halloween Flash Flood.

References:

1) Cao, G., T. Giambelluca, D. Stevens, and T. Schroeder 2007: Inversion variability in Hawaiian trade wind regime. J. Climate, 20, 1145-1160.
2) Cao, G., 2007: Trade Wind Inversion Variability, Dynamics and Future Change in Hawai'i. PhD Dissertation, University of Hawai'i, pp. 191.
3) Neiburger, M., D. S. Johnson, and C-W. Chien, 1961: Studies of the Structure of the Atmosphere over the Eastern Pacific Ocean in Summer: I. The Inversion over the Eastern North Pacific Ocean. Berkeley: University of California Press.
4) Riehl, H., 1979: Climate and Weather in the Tropics, Academic Press: London, pp. 623.
5) Schubert, W. H., P. E. Ciesielski, C. Lu, and R. H. Johnson, 1995: Dynamical adjustment of the trade wind inversion layer. J. Atmos. Sci., 52, 2941-2952.
6) von Ficker, H., 1936: Die Passatinversion. Veröff. Meteor. Inst. Berl. 4, 1-33.
7) Xie, S.-P., W.T. Liu, Q. Liu, and M. Nonaka, 2001: Far-reaching effects of the Hawaiian Islands on the Pacific ocean-atmosphere. Science, 292, 2057-2060.

Author and Contact: Guangxia Cao
Author: Duane Stevens
Organization: Dept. of Meteorology, University of Hawai'i at Manoa, 2525 Correa Road, Honolulu, HI, 96822.

Resources: IBM SP3/SP4 *Tempest* at MHPCC
Acknowledgement: The EWC fellowship, TMCF assistantship, and 2006 MHPCC Engagement Grant are acknowledged. Thanks to Drs. Tom Giambelluca for support and advice, and Sue Brown for administrative and logistics support.

Modeling Heterogeneous Patient Populations using Gene Expression for Colorectal Cancer Prognosis

Vidya Kamath, Lawrence O. Hall, Dmitry Goldgof, Rangachar Kasturi,
Timothy J. Yeatman, and Steven Eschrich

Cancer is a disease process that emerges from a series of genetic mutations causing seemingly uncontrolled multiplication of cells. Molecular genetics indicates that various combinations of genetic events or alternative pathways in cells may lead to cancer. The use of gene expression microarrays can interrogate much of the transcriptome (54,000 transcripts), allowing the possibility of identifying signature patterns that distinguish clinically relevant conditions in patient samples. Typically a small set of genes is chosen to use within a signature; explicitly modeling of the heterogeneous nature of cancer is not generally addressed. We have developed an algorithm for the exploration of gene expression data using multiple classifiers, to identify the multiple molecular signatures that describe prognosis of colorectal adenocarcinoma. Each classifier can potentially use a different set of optimal features, thereby allowing the exploration of many different aspects of the diseased tissue, rather than the traditional approach of a single signature.

Project Goals: As part of the Maui Institute for Molecular Medicine, our group is analyzing gene expression microarray data from colorectal adenocarcinoma patients. The overall goal for the project is to identify gene signatures that predict patient outcome. Although clinical staging provides the baseline for recommending treatment, many early-stage patients do not require treatment and many late-stage patients do not respond to current treatments. Therefore, a molecularly-guided approach to therapy and prognosis may provide significant advances over current clinical care.

The aim of this specific project is to explore the characteristics of gene expression data to understand the underlying heterogeneity of the disease. As demonstrated in other cancers (notably breast cancer), specific genetic mutations may indicate the success of a particular therapy (e.g., herceptin). When considering the existence of a heterogenous population of colorectal cancer patients, one mechanism for studying the different aspects of the population is to explore the space of possible solutions. The random subspace machine learning approach is used to generate information on the heterogeneity of classifier performance. We aim to create a large number of subspaces to explore the data in a more complete fashion, and combine the features and classifiers that perform well in this setup so as to create a generalized and efficient predictor of colorectal cancer prognosis.

We utilized 121 samples from colorectal adenocarcinomas from the Moffitt Cancer Center Tumor bank, collected under a protocol designed to address the question of clinical outcome (Yeatman). The patient samples were labeled as good prognosis (84) cases if the patient survived greater 36 months, and bad prognosis (37) if the survival was less than 36 months. Gene expression values for 54,675 probesets were measured using the Affymetrix Human Genome U133 Plus 2.0 GeneChip.

Research: Data mining classifiers such as decision trees, support vector machines and neural networks build decision boundaries to distinguish classes using the entire input feature set. Feature selection may be done prior to classification to select a small set of features that are important for classification. However, building a single classifier using a single set of features could lead to over-fitting, or optimization of the classifier on the training data. Further, only the specific characteristics of the data used for the feature selection will be incorporated into the classifier decision function. Alteration of the training configuration, such as using different criteria for feature selection, or using a different sample set for training, yields different classifier models. Some of these models could be very strong in predicting the classes, while others may have varying degrees of weaknesses.

Ensemble techniques, such as random subspaces, can be used to explore the characteristics of such classifiers. The random subspace approach, described in Figure 1, is simply a repeated random sample of features to construct an ensemble of classifiers. Ten thousand classifiers were created on the gene expression dataset, each using 200 randomly selected features. The classifiers were trained using 90% of the samples, and tested by predicting the classes of a held-out 10% of samples. If relatively few genes are useful for classification purposes, few classifiers created using a random subspace approach will be accurate on independent test samples.

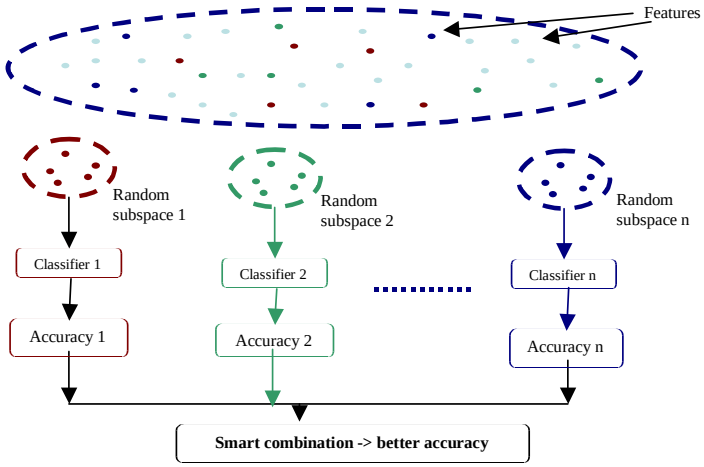

Figure 1: Basic random subspace approach.

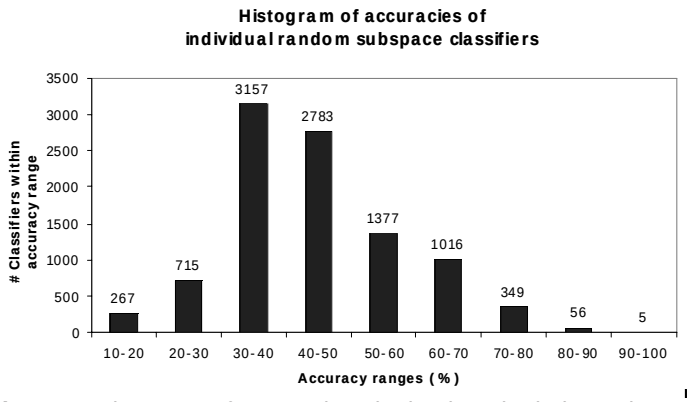

Figure 2. Histogram of accuracies obtained on the independent test samples.

Figure 2 shows a histogram of classifier accuracy, created using generally uninformative features on independent test samples. Less than 1% of the classifiers were found to have weighted accuracies (average of per-class accuracy) greater than 80%. In particular, classifier performance on the poor prognosis class is generally far worse in most subspaces. This suggests that the proposed approach could be used to test the hypothesis that multiple, heterogeneous tumor profiles lead to poor prognosis.

We first attempted to account for the heterogeneity of the classifiers by building an ensemble predictor that may be more general as well as accurate in predicting the individual prognostic classes than any of the individual classifiers. We selected all classifiers with accuracies greater than 80%, and used them in an ensemble. The ensemble was tested on a set of validation samples that had not been seen by the classifiers. Each selected classifier was used to predict the class of a validation sample and the final class assignment was determined by a weighted majority vote. Weights were computed as the prior probabilities of the voted class. This technique improved classification accuracy over using all possible classifiers however an ensemble of independent classifiers can be difficult to interpret.

As an alternative to constructing an ensemble model, we are currently developing a single model that utilizes genes identified from the random subspace classifiers. Our modified random subspace algorithm is therefore used as a multivariate feature selection tool. In the random subspace approach, one classifier is built on a single random subset of features. Subspaces yielding classifiers that perform with high accuracies on independent test samples were selected for building the final classifier model. All the features in the selected subspaces were pooled together and used as an input to the final predictor model. Therefore, the heterogeneous nature of the clinical samples is being captured through the random subspace sampling. The final classifier, built from the individually accurate subspace classifiers, has the potential to be at least as accurate as the least accurate subspace. The subspaces may be selected in various ways. The method of selection will dictate the quality of the final predictor. We chose to select subspace classifiers that had equally high accuracies in each class, or classifiers with a balanced accuracy of prediction. This ensures that the final predictor is not inordinately biased towards any one class, and is capable of describing both classes with sufficient accuracy.

Results: Our results suggest that colorectal tumors are heterogeneous in gene expression, particularly with respect to the prediction of outcome. Useful features can be found in large gene expression datasets by the means of random subspace techniques and assembly of the accurate subspaces in various ways leads to different types of classifiers. Groups of these predictors may, in fact, define unique signatures of colon cancer prognosis. These signatures maybe further refined by generating larger numbers of subspaces that explore the data in more detail. The outcome of these experiments will be a well-refined and generalized signature of colon cancer prognosis that encompasses this heterogeneity.

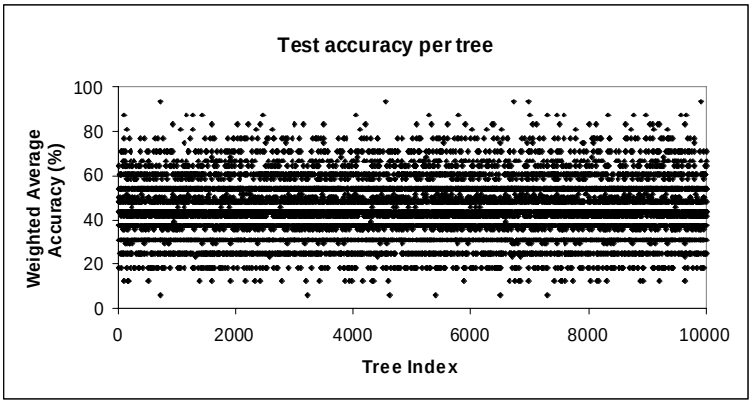

Figure 3. Prediction performance of random subspace classifiers on independent test samples.

Author and Contact: Vidya Kamath
Authors: Timothy J. Yeatman and Steven A. Eschrich
Organization: H. Lee Moffitt Cancer Center and Research Institute,12902 Magnolia Drive, Tampa, FL, 33612
Authors: Lawrence O. Hall, Dmitry Goldgof, and Rangachar Kasturi
Organization: Department of Computer Science, University of South Florida, 4202 E Fowler Avenue, Tampa, FL, 33620
Resources: *Minihui* at MHPCC and BRCA 12-node AMD Cluster at the University of South Florida
Sponsorship: National Functional Genomics Center, HLM and the Maui Institute of Molecular Medicine
Acknowledgements: This research was partially supported by the Department of Defense, National Functional Genomics Center project, under award number DAMD17-02-2-0051. The information contained in this document does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

Seismic Soil-Structure Interaction Analysis of the Kealakaha Stream Bridge on Parallel Computers

Seung Ha Lee and Si-Hwan Park

The Hawaii Department of Transportation is planning on constructing a new bridge on the Mamalahoa Highway over the Kealakaha stream, Island of Hawaii. The research team was faced with a computational model that is beyond the capability of a single-CPU machine as the model contains an extensive soil region in the vicinity of the bridge to accurately represent the seismic response of the bridge-soil system. The main objective of the current study was to develop a parallel computational framework for large-scale seismic soil-structure interaction analysis.

Introduction: The Hawaii Department of Transportation is planning on constructing a new bridge on the Mamalahoa Highway over the Kealakaha stream, Island of Hawaii. The three-span prestressed concrete bridge, depicted in Figure 1, will be designed to withstand the anticipated seismic activity, which is particularly important because of the very high seismic activity in the Island of Hawaii. In the UH Civil and Environmental Engineering Department, there is an ongoing research project on soil-structure interaction modeling of the Kealakaha Bridge. The research team was faced with a computational model that is beyond the capability of a single-CPU machine as the model contains an extensive soil region in the vicinity of the bridge to accurately represent the seismic response of the bridge-soil system. The main objective of the current study is to develop a parallel computational framework for large-scale seismic soil-structure interaction analysis that can be used on commonly available clusters.

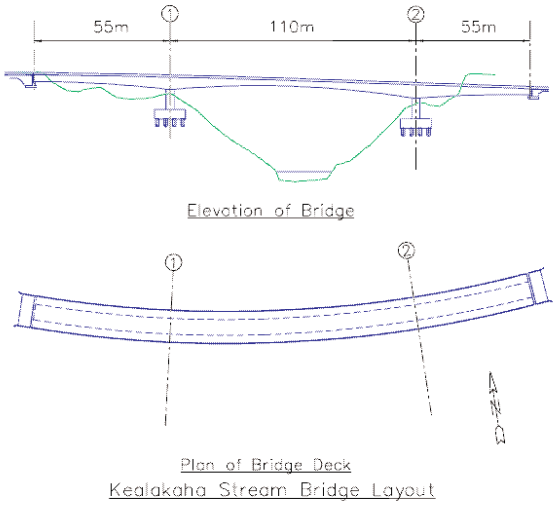

Figure 1. Layout of the Kealakaha Stream Bridge.

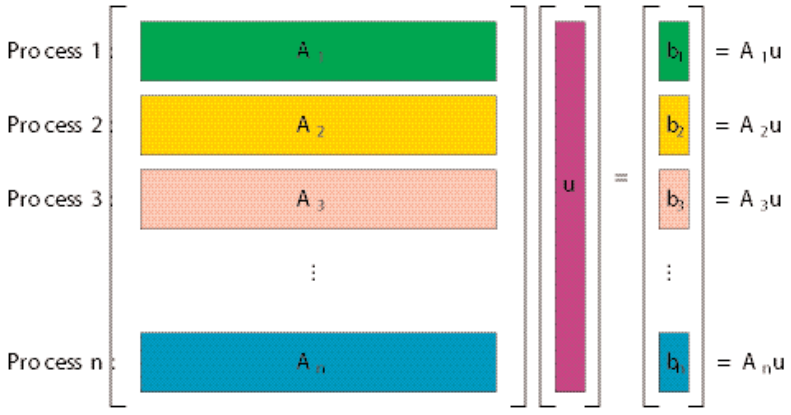

Figure 2. Compressed sparse row format for parallel matrix-vector multiplication.

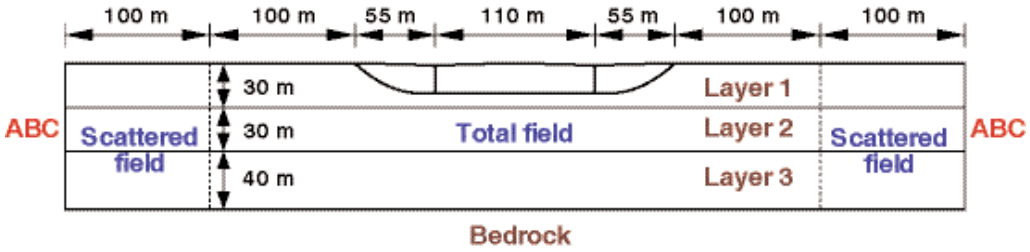

Figure 3. Computational model of the bridge-soil system.

Parallel Computational Framework:

We used the central difference method to solve the time-dependent equations of motion thus avoiding solving linear equations. The classical viscous absorbing boundary condition (ABC) was used to approximately simulate the unbounded extent of the soil region. It was recognized that matrix-vector multiplication is the most intensive part of the computation, and so our effort was concentrated on this aspect of implementation in parallelizing the code. We partitioned a matrix-vector multiplication $\boldsymbol{Au} = \boldsymbol{b}$ as depicted in Figure 2. Since matrices are sparse, the compressed sparse row format is used where process $\boldsymbol{i}$ contains only the non-zero entries of $\boldsymbol{A}_i$. In the MPI-based computational procedure, first $\boldsymbol{u}$ is broadcast and $\boldsymbol{A}_i\boldsymbol{u} \equiv \boldsymbol{b}_i$ is computed on each slave process. Then the result is sent to the master process, where $\boldsymbol{b}$ is obtained by assembling the results from all processes. The *Squall* system at MHPCC was used for the computations.

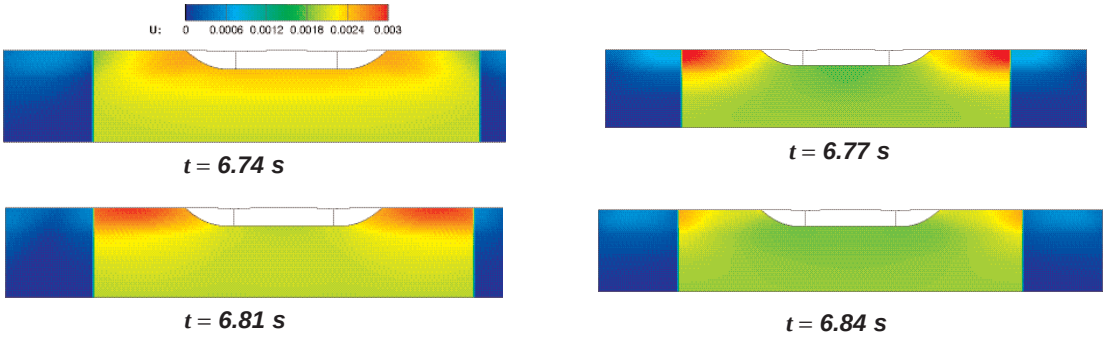

Figure 4. Displacement amplitude contour.

Computational Model and Numerical Results:
In the computational model shown in Figure 3, the soil consists of three layers, and a relatively large buffer region denoted as the scattered field was used to help the viscous ABC absorb outgoing waves. 5,376 solid and 56 frame finite elements were used to discretize the model. To satisfy the stability criterion, $\Delta t = 2 \times 10^{-5}$ s was used, which necessitated undergoing 1.5 million time steps for the response analysis due to a 30 second-long earthquake recorded at the Pahala station on the Island of Hawaii. It took about 4 hours and 30 minutes on the *Squall* cluster using 16 processors. Figure 4 displays the displacement amplitude contour on the soil region at a few different time instants. Propagation of waves from the interior to the outer edge of the domain is clearly seen.

Conclusions:
In this work we have established a computational framework for large-scale soil-structure interaction analysis. In the future, we would like to consider a 3D model of the bridge-soil system using a more efficient parallel implementation.

Author and Contact: Seung Ha Lee
Author: Si-Hwan Park
Organization: Department of Civil and Environmental Engineering, University of Hawaii at Manoa, Honolulu, Hawaii, 96822
Resources: IBM P3 *Squall* system at MHPCC
Sponsorship: University of Hawaii
Acknowledgement: The authors would like to thank the University of Hawaii and the Maui High Performance Computing Center for their support of this project.

UHF Track-Reactive Simulation

Donald J. Fabozzi II

Traditional airborne surveillance simulations have been limited in either spatial or temporal fidelity due to the expensive software and hardware requirements. Recently, advances have been made which provide the rapid deployment of high-fidelity simulation through a modular visual programming environment on a High Performance Computer (HPC). Based on the visual programming environment Khoros, the Radar Analysis Simulation Tool (RAST-K) is a flexible simulation for quickly prototyping airborne surveillance configurations containing radar system features, point targets, and USGS maps. Additionally, RAST-K has been ported to a Linux cluster to simulate realistic flight scenarios. As flight scenarios involve parametric changes between Coherent Processing Intervals (CPIs), additional interfaces were developed to control platform, target, and environmental attributes to correctly model realistic airborne radar surveillance. Interfaces were then developed which partition the simulation, signal processing, and visualization across processors. The recent phase of development has expanded the simulation to incorporate closed loop processing including Doppler processing, Kalman filter tracking, and resource management, as shown in Figure 1.

Research Objectives: The primary objective of this research effort is to establish a simulation and modeling testbed to support the UESA radar system at the Makaha Ridge Test Facility, Kauai, HI. This testbed was established to model the high-fidelity single CPI (pulse-to-pulse) radar returns and investigate multi-CPI platform deployment characteristics.

Methodology: The methodology involved:

1) The establishment of a high-fidelity single CPI pulse-to-pulse simulation capability. Radar Analysis Simulation Tool for Khoros (RAST-K) software was selected to meet this requirement, as shown in Figure 2.
2) The development of a capability that would replicate single CPI's and apply real-world, time-dependent principals.
3) The development of the execution and visualization interfaces for the multi-CPI system to run on a High Performance Computer (HPC).

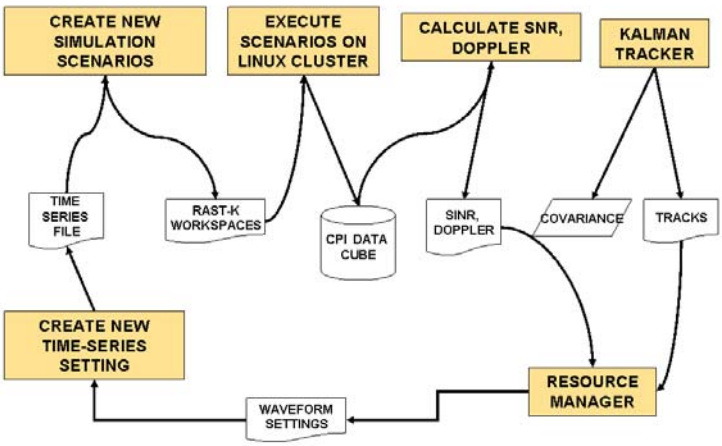

Figure 1. Closed Loop Simulation Architecture.

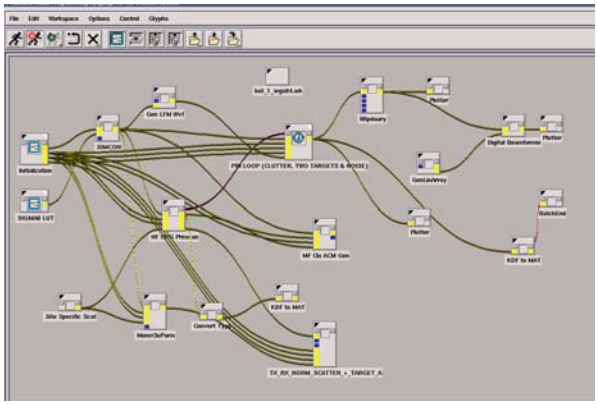

Figure 2. UESA single CPI pulse-to-pulse simulation.

Results: AFRL/MHPCC fielded a distributed computing system consisting of a Linux application and file server and a 16-node cluster to support the simulation environment.

The developed system generates multiple CPI radar returns, each representing a realistic radar return, in near real time. The system has been tested in the Socorro, NM, Makaha Ridge, HI, and China Lake, CA regions with various excursion types at each location. A typical use is to investigate the degree of shadowing and reflectivity from a position in space, as pictured in Figure 3. The left-hand figure shows the platform positioned at the center of the picture emanating due south over the San Gabriel Mountains in California. The right-hand picture in Figure 3 shows the labeled GIS terrain map. As shown, the ability to quickly associate surveillance radar returns with digital maps provides the analyst with advanced decision making tools.

The RAST-K system is a considerable time saver for multi-CPI simulation in that the generation of 162 beam positions, which normally takes about 45 minutes on a single processor desktop, now takes about 4 minutes. Further, this data partitioning model is highly efficient at nearly 70-90%, taking into account the node-scheduler availability.

Significance: The U.S. Navy is currently investigating a host of radar surveillance concepts that resolve thousands of targets, while maintaining a scanning mode. As such, the development of actual, as well as simulated, systems is providing the leading capabilities to meeting those goals. The RAST-K simulation system at MHPCC is currently unique in that it provides both high-fidelity pulse returns, as well as multi-CPI realistic radar returns. The impact of this simulation capability to DOD surveillance research and development efforts can potentially be very significant.

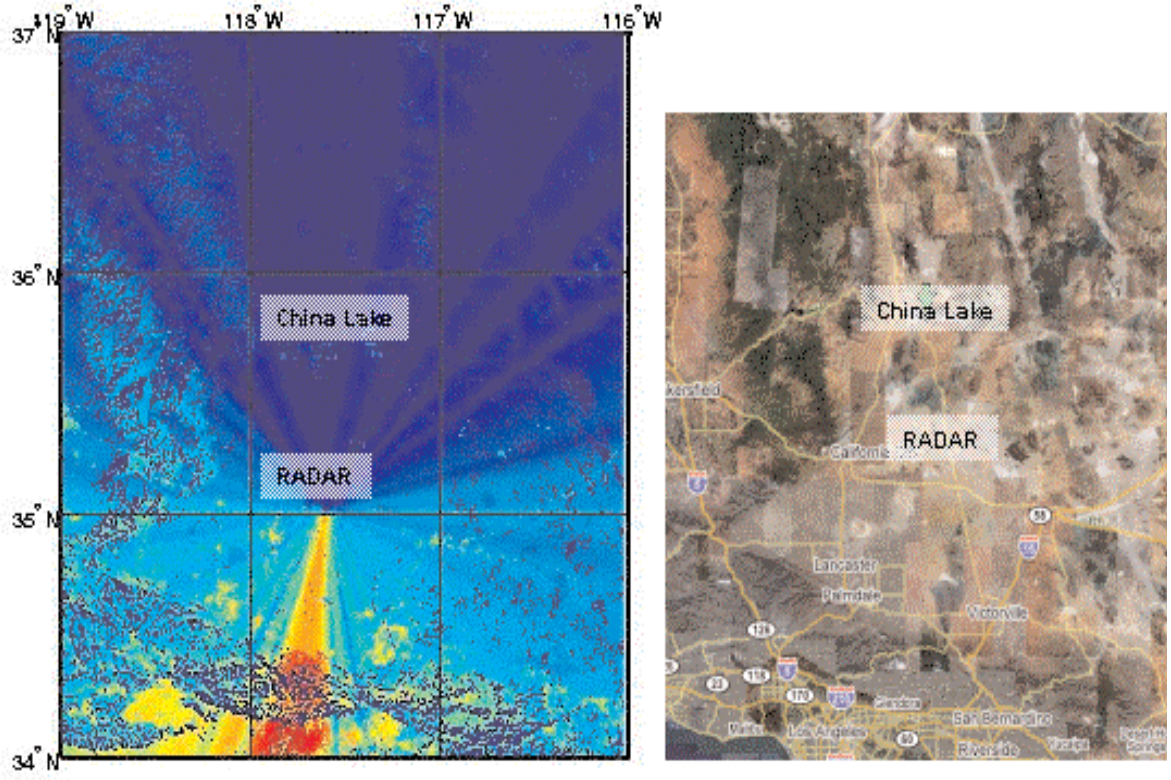

Figure 3. UESA "Searchlight" beam illuminating the China Lake, California region, along with accompanying map. This was modeled with the RAST-K system at MHPCC.

Author and Contact: Donald J. Fabozzi II
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: HP ProLiant DL360 G4p server and a 16 node/32 Processor Linux Cluster, 48 GFLOP peak
Sponsorship: Office of Naval Reseach
Acknowledgement: The author would like to acknowledge Dr. Mike Pollock, ONR, and Dr. Brian Freburger, NRL.

PCID and ASPIRE 2.0 - The Next Generation AMOS Image Processing Environment

Charles L. Matson, Charles C. Beckner Jr., Kathy Borelli,
Tom Soo Hoo, Shiho You, Brandoch Calef, Maria Murphy, and Ron Viloria

One of the missions of the Air Force Maui Optical and Supercomputing (AMOS) site is to generate high-resolution images of space objects using the Air Force telescopes located on Haleakala and the computers at the Maui High-Performance Computing Center (MHPCC). Two routinely used methods for overcoming the atmospheric blurring resulting from ground-based telescopes are the real-time use of adaptive optic systems and the use of image restoration algorithms on short-exposure images of the space object. Recently, a multi-frame blind deconvolution (MFBD) algorithm called Physically-Constrained Iterative Deconvolution (PCID) has been efficiently parallelized and is able to produce image restorations in as little as a few seconds. Because the algorithm can be complicated to use, a GUI is being developed to be the front end to the PCID algorithm. This interface, called the Advanced SPeckle Image Reconstruction Environment (ASPIRE) version 2.0, is the next generation of the current ASPIRE GUI used as a front end to the bispectrum algorithm. ASPIRE 2.0 will be the front-end GUI to PCID, the bispectrum algorithm, and the AMOSphere database. In this paper we describe ASPIRE 2.0 and PCID and how they can be used to obtain high-resolution images.rational space catalog.

PCID: The PCID algorithm is a MFBD algorithm that jointly estimates the pristine image and the point-spread functions (PSFs) that generated the sequence of blurred images used as input to the algorithm. The PCID algorithm has been parallelized to run efficiently on commodity clusters (i.e., clusters consisting of computers that do not have specialized hardware in them, connected by standard networks such as Ethernet, Infiniband, or Myrinet). We have demonstrated excellent scalability for tens to hundreds of nodes, depending on the speed of the interconnect and the sizes of the raw image arrays. Image reconstruction times range from a few seconds to hours, depending on the number of frames of data used, the number of conjugate gradient iterations, and the sizes of the data arrays.

PCID and ASPIRE 2.0 are hosted on two AMD-Opteron-based clusters at the MHPCC. These two clusters started out as a single Cray XD1 cluster, but were obtained by splitting the single cluster into two parts. One of these two clusters, called *Polaris*, is used for routine AMOS data reduction while the second, called *Hoku*, is used for research and development supporting the software on *Polaris* as well as other R&D tasks.

ASPIRE 2.0: Effective image processing at AMOS involves more than just employing image restoration algorithms that produce high-quality restorations. It is also necessary to be able to access the raw data to be processed and the parameters that describe how the raw data was collected, to calibrate the raw data, to submit multiple image processing jobs to the host computer, to visualize and store the restored images, and to accomplish all of these tasks in an intuitive and easy-to-use manner. The ASPIRE 2.0 environment is being created to achieve all of these goals.

The next generation ASPIRE 2.0 image processing environment consists of a number of proven, leading edge technologies and languages such as ColdFusion, CFML, XHTML, CSS, Perl, Java, JavaScript, MySQL, and Apache working in unison to provide an effective, functional, and pleasant user experience through the ubiquitous use of the World Wide Web. Innovative techniques have been implemented to empower the ASPIRE 2.0 user to focus on tasks at a higher level without the burden of dealing with the infrastructure at a lower level.

A key feature of ASPIRE 2.0 is that it will be seamlessly integrated with AMOSphere, the AMOS site data repository and the location from which the vast majority of data that is generated and processed by the site is disseminated to customers. The ASPIRE home page is shown in Figure 1.

Figure 1. The ASPIRE home page.

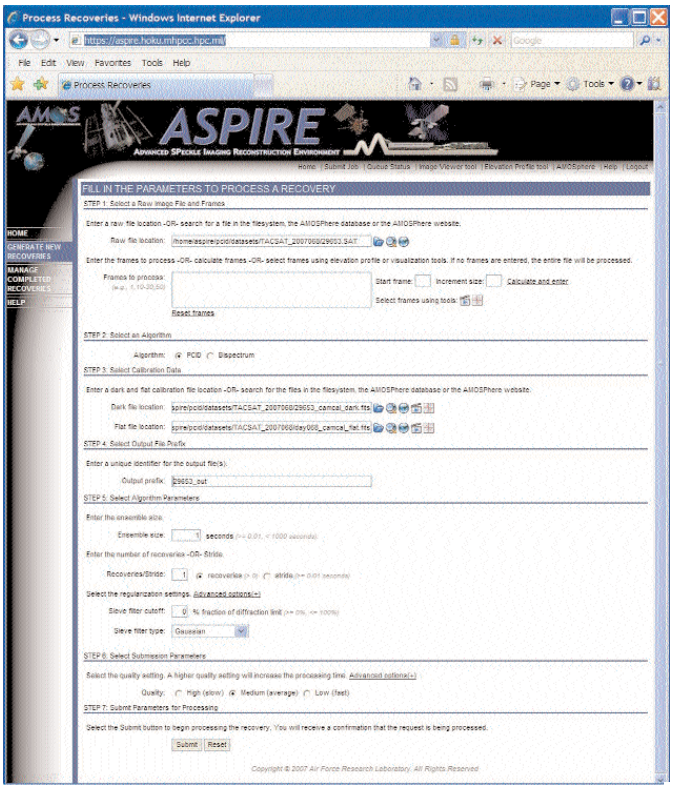

Figure 2. The ASPIRE "Generate New Recoveries" page.

The ASPIRE web interface provides access to necessary tasks such as generating, managing and viewing new image recoveries. From the Generate New Recoveries page shown in Figure 2, the user is guided through the seven steps that must be completed to successfully generate some image recoveries. As many of the inputs as possible are filled in automatically for the user, based upon the data itself. A dataset can be selected from AMOSphere or the file system using search and browsing icons, and tools that allow the user to browse through the raw images in the data file and view the observation parameters are accessed via the Select Frames icons. The user selects the desired algorithm, the calibration data files for the raw image, and a prefix to prepend to all of the output files to identify. All parameters relevant to the algorithm chosen are presented to the user. The user can select a desired image quality based upon the amount of time required to carry out the processing. A job is sent to the batch queue on the processing computer when the user submits the web form data.

The 'Manage Completed Recoveries' page, shown in Figure 3, displays a list of jobs that are either running or have finished running but have not yet had their results analyzed and either stored or discarded. Any running job can be canceled by the owner of the job. Any job still in the list, running or completed, can be resubmitted if desired. Finally, for all the jobs that have been completed, the recoveries can be viewed with a movie viewer, and/or can be stored into AMOSphere, and/or can be deleted without storing into AMOSphere. The middle portion of the page displays the details of the particular job selected in the list and the bottom region shows all the jobs in the queue. The queue information is particularly useful for users who want to submit one or more jobs to get an idea of who is ahead of them in the queue and what kind of turn around time might be expected.

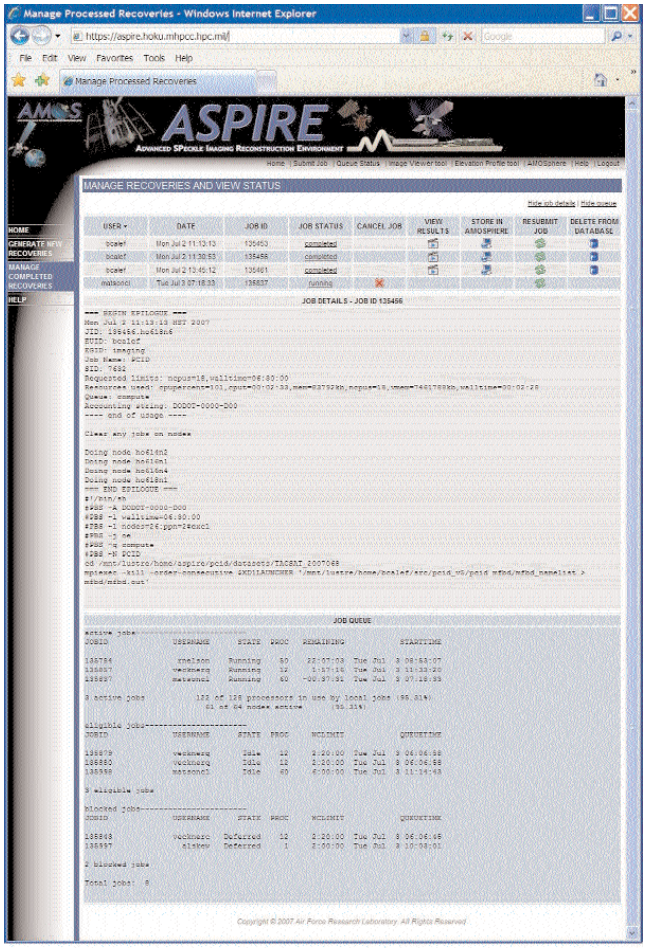

Figure 3. The ASPIRE 'Manage Completed Recoveries' page.

The 'Image Viewer' tool (Figure 4) can be used to assist in the data reduction process. The user can review the data in the 'File Data' column to get an initial idea of the quality of the raw data in order to determine whether or not it is worth processing to generate recoveries. If the dataset looks promising from these numbers, the user can then use the 'Image Visualization' tool to play a movie of all the recoveries and visually assess the raw data. The 'Image Viewer' tool can also be used to analyze the quality of the image recoveries generated by PCID or the bispectrum algorithm. The user can click on the 'View Results' icon for the desired set of recoveries and view them with this tool. Based on the quality of the images seen using the 'Image Visualization' movie viewer, the analyst can decide whether or not to store the recoveries in AMOSphere.

Conclusions and Future Work: A new generation of image processing software, consisting of the PCID algorithm and the ASPIRE 2.0 environment, is being developed and transitioned for routine processing of AMOS data. The PCID image processing algorithm is an MFBD algorithm that closely approaches the theoretical limits to image quality and has been to generate image recoveries in seconds to minutes. A significant upgrade of the current ASPIRE 1.0 software environment, ASPIRE 2.0, is being developed that will provide an easy-to-use interface to the PCID algorithm as well as the current operational bispectrum algorithm. ASPIRE 2.0 will also be seamlessly integrated with AMOSphere to provide a site-wide image processing and data management capability. It will provide suggested values for parameters needed to run PCID or bispectrum that can be overridden, if desired. In addition, it will provide easy access to AMOSphere or the processing computer's file system to select raw data files and to store image recoveries.

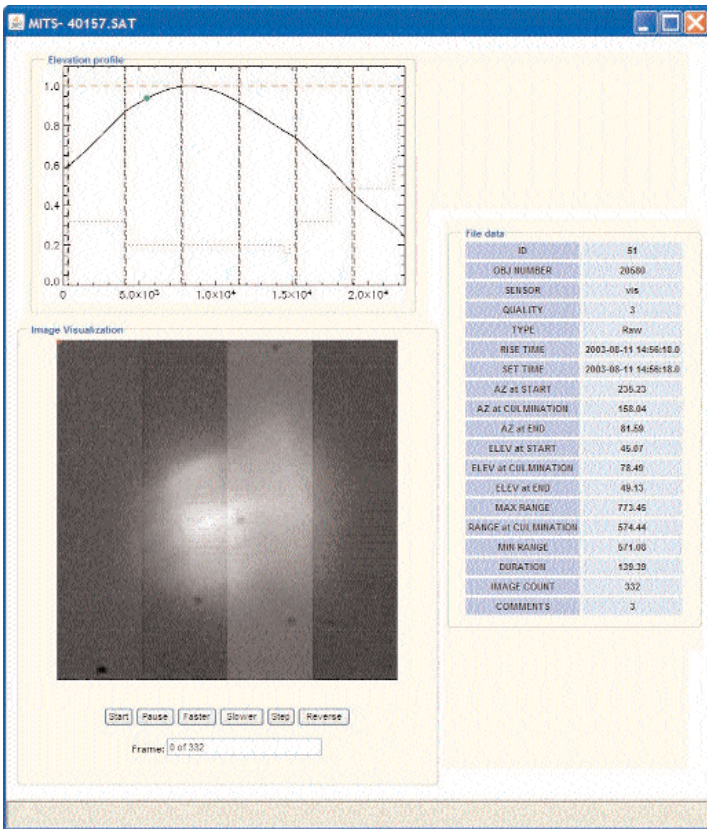

Figure 4. The ASPIRE 'Image Viewer' page.

Author and Contact: Charles L. Matson
Author: Charles C. Beckner, Jr.
Organization: Air Force Research Laboratory/DESA, Kirtland AFB, NM, 87117-5776
Author: Kathy Borelli
Organization: KJS Consulting, 71 Awalau Road, Haiku, HI
Authors: Tom Soo Hoo, Shiho You, and Brandoch Calef
Organization: Boeing LTS, 535 Lipoa Parkway, Kihei, Maui, HI, 96753
Authors: Maria Murphy and Ron Viloria
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Resources: Cray XD-1 (*Hoku* and *Polaris)* at MHPCC
Sponsorship: AFRL, AFOSR, and HPCMO were the sponsors of this research.
Acknowledgements: The authors wish to thank the Air Force Research Laboratory's Directed Energy Directorate, the Air Force Office of Scientific Research, and the DoD High Performance Computing Modernization Office for their financial support that has made this work possible.

A Comparative Study on Boron and Olefin Metathesis[1]

Eluvathingal D. Jemmis, Susmita De, and Pattiyil Parameswaran

Boron metathesis is triggered by the attack of the substrates at the positively charged B atom of the Fe-borylene complex and proceeds via the preferred acyclic intermediate. In contrast, olefin metathesis prefers a four-membered intermediate resulting from the interaction of the olefin with the metal of the carbene complex. Polar substrates with low-lying σ*-MO (weak σ-bond) are preferred in boron metathesis reaction.

Research Objectives: The first metathesis reaction[2] of transition metal borylene complex $[CpFe(CO)_2BN(i\text{-}Pr)_2]^+(BAr^f{}_4)^-$ (Ar^f = 3,5-$(CF_3)_2C_6H_3$) with AX [A is Ph_3P and Ph_3As, and X is O, S] has been reported recently. We study this reaction to obtain the variations possible in the metal-borylene complex and the substrates, and the mechanism of the reaction in relation to olefin metathesis (Nobel Prize in Chemistry in 2005).[3] This also forms a part of our attempt to find analogies between carbon and boron.[4]

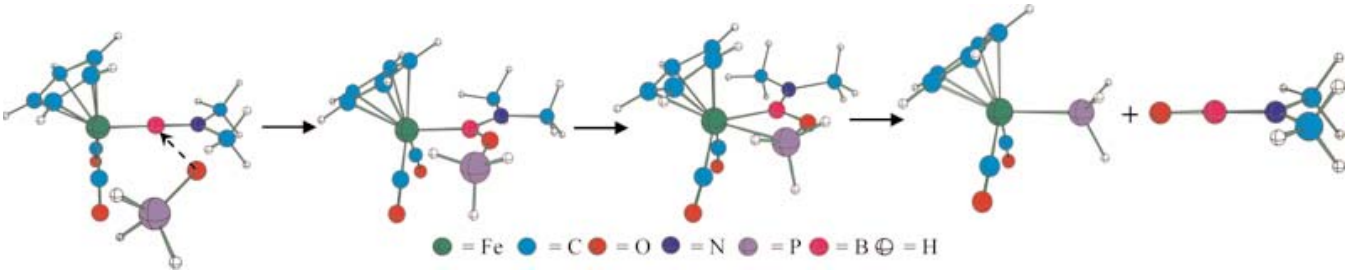

Figure 1: The Mechanism of Boron Metathesis Reaction of $Cp(CO)_2FeBN(CH_3)_2^+$ with H_3PO.

Methodology: The Gaussian 03 Program package[5] at MHPCC was used for all calculations. Geometry optimization was carried out at B3LYP/LANL2DZ and BP86/ZORA/TZ2P levels of theory. All the structures were characterized by frequency analysis.

Results: The analysis of bonding and charge distribution shows that Fe-borylene complex is a Fischer-type carbene analogue. The initial attack of the substrates takes place at the positively charged B of the Fe-borylene complex and forms the preferred acyclic intermediate (Figure 1). The energetics involved in the boron metathesis is comparable to that of the olefin metathesis. The boron metathesis is more favorable for those substrates which have low lying σ*-MO (weak σ-bond). The β-hydride transfer is a competitive reaction to boron metathesis for the substrates having low lying π*-MO. The relative stability of the products is controlled by the strength of both Fe-E and B-X bonds.

References:

1) De, S.; Parameswaran, P.; Jemmis, E. D. Inorg. Chem. 2007, 46, 6091.
2) (a) Kays, D. L.; Day, J. K.; Ooi, L.; Aldridge, S. Angew. Chem. Int. Ed. Engl. 2005, 44, 7457. (b) Kays, D. L.; Rossin, A.; Day, J. K.; Ooi, L.; Aldridge, S. Dalton Trans. 2006, 2, 399. (c) Kays, D. L.; Day, J. K.; Aldridge, S.; Harrington, R. W.; Clegg, W. Angew. Chem. Int. Ed. Engl. 2006, 45, 3513.
3) (a) Chauvin, Y. 3824; Angew. Chem. Int. Ed. 2006, 45, 3741. (b) Grubbs, R. H. 3845; Angew. Chem. Int. Ed. 2006, 45, 3760. (c) Schrock, R. R. Angew. Chem. Int. Ed. 2006, 45, 3748.
4) (a) Jemmis, E. D.; Jayasree, E. G. Acc. Chem. Res. 2003, 36, 816. (b) Jemmis, E. D.; Jayasree, E. G.; Parameswaran, P. Chem. Soc. Rev. 2006, 35, 157.
5) M. J. Frisch *et al*, Gaussian, Inc., Pittsburgh, PA, 2003.

Author and Contact: Eluvathingal D. Jemmis
Authors: Susmita De and Pattiyil Parameswaran
Organization: Indian Institute of Science, Bangalore, India, 560 012
Resources: *Squall* (IBM SP3) at MHPCC
Sponsorship: This research was supported by the Board of Research in Nuclear Sciences, Mumbai.

Theater UnderSea Warfare (TUSW)

Carl Holmberg and Robert Dant

The Air Force Research Laboratory Maui High Performance Computing Center (AFRL/MHPCC) maintains a Linux cluster for dedicated Theater UnderSea Warfare (TUSW) use. This TUSW cluster provides high performance computing resources to TUSW users for computationally intensive undersea warfare (USW) simulations. TUSW software has been integrated to demonstrate "reachback" high performance computing during fleet exercises. TUSW resources have been successfully demonstrated during Silent Fury (SF), RIMPAC, and Undersea Dominance (UD) exercises in the Pacific Theater, as well as during several "Local" CTF-12 (Pearl Harbor, Hawaii) training activities.

As a part of the follow-on Theater UnderSea Warfare Initiative (TUSWI), MHPCC was tasked by the Office of Naval Research (ONR) to conduct a literature survey of the products and implementations currently available to determine the feasibility of a global computing grid to enhance the situational awareness of the warfighter.

Research Objectives: The intent of the literature survey of grid solutions is to determine to what degree the products and deployed architectures currently available can cope within a set of operational assumptions regarding the DOD's existing IT environment, while providing significant added value to the warfighter using a global information grid concept (Figure 1). A number of these assumptions are common to the civilian IT arena and are likely addressed by many COTS products, while others are unique to the DOD. In parallel with these assumptions, the survey also considers a number of technical trade-offs..

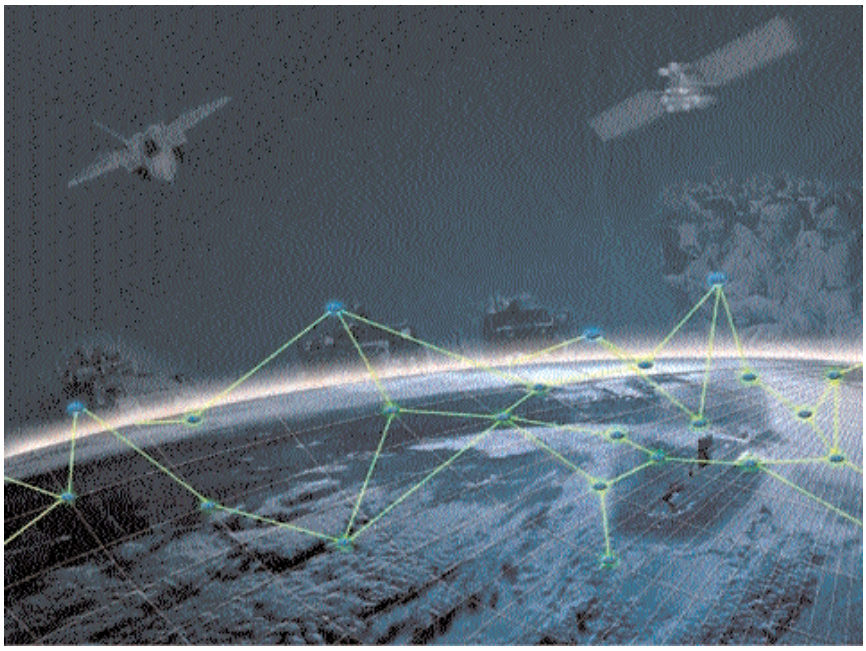

Figure 1. A Global Information Grid Tied to the Warfighter.

Methodology: Material for the survey was collected from literature provided on-line by systems vendors, military and civilian technology research organizations, and existing grid consortiums. Among the vendor sites included were Apple, BEA, HP, IBM, Oracle, Sun, and Veritas. Non-vendor sites included Defense Link, the Globus Consortium, IEEE, Mitre, and the Open Grid Forum. Particular focus was given to products and implementations that were Open Grid Services Architecture (OGSA) or Web Services Resource Framework (WSRF)-compliant. Grids following these reference architectures lend themselves to Service Oriented Architecture (SOA) implementations, and therefore ensure a better fit with the requirements of the DOD's Global Information Grid (GIG).

Results: A number of large and mid-level vendors have released grid-focused product lines. These product lines include complete (or nearly complete) OGSA-compliant suites, software infrastructure components, mobile device tools, and network hardware.

i

Only in the last five years compute grids have begun to transform from glorified clusters into service architectures, but a few case studies are beginning to appear in the literature. Three cases were selected based on their varying degrees of success at implementing a Web Services or SOA application interface. They include: 1) a bioinformatics project, 2) a hyperspectral imaging tool, and 3) the Open Science Grid Project (see Figure 2).

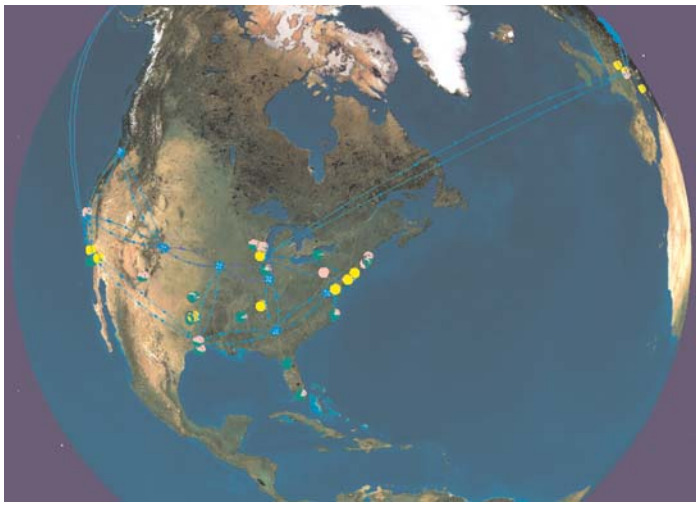

Figure 2. The Open Science Grid Project, Linking 10,000 CPUs around the Globe.

Conclusions/Significance: The products and cases studies described in the survey lead to the conclusion that the components and concepts from which to construct a robust and useful global DOD intergrid are currently available. As they become deployed more widely in the commercial enterprise, we can expect additional grid enabling tools to reach market and, based on the hard-won experience of additional users, a number of best practices to be published. Recommendations included:

- Conduct an in-depth study of the design and operation of the Open Science Grid Project. Discussion with the administrators and primary users will give a better understanding of potential pitfalls and trade-offs in the implementation and use of a large compute grid.
- Develop a comprehensive census of likely applications and data sets to initially be hosted on a DOD grid, potential resource providers, and major users.
- Design and implement a prototype grid to integrate within a selected DOD command, using available networked resources such as those clusters hosted at MHPCC and other DOD computing centers in the Pacific theater (see Figure 3).

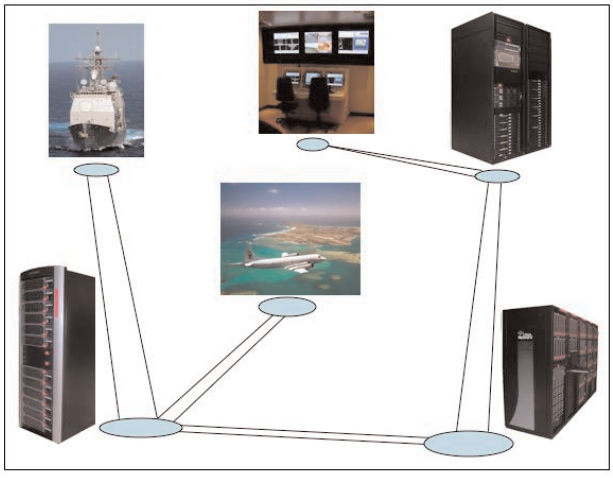

Figure 3. Prototype Grid Concept within a DOD Command.

Author and Contact: Carl Holmberg
Author: Robert Dant
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Resources: TUSW Linux Cluster at MHPCC
Sponsorship: Office of Naval Research

INDEX OF AUTHORS

B

Maxim Ballmer 10
Charles C. Beckner, Jr. 30
M. T. Bettencourt 20
Todd Anthony Bianco 10
Kathy Borelli 2, 30
Matthew Burnham 16

C

Brandoch Calef 30
Guangxia Cao 22
Connor Caples 8
K. L. Cartwright 20
Francis Chun 2

D

Robert Dant 18, 34
Susmita De 33
Bruce Duncan 2

E

Steven Eschrich 24

F

Donald J. Fabozzi 28
L. Neil Frazer 4
Arthur J. Freeman 12
Gabriel I. Font-Rodriguez 8

G

Concetto Giuliano 2
Dmitry Goldgof 24

H

Cordell Hachinsky 8
Shawn Hackett 8
Lawrene O. Hall 24
Nicholas B. Hatcher 12
LuLu Huang 1
Carl Holmberg 18, 34

I

Garrett Ito 10

J

Eluvathingal D. Jemmis 33

K

Vidya Kamath 24
Jerome Karle 1
Rangachar Kasturi 24
Oleg Yu. Kontsevoi 12

L

Seung Ha Lee 26

M

John Mahoney 10
Partha Mandayam 18
P. J. Mardahl 20
Lou Massa 1
Charles L. Matson 2, 30
Maria Murphy 30

N

Dave Norman 18
Eva-Marie Nosal 4

P

Pattiyil Parameswaran 33
Si-Hwan Park 26

R

Kevin Roe 2, 6

S

Paul Schumacher 2
Rong C. Shieh 14
Tom Soo Hoo 30
Duane Stevens 22

V

van Hunen 10
Ron Viloria 30

W

Shannon Wigent 16

Y

Timothy J. Yeatman 24
Shiho You 30