



Army High Performance Computing Research Center

AHP CRC Bulletin



Volume 3 Issue 1

Distribution Statement A: Approved for public release; distribution is unlimited.

Moving Forward: New and Evolving Projects

INSIDE THIS ISSUE


AHP CRC News	2
Education and Outreach	3
Flapping Wing MAVs	5
Materials Modeling	10
HPC On Board	13
The Researchers	18
Project Briefs	19
Publications, Presentations	26

The Army High Performance Computing Research Center, a collaboration between the U.S. Army and a consortium of university and industry partners, develops and applies high performance computing capabilities to address the Army's most difficult scientific and engineering challenges.

AHP CRC also fosters the education of the next generation of scientists and engineers—including those from racially and economically disadvantaged backgrounds—in the fundamental theories and best practices of simulation-based engineering sciences and high performance computing.

AHP CRC consortium members are: Stanford University, High Performance Technologies Inc., Morgan State University, New Mexico State University at Las Cruces, the University of Texas at El Paso, and the NASA Ames Research Center.


<http://www.ahpcrc.org>




AHP CRC research efforts are evolving in response to the Army's top priorities for supporting and protecting soldiers and the stated needs of our colleagues at the Army laboratories. At this stage of the program, several projects are delivering new technologies and capabilities to their Army counterparts, for use in warfighter-directed applications.

AHP CRC researchers are integrating computational fluid dynamics, structural design, algorithm development, and mechanical modeling to develop a flapping-wing micro-aerial vehicle. They use quantum mechanics and genetic algorithms to develop materials for lightweight batteries and phased-array antennas. Power and efficiency studies evaluate on-board HPC systems and advanced image processing applications.

2010 marked the first year for AHP CRC summer interns at ARL and the second year for the AHP CRC Summer Institute. AHP CRC-funded graduate students and postdocs gave live demonstrations of their projects at Supercomputing 2010 and the 27th Army Science Conference.





Clockwise from top left: flapping wing model, Army battery needs, SAR image processing, network configuration.



This issue of the AHP CRC Bulletin features a selection of the research projects that came into their own in 2010. ★

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUL 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011	
4. TITLE AND SUBTITLE AHPCRC (Army High Performance Computing Research Center) Bulletin. Volume 3, Issue 1, July 2011				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Army High Performance Computing Research Center,c/o High Performance Technologies, Inc,11955 Freedom Drive, Suite 1100,Reston,VA,20190-5673				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 32	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			



AHPCRC Researchers Display Their Work

AHPCRC was a very visible presence at Supercomputing 2010 (November 15–18, New Orleans LA) and the 27th Army Science Conference (November 29–December 2, Orlando FL).

Technology demonstrations at Supercomputing included:

- Heterogeneous HPC in Field-Deployable Systems (power vs. quality for graphics processing units in mobile tactical radar, presented by Ricardo Portillo; a comparison of three approaches for programming general-purpose graphics processing units, presented by Yipkei Kwok, both graduate students with Pat Teller, UTEP)
- Field-Deployable and On-Board Multicore Processor Systems (an automated target-finding algorithm, presented by Tomasz Tuzel and Soumik Banerjee, graduate students with Jeanine Cook, NMSU)
- Hybrid Optimization for Parameter Estimation Problems (finding a global minimum in a system with many local minima, presented by Miguel Hernandez IV and Reinaldo Sanchez Arias, graduate students with Leticia Velázquez and Miguel Argáez, UTEP)

The following demonstrations were presented at both conferences:

- Real-Time Finite Element Modeling on a Graphics Processor (rational engineering design with real-time feedback, presented by Cris Cecka, graduate student with Eric Darve at Stanford)
- Image Webs for Collaborative Information Correlation (extracting useful information from large image databases, presented by Omprakash Gnawali and Zixuan Wang, postdoc and graduate student, respectively, with Leonidas Guibas, Stanford)
- Online Reduced-Order Models for Mobile Devices (an iPad app for facilitating decision-making in the field, presented by Mark Potts, senior staff scientist, HPTi, working in collaboration with Charbel Farhat, Stanford) ★

Top two photos: AHPCRC exhibit at Supercomputing 2010:

Tomasz Tuzel (NMSU), Omprakash Gnawali (Stanford), Barbara Bryan (HPTi), and Cris Cecka (Stanford)

Miguel Hernandez, IV and Reinaldo Sanchez Arias (UTEP)

Bottom two photos: AHPCRC exhibit at the 27th Army Science Conference:

Cris Cecka (Stanford) demonstrates finite element modeling using a GPU to visiting middle school students (Army RDECOM photo)

Dr. Marilyn M. Freeman, Army Deputy Assistant Secretary for Research and Technology, discusses the GPU demonstration with Cris Cecka.

(All photos by Nancy McGuire, HPTi, unless otherwise noted.)

AHPCRC Supports NMSU PREP Summer Program

New Mexico State University's Pre-Freshman Engineering Program (PREP) celebrated its 14th year in the summer of 2010, with a record 176 students completing the program. AHPCRC is a key funding agency for PREP, and through AHPCRC's commitment to excellence and generous financial support, PREP has continued to grow.

PREP is administered through the New Mexico Alliance for Minority Participation for the NMSU College of Engineering. This program recruits achieving pre-college students from the three school districts in Doña Ana County for a six-week, academically intense summer program with the goal of preparing these students for careers in science, technology, engineering, and mathematics (STEM). Students take courses in logic, algebraic structures, technical writing, engineering, computer science, and physics. The goals are to stimulate participants' interest in higher mathematics and science and to provide problem-solving sessions to equip them with the necessary tools and the desire to pursue a career in STEM.

Friday field trips and Career Awareness Seminars provide students with opportunities to meet and interact with professionals who instill the vision and passion to become the scientific leaders of tomorrow. The participants may begin the program as early as sixth grade and attend for four years prior to high school graduation. Although PREP is open to everyone, the program focus is on female and minority populations traditionally underrepresented in the STEM fields.



PREP 4, which offers college-credit courses to students in their fourth year of PREP, completed its second year. PREP 4 students toured the White Sands Test Facility's new state-of-the-art Range Launch Complex control room and control tower.

They also learned about robotics and programming while assembling and programming Boe-Bot robot kits.



Article and photographs provided by Erica Voges, program director.

PREP 3 students built solar cars and ran experiments to learn about the advantages of solar power. PREP 2 students toured and received briefings at Holloman Air Force Base (Alamogordo, NM). While there, the students worked with the Explosive Ordnance Devices Division, the High Speed Track, the T-38 Aircraft Training Facility, and Heritage Park. They spoke with two female pilots about what it's like to be a pilot in the Air Force. PREP 1 and 2 students designed, built and launched multiple single and double-stage rockets. PREP 1, 2 and 3 students interacted with guest lecturers Dr. Stephen Kanim (Physics Professor); Dr. David Voelz (Electrical Engineering Professor); and Dr. Ricardo Jacquez (a civil engineer and Dean of the College of Engineering) during the Career Awareness component for PREP.

PREP 1, 2, 3, and 4 students visited the International Space Museum, which educates visitors from around the world on the history, science, and technology of space. During their visit, they observed NASA technology and multiple rocket launches and took part in a "physics magic show." In addition, they worked with computers to learn about basic hardware and software components, development of algorithms through flowcharts, BASIC programming, Visual C++, Web Design, Microsoft Office, and MatLab.

On the campus of NMSU, students viewed the large wind tunnel that is used for research by the Mechanical and Aeronautical Engineering Departments. A guest speaker from the Mechanical Engineering Department spoke to all the PREP students about the Unmanned Aerial Systems Technical Analysis and Applications Center designed to promote safe integration of the unmanned systems in the National Airspace System. ★

more on page 4

Education and Outreach

continued from page 3

AHPCRC Summer Institute 2010

The AHPCRC Summer Institute marked the completion of its second successful season at Stanford University on August 13, 2010, with a morning of student project presentations. After a welcome by AHPCRC Center Director Charbel Farhat and introductory remarks by Army Research Laboratory Director John Miller, each student gave a 15-minute presentation on his or her research to an audience of ARL representatives, Summer Institute mentors, and their fellow students. AHPCRC ARL Cooperative Agreement Manager Raju Namburu; ARL Minority Outreach Program Manager Vallen Emery, Jr.; and John Miller offered closing remarks and presented awards. Students, mentors, and research topics are as follows:

- Matthew Zahr, UC Berkeley (mentors: David Amsalem, Kevin Carlberg, Charbel Farhat). “Comparison of Model Reduction Techniques on High-Fidelity Linear and Nonlinear Electrical, Mechanical, and Biological Systems”
- Xiao Ying Zhao, Stanford University (mentors: David Powell and Charbel Farhat). “The Effect of Environmental Degradation on the Ballistic Resistance of High Strength Fabric”
- Jennifer Kuchle, University of Texas, El Paso, and Eduardo Vega, New Mexico State University (mentors: Charbel Bou-Mosleh and Charbel Farhat). “Aerodynamic Analysis of the NASA Common Research Model (CRM) Wing-Body-Tail”
- Oscar Octavio Torres-Olague, New Mexico State University (mentors: Ramsharan Rangarajan, Raymond Ryckman, Pablo Mata Almonacid, and Adrian Lew). “Simulation of Ballistics Gel Penetration under Axisymmetric Conditions”
- Vivian Nguyen and Juan Pablo Samper Mejia, Stanford University (mentors: Cris Cecka and Eric Darve). “Real Time Finite Element Analysis of Dynamic Problems Using GPUs”

- Joshua McCartney, University of Texas, El Paso, and Sabin Pokharel, Morgan State University (mentors: Sylvie Aubry and Wei Cai). “Extension of Dislocation Dynamics for Semi-conductor Materials.”
- Prakash Sharma, Morgan State University (mentors: Omprakash Gnawali, Kyle Heath and Leonidas Guibas). “Evaluation of RASL Image Alignment Algorithm”
- Esthela Gallardo and Edgar Caballero, University of Texas, El Paso (mentors: Nick Henderson and Walter Murray). “Visualizing Iterates of Optimization Algorithms”
- Greg Romero and Rovshan Rustamov, New Mexico State University (mentors: Jared Casper and Kunle Olukotun). “Commodity CPUs and a Tightly-Coupled FPGA” ★

By the Numbers

The AHPCRC Summer Institute for Undergraduates is an 8-week program, held annually at Stanford University since 2009. Each year, 14–16 students from several universities work with 17–21 Stanford professors, research associates, postdocs, and graduate students, who serve as instructors and mentors.

Of the 16 2009–2010 students who returned follow-up surveys:

- 8 have interned or will intern at ARL in summer 2010 and 2011
- 1 is employed at ARL at the White Sands Missile Range in New Mexico
- All 5 students who have graduated are either attending graduate school or have applied for Fall 2011
- 10 of the 11 who have not yet graduated state that participating in the program made it more likely to apply to graduate school
- 6 students stated that they are more likely to pursue computational science
- 2 of the 2010 students plan to apply for Army SMART scholarships in 2011 toward their PhDs

Simulations of Flapping Wings

Recent advances in reconnaissance and surveillance technology have included means of accessing confined or hazardous spaces while minimizing risk to soldiers. Micro-aerial vehicles (MAVs) carrying cameras, sensors, and communications devices show great potential toward meeting these criteria, but they also present a host of design problems not encountered by engineers who design more conventionally-sized aircraft. Because small flyers must provide proportionally more thrust than their larger counterparts to overcome the air's viscosity, they must either be equipped with more powerful propulsion systems (requiring more weight and fuel) or they must move their wings to produce this thrust, in much the same fashion as insects and small birds. MAVs with flexible flapping wings promise greater stability and maneuverability than their fixed-wing or rotary-winged counterparts, but the design of these wings is at a much earlier stage than for the other technologies.

Compared with the traditional approach of designing aircraft with rigid structures and for steady aerodynamics, designing MAVs to exploit flexibility

The Researchers

AHPCRC researchers are collaborating to develop the efficient numerical simulation capabilities needed to analyze and optimize air flow past flapping-wing MAVs in hover or low-speed forward flight:

NASA Advanced Supercomputing Division (Ames Research Center)

Terry Holst, Tom Pulliam, Piyush Mehrotra, Dennis Jespersen, Steve Heistand

Stanford University


Professors Charbel Farhat and Antony Jameson
Graduate students Matthew Culbreth and Joshua Leffell

New Mexico State University

Professors Mingjun Wei, Banavara Shashikanth, Fangjun Shu

They are joined in this effort by Stanford professors Walter Murray and Michael Saunders, who are developing the optimization algorithms needed for improved efficiency and reduced computational costs.

To read more about the MAV work at Stanford and NMSU, see page 9 of this issue and *AHPCRC Bulletin* Vol. 1 No. 3 (www.ahpcrc.org/publications.html).



Grid setup for modeling interactions between the wing and the surrounding air flow. Graphic courtesy of Charbel Farhat, Stanford.

and unsteady aerodynamics will be very difficult. Very few of the efficient computational design tools used for large aircraft design can be used in the unsteady, low Reynolds number (viscous flow) regime, and designers must fall back on costly unsteady numerical flow simulations and experiments as the primary design tools. Addressing this issue requires the development of efficient, physically accurate massively parallel computational fluid dynamics (CFD) tools that incorporate aeroelastic effects and large motions associated with flapping wings. The design degrees of freedom increase significantly when considering a flexible wing in a generalized periodic flapping motion. Although much has been learned from observing the flight of birds and insects, it is still far from clear how to couple wing flexibility and flapping motion in an optimal way for a given flight performance metric.


Analyzing Flapping Wing Flows

The NASA group is working to improve the simulation accuracy associated with the analysis of flapping-wing MAVs in hover or low-speed forward flight, using NASA's OVERFLOW Navier-Stokes flow solver. This flow solver has been developed extensively over the past two decades for three-dimensional applications. The group is placing paramount importance on getting the best tradeoff between accuracy and efficiency.

continued on page 6

Flapping Wings

continued from page 5



Time sequence simulation of flapping fruit fly wings.
Left column: downstroke. Right column: upstroke.
Graphic courtesy of Terry Holst, NASA Ames.

Because the Reynolds number of MAV flows is small (air flow is more laminar than turbulent), the effect of turbulence and transition will not be studied in the context of the effort. The first simulations addressed an airfoil in pitching-plunging motion, followed by wings in rigid body motion—sinusoidal flapping with

a superimposed twisting motion. General motion of a flapping-deforming wing will be studied using the shape discretization procedure that is developed by the Stanford group (*illustration, previous page*). Various flapping-wing geometry discretization methodologies are being assessed, with emphasis on achieving a fully defined design space while minimizing the number of decision variables.

A head-to-head comparison with the AERO code used by the Stanford group showed good agreement for calculations of lift and drag as functions of time. The OVERFLOW code has been tested on several high performance computers to see how varying the machine architecture and number of cores influences the run time of the code. Results showed that the code scaled linearly over the range of cores tested; for all machines tested, increasing the number of cores used for the calculations produced a corresponding speed-up in run times, which is a desired outcome.

Optimizing Flapping Wing Flows

Coupling wing flexibility and flapping motion in an optimal way for a given flight performance metric remains a poorly understood problem. This motivates the use of numerical optimization techniques coupled with unsteady flow simulations to obtain the periodic wing motions and deformations that best suit different types of flight regimes such as hovering and forward flight. A 3D unsteady viscous flow solution for a wing requires on the order of 10 CPU hours. Flapping wing optimization will require thousands or tens of thousands of these flow solutions, making the task essentially infeasible without massively parallel algorithms and hardware. In addition, appropriate objective functions for flapping flight are not as clear as for the steady case. The lift, thrust, and required power are some of the metrics of interest. Different modes of flight also apply, including hovering, steady forward flight and maneuvering. Few studies of flapping optimization have been reported in literature. The problem of optimal generalized trajectories and deformations for 3D flapping have not been addressed. The Stanford research group is developing 3D unsteady optimization results, starting with optimizations of rigid motions with increasing degrees of freedom.

Additional degrees of freedom are introduced to parameterize deformations to the flapping wing surface. A new 3D unsteady flow solver has been tailored specifically for simulations of flapping wings, specifically in its ability to accommodate moving meshes. These and other revisions to the computational methods have accelerated the solution of steady-state problems by a factor of 60, and unsteady-state problems by a factor of 20 over previous implementations.

Advanced Optimization Algorithms and Software

At the heart of massively parallel MAV flight simulation programs are algorithms that solve thousands of partial differential equations (PDEs) describing wing motion and the resulting air flow patterns. A simulation's demand for computer resources and the amount of time it takes to execute depend strongly on the ability of each algorithm to arrive at accurate results in the least amount of time.

Murray and Saunders are developing such algorithms to assist in finding optimal wing shapes and motions for these small flying vehicles, in collaboration with Jameson, Culbreth, and several of their own Institute for Computational and Mathematical Engineering (iCME) students. They are using the Stanford-developed PDE-solver AERO-F and the large-scale optimization solvers SNOPT and SQOPT. Their intention is to generate a version of AERO-F that can be used efficiently within an optimization algorithm and that can be used to estimate derivatives efficiently using the finite differences method. They are also developing a significantly improved version of SNOPT that can link with AERO-F for the solution of complex problems, including those encountered in the design of MAVs.

Currently, the optimization solver sits on top of the AERO-F PDE solver and makes repeated calls with different values of the optimization parameters (which may be functions rather than parameters). High performance computing (HPC) is essential for solving the PDE many times. Indeed, the number of calls to the PDE solver is so large and the cost of a single



Flexible flapping wing form. Graphic courtesy of Charbel Farhat, Stanford.

call so high that an accurate PDE solution may not be practical for each call. Thus, the first objective is to improve the optimization solver SNOPT so that it requires fewer calls to AERO-F.

Discretization

One key focus area is the transformation of the original problem of finding the shape of a wing and its location in time into an optimization problem in a finite set of variables. There are many approaches, each affecting the efficiency in terms of the number of PDE solutions required by the optimization algorithm. Of particular importance is that problems be well scaled. Also, different schemes introduce different structures within the optimization. Typically, the simpler schemes induce simple structures but require more variables, which usually reduces efficiency.

For example, it may be better to discretize motion differently from the shape of the wing. Moreover, since structure varies with time, it may be better for structure and motion variables to be perturbations of a single structure and a single path.

Global optimizers

Given the variety of wing shapes and motions in nature, it seems likely that the flapping wing problem is non-convex and has multiple local minimizers.

continued on page 8

Flapping Wings

continued from page 7

(Some examples of non-convex functions are wavy lines and curves that double back on themselves.) At present, efficient software exists only for finding local minimizers. Using multiple initial estimates is a simple approach to finding global (or good local) minimizers, but the problem can grow very large as the number of variables increases.

An alternative is to re-frame the problem in terms of a mathematical function that is easier to work with. In mathematical parlance, a homotopy method is used, in which a strictly convex function is added to the problem. Such methods serve a dual purpose, because the resulting problem is better conditioned and hence easier to solve. This keeps the problem to a reasonable size and rules out poor local minimizers.

Another issue is how to regularize the problem should it prove to be ill-conditioned; that is, limitations in precision or small errors in the data cause the calculations to fail or produce large errors in the solution. Regularizing an ill-conditioned problem typically involves some type of reformulation or adding additional assumptions (for example, setting upper and lower limits on some parameters). Real-world problems contain subsets of variables that are correlated with each other (they increase or decrease together), and that feature must be incorporated within the model formulation. This restricts the search space and enhances the efficiency of the optimizer. Such modifications also improve the condition of the problem, further contributing to efficiency.

Typically, there is little point in solving a problem to greater accuracy than the system models the real world. However, within an optimization algorithm, the accuracy of the fit to the model is not as important as the consistency of the error: It is better to make a larger consistent error than an irregular but smaller error.

An interesting feature of this sequence of problems is that the solution of an early problem need have no relationship to the real model. What is important

Using Advanced Architectures Efficiently


After nearly a decade of per-node computing performance being driven by steadily increasing CPU clock-speeds, energy consumption and other factors are forcing CPU makers to look for other means to improve performance. Clock-rates have essentially flat-lined, but the number of transistors per processor has continued to increase following Moore's law, and manufacturers have been using them to create additional CPU cores in the same substrate. Meanwhile, gains in memory, cache and I/O performance for CFD codes have nearly stagnated. Currently, processors with six cores are widely available, and chips featuring many tens (and even hundreds) of cores are in the offing. In such processors, the cores share the same cache, memory, and network bandwidth as used to be dedicated to a single core. Thus, the high performance computing (HPC) industry has seen a radical decrease in the amount of bandwidth available to each core. This decrease promises severe consequences for the performance of traditional computational fluid dynamics (CFD) algorithms, because bandwidth has been the traditional bottleneck for CFD algorithms on parallel HPC hardware. Thus, it is vital to maximize the efficiency of each analysis run. The NASA research group is exploring several computer science aspects of CFD code execution that could provide a significant reduction in computational cost, including algorithms with a better native fit to the strongly hierarchical memory access of modern compute platforms. In addition, the NASA group is working to improve execution efficiency for the optimization portion of this effort.

is the relationship of its solution to that of the next discretization.

Continuing work for each of these projects focuses on adding a greater degree of sophistication to the models, making the calculations more efficient, and verifying the various modeling methods against each other and against experimental data from mechanical models. ★

Recent Successes:

Prof. Charbel Farhat and co-workers at Stanford University are developing a state-of-the-art computational fluid dynamics (CFD)-based coupled fluid–structure analysis capability for use with high performance computing (HPC) resources. They are determining whether it is better to flap a wing in pure plunge, pitch, or twist motion or to use combined motions in order to provide additional thrust. They are also focusing on determining the optimal amplitudes and frequencies of the flapping and twisting motions. This capability, which features an advanced embedded boundary method for viscous CFD, is being built into Stanford’s multidisciplinary code AERO, which was demonstrated for the Vehicle Technology Directorate (VTD) of ARL at the beginning of 2010 in order to spur technical collaboration. A first release of this upgraded code featuring the new embedded viscous CFD method will be delivered to ARL scientists, together with a report discussing its impact on the simulation of flexible flapping wings in the low Reynolds number regime. This Fluid–Structure Computational Technology developed for Flapping Wings has been adopted by Boeing for the analysis of High Altitude Long Endurance (HALE) Systems.




CFD simulation for flapping wing using Stanford’s AERO code.

NMSU FlexSI simulation of flexible flapping wings in forward flight. Fluid flow is marked by vortex shedding; inset shows detailed structural deformation.

Prof. Mingjun Wei and co-workers at New Mexico State University are developing numerical algorithms, incorporated into their FlexSI code, that exploit HPC technologies to simulate flapping and twisting aeroelastic wings with fully coupled interaction between fluid flow and wing structure. They are investigating the influence of structural flexibility on the active and passive motions of flexible wings in plunging, pitching, twisting, and root-flapping motions to understand the mechanisms and maximize the propulsive efficiency. To validate their models, they are conducting laboratory experiments and theoretical modeling concurrently. FlexSI solves the whole problem monolithically and reduces the computational cost tremendously by avoiding iterations commonly existing in fluid–structure interaction problems. NMSU FlexSI code has recently been extended to three dimensions, and the parallel version is in development. Recently, the NMSU group has begun to study the effects of wind gusts on the leading and trailing edge vortices, which are related to hovering stability.

The NMSU group is making experimental measurements using robotic models with rigid and flexible wings to simulate hovering and forward-flying birds. Fluid–structure interactions are being investigated in a wind tunnel, a water channel, and an oil tank.



Multiscale Modeling of Materials

The rotating reflector antenna associated with airport traffic control systems is giving way in some applications to a newer technology called the phased array antenna system (sometimes called a beamformer, *example shown at right*). Thousands of closely spaced, individual radiating elements produce a composite beam that can be shaped and directed electronically in microseconds, enabling it to track hundreds of targets simultaneously. Using a phased array, one radar system can be used for both missile guidance and target detection/tracking functions, rather than requiring separate dedicated systems, and no moving parts are required.

Each antenna element requires its own associated radio frequency (RF) phase shifter. The cost of the phase shifters is a limiting factor to the adoption of phased array antenna technology: the price of a system with thousands of elements can be prohibitive.

Phase shifters may be made from tunable dielectric materials, including the ferroelectric material barium titanate (BTO, or BaTiO_3). The magnitude and direction of electrical polarization in ferroelectric materials can be manipulated using an applied electrical field. (*See box, next page.*) Thus, the relative phases of the respective signals feeding the antennas can be varied so as to reinforce or suppress the effective radiation pattern of the array in specific directions.

Eric Darve, mechanical engineering professor at Stanford University, and his students, are developing new modeling methods in order to facilitate the rational design and evaluation of ferroelectrics and related materials. These methods combine a modeling component at the atomistic scale and a numerical component, with techniques to solve linear systems arising from finite-element and finite-volume analyses at the macroscopic scale.

Perovskite-type crystal structure typical of barium titanate. Atoms at corners of cube: barium; white atom: titanium; red atoms: oxygen. The five atoms that form the unit basis are shown in the inset. (Graphic courtesy of José Solomon, Stanford.)




The Cobra Judy phased-array radar system on the missile range instrumentation ship USNS Observation Island. (Wikimedia Commons)

Material models

Molecular dynamics (MD) methods can be used to model BTO crystals, using existing computational force fields. However, the existing force fields and parameterization cannot accurately model the crystal in different solid phases at various temperatures, making it impossible to predict accurately the most stable phase at a given temperature. In addition, no force field is currently capable of modeling defects in the crystal. This significantly restricts the ability to model and predict the behavior of BTO.

Developing force fields is a challenging task; Darve's group has created a unique set of tools to address this challenge. The first goal of the project is to apply genetic algorithm (GA) approaches to create a novel parameterization of the shell model potential as applied to BTO perovskite-structured crystals (*illustrations, this page and next page, and explanation, next page*).




By re-parameterizing the shell model potential to tailor it specifically for BTO, the behavior of the crystal can be modeled with a high degree of fidelity under a diverse range of physical and electrical loading conditions. Predicting the behavior and response of materials often requires coupling atomistic and macroscopic models—for example, quantum models, molecular dynamics, Monte Carlo, and finite-element analysis. The resulting mathematical systems can be difficult to solve. The geometry of the domain is often very complicated.

In order to derive the BTO-specific parameterization, density functional theory (DFT) analysis is coupled with a GA-based technique by which the multidimensional parameter space can be explored efficiently. First, DFT is used to perform a single-point energy calculation on numerous deformed configurations of BTO's unit basis (*figure, below left*). The numerous configurations examined with DFT are various deformations from the stable state of the perovskite structure in each of its four crystalline phases, (*cubic, tetragonal, orthorhombic, and rhombohedral; see figure, page 12*). These calculations provide various energy configuration curves, which the GA then uses as a reference database.

The GA uses evolutionary algorithms derived loosely from Darwinian concepts. The starting point is a population of parameter sets for the functional form of the potential. The potential is evaluated using each population member (i.e., parameter set), and the resulting difference between the energy produced from the potential and that of the reference DFT database is used to establish the fitness of the given member. Members of a population are consequently combined, using either crossover or mutation techniques, with the objective of creating offspring with increased fitness. As the GA produces generation after generation, the algorithm strives to produce parameter sets with ever-greater fidelity to the quantum energy calculations.

Validating the Technique

In order to validate their technique, as well as optimize the selection of which crystal geometries to explore with the DFT, a series of synthetic energy curves was



Ferroelectric Materials

Ferroelectric materials exhibit a nonlinear response to an applied electrical field. That is, the electrical polarization of a ferroelectric material increases or decreases suddenly at a particular electrical field strength. This transition point is characteristic of the specific material. Slight alterations in material composition or other characteristics can “tune” the transition point to a desired electrical field strength.

Additionally, ferroelectrics exhibit hysteresis—the transition point in an increasing applied field is not the same as the transition point for a decreasing applied field. This property is useful for application in memory devices.

The polarization effects in BTO arise from the distortion of the titanium oxide sublattice due to the large size of the barium ions that occupy the large cavities in the lattice (*silver atoms in the figure above*). The titanium atoms are forced off-center in the TiO_6 octahedra (*blue and red atoms in the figure*), producing an uneven distribution of electrical charges, creating a dipole effect. Applying pressure forces the titanium atoms back toward the centers of the octahedra, reducing the dipole strength. This sensitivity to pressure is called piezoelectricity.

Shell Model Potential

The shell model describes the deformation of the electronic structure of an ion (electrically non-neutral atom) as a result of interactions with other atoms. Each atom in a solid is described in terms of a massive core and a massless shell. Core-shell displacements create dipole moments (uneven distribution of electrical charges) that serve as a means of storing energy (potential).

produced from a known shell model parameter set for BTO. The researchers hoped to reproduce the original parameter set that was used to create these synthetic curves, and thus have a direct quantification of the accuracy and effectiveness of the technique. A series of two-ion displacements was performed on the crystal

continued on page 12

Multiscale Modeling of Materials


continued from page 11

unit basis of the lattice to determine the symmetries inherent to the geometry. Each ion was displaced from $-0.5 \text{ \AA}$ to $+0.5 \text{ \AA}$ from its empirically determined stable configuration in the cubic phase.

Ninety displacements were compared, comprising a total of ten two-ion combinations displaced in nine orthogonal symmetries. A subset of 18 displacements was identified that covers the entire set of displacements. Energy curves produced from these 18 displacement configurations were fed to the GA. This is a fairly stringent test, because fitting the GA force field parameters to this data set guarantees that all thermal fluctuations of the system around the equilibrium point are reproduced correctly by the force field. Using this geometry, a series of evolutions was performed on the Buckingham parameters A and ρ . The agreement between the exact and evolved parameter sets was close over a wide energy range. The general results of the current GA analysis are satisfactory and served as a general validation of this approach.

In parallel to efforts to optimize the GA technique, various energy configuration curves were created for use with the novel parameter derivations. Using the DFT code Abinit, the same coupled displacements that were explored in the synthetic analysis are being reproduced within the quantum framework. In fact, a key benefit of the synthetic study was to identify equivalent displacement configurations so as to minimize the computational overhead required in the fitting procedure with the quantum data.

Darve's group has begun to implement a "hybrid" technique, by which a steepest descent optimization method (such as the conjugate gradient, CG) is coupled to the GA algorithm. In this novel approach, the GA will begin a partial evolution process followed by application of CG iterations. By iterating this procedure, they hope to improve the accuracy of GA still further.



Crystal phases for the perovskite-type crystal structure typical of barium titanate (clockwise from top left): cubic, orthorhombic, rhombohedral, and tetragonal.

Moving Ahead

The initial objective of the force field development is to model the phase transition cycle of the crystal lattice as the BTO is heated thermally, capturing the transition temperatures between phases with a high degree of fidelity to empirical results. Experimental measurements are available to verify and validate the code. In addition, they will tune the force field to reproduce properly the formation of an oxygen vacancy, which is of pivotal importance. This allows for MD simulations of oxygen vacancy diffusion in the crystal. Further objectives include the modeling of dislocation energies of all three atomic species in the perovskite, along with the ability of an MD calculation to determine the dielectric constant of the crystal in response to an applied electric field.

Army researchers familiar with this work have complimented the quality of the math and science behind it. The electronic potentials developed here are of particular interest, and will assist in developing more cost-effective phased-array antenna components. ★

HPC in Image Processing: Putting the Technology on Board

Size, weight, and power constraints greatly limit the amount of computer hardware that can be carried on an aircraft or land vehicle. This, in turn, limits the extent of image processing that can be performed on board. Often, less compute-intensive algorithms are used at the expense of design flexibility and image quality. This has changed in recent years with the advent of small, powerful terascale processing units, such as graphics processing units (GPUs), which are more capable of executing compute-intensive algorithms.

Hardware accelerators, including GPUs, field-programmable gate arrays (FPGA), and solid state devices, are gaining acceptance in on-board systems because of their efficiency in performing the repetitive, specialized tasks typical of radar processing and machine vision, freeing up general-purpose central processing units (CPUs) to concentrate on the data-dependent control operations at which they excel.

Combining two or more specialized hardware components to create a heterogeneous high performance computing (HHPC) system enables even further advances in processing capability. Greater computational power, decreased size, and a focus on reduced power consumption makes it possible to integrate CPUs and hardware accelerators into small systems with enough compute capability to execute critical military applications in the field, in real or near-real time.

Under the AHPCRC program, a group at the University of Texas at El Paso (UTEP) led by Pat Teller and Sarala Arunagiri and Jeanine Cook's group at New Mexico State University (NMSU) are working to evaluate the capabilities of HHPC systems and how they can be used in the field. Currently, the UTEP group is evaluating the precision and power consumption characteristics of various types and levels of image processing calculation algorithms for HHPC systems, in particular, multi-core/GPGPU systems. Meanwhile, the NMSU group is assessing the capa-



Synthetic-aperture radar processes and integrates multiple conventional radar images to produce one high-resolution image. (NASA JPL graphic)

bilities and field applications of HHPC systems that use FPGAs. Both groups are collaborating with researchers at the Army Research Laboratories (ARL) at Adelphi and Aberdeen, MD. The UTEP group has also started a collaboration with researchers at ARL-White Sands, on applications of interest to the Army.

On-Board and Field-Deployable Systems

To determine the appropriate architecture of a field-deployable system, constraints in power, execution time, accuracy, size, and weight must be considered. Many military-related applications depend on real-time production of results, requiring a consideration of operating system performance in terms of both execution time and power.

Determining the best hardware and software configurations for on-board systems requires that several questions be addressed. What are the characteristics of these applications with respect to memory footprint, cache behavior, execution unit utilization, and number of FLOPS (floating-point operations per second)? What performance trade-offs (execution time, power, size, weight, precision) are associated with implementing the same application on different architectures? Can mathematical accuracy and/or precision be

continued on page 14

Field-Deployable Systems

continued from page 13

modified to increase performance and decrease energy dissipation? What power characteristics are associated with various accelerator devices, particularly FPGAs and GPUs? What are the characteristics of power consumption at the instruction level? Can low-power instructions be used without significant performance loss? What application characteristics determine optimal performance on different architectures? What execution phases of an application map best to a given architecture? How do the performance and power consumption of various real-time operating systems differ?

By answering these questions, proposed systems can optimize the constraints for Army applications and potentially drive future development of such systems. Both research groups are answering these questions primarily through experimentation and measurement, supplemented with modeling. They use existing and newly developed techniques to answer questions about execution time and power performance. Initial experiments are being done on the AHPCRC-sponsored Chimera heterogeneous machine at the University of Texas at El Paso.

Performance Measurement

Cook's group is analyzing the performance of Army-relevant algorithms and applications to determine the potential benefit of implementation on FPGAs in both serial and parallel processing contexts. The tools used to measure performance in an HHP system differ from those traditionally used to measure CPU performance. Typically, an application programming interface (API) is used in a CPU to access on-chip performance counters to collect data on total cycle counts, instruction counts, cache miss behavior, and branch prediction accuracy. GPUs also generally have user-accessible performance counters to evaluate performance. FPGA development environments typically include an accurate simulator to evaluate design performance.

Measuring power consumption in HHP systems also

Synthetic Aperture Radar

Computationally intensive image processing is required by synthetic aperture radar (SAR), which uses radar detectors mounted on an aircraft or land vehicle to collect a series of low-resolution images. These images are computationally formed and integrated, or back-projected, to produce a final high-resolution image. This process produces images that resemble images that might have been produced by a much larger single-aperture device.

Backprojection-based image formation algorithms for SAR and other types of radar systems may be able to take advantage of emerging compute technologies to provide faster and more power efficient real-time or near real-time solutions. Field-deployable radar systems are often self-contained and must run for long periods of time without any connection to an external power source. Therefore, FPGAs and GPUs may enable a better optimization of power and performance than traditional CPU-based systems, particularly for systems to be deployed in the field.

requires a suite of tools for the various technologies on which the application is distributed. Because a CPU will be used in the final mobile system, it is necessary to characterize and optimize CPU power consumption. Cook's group has implemented a testbed that includes a data logging machine, a digital multimeter, and a clamp-on ammeter to measure the dynamic power consumption of codes executing on CPUs and GPUs. Accurate measurements of CPU power is also done using performance counters in conjunction with analytic models. These models can also be used to make real-time scheduling decisions to reduce power consumption.

FPGA power is typically estimated and optimized using vendor tools that are generally integrated into the development environment for a particular device.

Performance Modeling

Previously, Cook's group has worked to develop and enhance a Monte Carlo processor modeling (MCPM) technique. They have several very accurate single- and


multicore models of contemporary processors, including the IBM Cell BE, Intel Itanium 2, Sun Niagara 1 and 2, and AMD Opteron. The Cell BE and Opteron models are used to predict and understand the performance of Army applications. Most of these models have been released through open-source licensing.

Using models is faster than traditional performance analysis using cycle-accurate simulators. The architectural components in the models can easily be changed through a configuration file, making it possible to study the performance effects of component enhancements. Using the models will improve understanding of the application-to-architecture mapping of different portions or phases of the applications. The Monte Carlo processor modeling approach will eventually be extended to include power models for predicting power consumption and energy dissipation.

Cook's group is developing a Monte Carlo processor model of a GPU. This is a bit more difficult than modeling a traditional CPU because of the limited availability and capability of existing performance tools. However, such a model will be very useful in determining an optimal GPU architecture and application-to-architecture mapping for Army-relevant applications, including the SAR backprojection algorithm discussed below.

SAR Backprojection Application

Army researchers have expressed an interest in applying this work to compute-intensive image processing applications, including synthetic aperture radar (SAR) backprojection. (*See box, Synthetic Aperture Radar, on previous page.*) An ARL-developed backprojection algorithm has been implemented on a Xilinx Virtex II FPGA development board that is interfaced to DRAM (dynamic random-access memory). Single-precision floating point and integer versions have been implemented. The current FPGA implementation of the ARL backprojection algorithm is very area-efficient: it occupies only a small portion of a relatively small FPGA. However, very irregular memory access patterns significantly degrade the execution time performance. Cook's group is currently working toward



Comparison of single- and double-precision synthetic-aperture radar images. (Courtesy of Pat Teller, UTEP.)

porting this implementation to an FPGA board that interfaces to faster memory.

A parallel implementation of the ARL backprojection algorithm is currently being developed and implemented on an FPGA. Alternative serial algorithms are being evaluated to determine their performance and the efficiency of parallel implementation.

Matched filtering, a signal reconstruction algorithm that extracts the echoed radar signal from the return signal, efficiently removes white noise from the signal to produce the echo signal. This signal is then sent to the processing algorithms before it is sent to the backprojection core. Cook's group has implemented a matched filter on a Tesla GPU, and they are measuring its performance. An implementation of matched filtering has also been employed on an FPGA, and the two implementations are currently being analyzed and compared to determine the performance and power benefits of each.

Cook's group is also currently studying image construction algorithms for various types of radar systems such as SIRE (Synchronous Impulse Reconstruction) in order to develop a more generic back-end imaging platform that can be integrated with various front-end radar systems.

continued on page 16

Field-Deployable Systems

continued from page 15

Precision Requirements

The UTEP group is evaluating trade-offs between precision and power/energy consumption by evaluating the relative power/energy consumption and quality of SAR images produced using the less energy-demanding single-precision computations, compared with double-precision image formation (IF). A preliminary study evaluated Fourier (frequency domain)-based and backprojection (time domain)-based SAR IF techniques on a CPU using codes written in the C programming language. These techniques differ in computational intensity and potential radar imaging capability.

This research was conducted in simple experimental and power measurement environments with a set of image quality metrics. (SAR image quality can be judged automatically using established image comparison metrics.) Results of the first phase of this research showed that image quality for single-precision IF often is comparable to that for double-precision IF. This facilitates the implementation of a power manager that takes image quality constraints into account. For this experimental environment, even though single precision did not offer any power benefits over double precision, it consistently reduced IF execution time, reducing total energy consumption by 14–51% as compared to double precision.


The second phase of this research, currently in progress, targets a high-end GPU testbed that resembles an embedded, field-deployable, HHPC environment. The testbed is set up for fine-grained direct measurement of GPU power to investigate the effect of algorithm

mic changes to GPU programs on power and energy consumption. The tradeoffs between power and energy consumption and SAR output quality are being evaluated for two GPU SAR backprojection implementations, which are often run on field-deployable systems and can benefit from HPC. The first, OSUBP, is a publicly available code that is based on AFRL-developed code and processes a realistic data set provided by AFRL. The second, SIRE/RSM, was developed by ARL and processes a simulated data set provided by the staff at the Adelphi Laboratory Center. The output quality metrics include common image processing metrics such as Peak Signal to Noise Ratio (PSNR), Image Fidelity (IF), and Mean Structural Similarity (MSSIM), as well as more radar-centric metrics such as Impulse Response, Impulse Response Width Resolution (IWR), Peak to Sidelobe Ratio (PSLR), and Integrated Sidelobe Ratio (ISLR). (*For more information, see references 1 and 2.*)

Stereo Vision

On-board image processing can be applied to other image processing applications as well. For example, stereo correspondence has traditionally been, and continues to be, one of the most heavily investigated topics in computer vision, and many algorithms for stereo correspondence have been developed.

Many autonomous and robotic systems use stereo vision to extract information about the relative position of 3D objects in their vicinity. Robots can use stereo vision to recognize and distinguish objects, using depth information to distinguish between similar objects placed one in front of another. Stereo vision can be used to extract information from aerial surveys, for calculation of contour maps, and in geometry extraction for 3D building mapping.



Left and right stereo images, and a disparity map containing depth information (center). Next page: graph cut and simulated annealing methods. (Stereo images: Middlebury Stereo Data Set; disparity maps: Pat Teller, UTEP)

Generally, stereo algorithms perform any or all of the following four steps:

1. matching cost computation;
2. cost (support) aggregation;
3. disparity computation / optimization; and
4. disparity refinement.

Depending on how they perform step 3, stereo matching algorithms can be classified as local or global.

The Army Research Lab provided a reference code to the UTEP group that was based on a global algorithm using simulated annealing. The performance of this reference code was taken as a reference point with which to compare alternate global algorithms that provide relative performance improvements. The main distinction between these algorithms is the minimization procedure that is used. The reference code uses single-precision simulated annealing as the minimization procedure, while other algorithms use probabilistic (mean-field) diffusion algorithms and graph cut algorithms. Because it is known to have better relative execution time performance and image accuracy, a graph cut algorithm was compared to the simulated annealing algorithm used in the reference code. Given that the graph cut code is known to have a better relative execution time, it is likely to also have lower energy consumption than its simulated annealing counterpart.

Results of comparative performance studies using optical stereo image pairs demonstrate that the graph cut code runs faster, produces a lower minimized energy function, and executes with marginally lower power consumption than the simulated annealing code. The average percentage improvements for 256×256 , 512×512 , and 1024×1024 pixel images were 182%, 140%, and 85%, respectively, for execution time; 20%,

29%, and 27% for the minimization function; and 1.05%, 0.69%, and 1.24% for power consumption. (See references 3 and 4, below.)

If HHPC is to be adopted widely for on-board image processing, a systematic cost-benefit evaluation is essential. Understanding the balance between image quality, power consumption, and execution time provides necessary guidance in the development of hardware and software capabilities. This knowledge also can drive rational decision-making in the acquisition and configuration of on-board systems. ★


For More Information:

[1] Portillo, R., S. Arunagiri, and P. Teller. Power vs. Performance Evaluation of Synthetic Aperture Radar Image-Formation Algorithms and Implementations for Embedded HEC Environments (Ongoing Study). Technical Report, UTEP-CS-10-47, Department of Computer Science, The University of Texas at El Paso, El Paso, TX, October 2010. <http://www.cs.utep.edu/vladik/2010/tr10-48.pdf>

[2] Portillo, R., S. Arunagiri, P. Teller, L. H. Nguyen, S. J. Park, D. R. Shires, and J. C. Deroba. "Power versus Performance Tradeoffs of GPU-accelerated Backprojection-based Synthetic Aperture Radar Image Processing." To appear in the *Proceedings of the Modeling and Simulation for Defense Systems and Applications VI Conference*, part of the SPIE Defense, Security, and Sensing Conference, Orlando, FL, April 25–29, 2011.

[3] Arunagiri, S., V. Barraza, and P. Teller, "Stereo Matching: Performance Study of Simulated Annealing vs. Graph Cut Algorithms Using Optical Images." Technical Report, Department of Computer Science, The University of Texas at El Paso, El Paso, TX, in preparation.

[4] Arunagiri, S., V. Barraza, P. Teller, J. C. Deroba, D. R. Shires, L. H. Nguyen, and S. J. Park, "Stereo Matching: Performance Study of Two Global Area-Based Algorithms." To appear in the *Proceedings of the Radar Sensor Technology XV Conference*, part of the SPIE Defense, Security, and Sensing Conference, Orlando, FL, April 25–29, 2011.



AHPCRC Featured Principal Investigators

Numerical Simulation of Flapping Flows



Terry Holst
NASA Ames Research Center,
Moffett Field, CA
Terry.L.Holst@nasa.gov



Piyush Mehrotra
NASA Ames Research Center,
Moffett Field, CA
Piyush.Mehrotra@nas.nasa.gov
(650) 604-5126



Tom Pulliam
NASA Ames Research Center,
Moffett Field, CA
Thomas.H.Pulliam@nasa.gov

Advanced Optimization Algorithms and Software

Walter Murray
Management Science and
Engineering (Research),
Director of Student Affairs for
the Institute for Computational
and Mathematical Engineering,
Director of the Systems
Optimization Laboratory,
Stanford University
walter@stanford.edu
(650) 723-1307



Michael Saunders
Management Science and
Engineering (Research),
Systems Optimization
Laboratory,
Stanford University
saunders@stanford.edu
(650) 723-1875



Multidisciplinary Parametric Modeling and L/D Quantification and Optimization



Antony Jameson
Aeronautics and Astronautics
Department
Stanford University
jameson@baboon.stanford.edu
(650) 725-6208

Multiscale Modeling of Materials



Eric Darve
Mechanical Engineering Dept.,
Institute for Computational
and Mathematical Engineering,
Stanford University
darve@stanford.edu.
(650) 725-2560

Evaluating Heterogeneous High Performance Computing for Use in Field- Deployable Systems



Jeanine Cook
Klipsch School of Electrical and
Computer Engineering
New Mexico State University
jecook@nmsu.edu
(575) 646-3153

Specifying Computer Systems for Field- Deployable and On-Board Systems



PATRICIA J. TELLER
Professor, Computer Science
The University of Texas at El Paso
(915) 747-5939
pteller@utep.edu



SARALA ARUNAGIRI
Research Specialist, Computer
Science,
The University of Texas at El Paso
(915) 747-8819
sarunagiri@utep.edu

Ballistic Impact Simulations

Adequately protecting soldiers and equipment from ballistic impact damage, without adding excessive weight, requires an understanding of the physics involved in a ballistic impact event.

Environmental Degradation

Professors Charbel Farhat (Stanford University) and Tarek Zohdi (The University of California, Berkeley), and their co-workers are modeling fiber-based composite materials and the effects of moisture absorption, heat, and mechanical damage on ballistic fabrics in laminated and metal-substrate systems. The HPC technology they are developing addresses the micro-mechanisms that control the fabric's service life under aggressive environmental conditions, and the impact of degradation on ballistic resistance.

Laboratory tests are being run at Berkeley to validate the ballistic fabric simulations. Zohdi's group currently has pneumatic and powder guns available for firing simulated fragments. A high-speed video camera captures the impact at up to 10,000 frames per second at a resolution of 64 by 16 pixels.

Farhat and Zohdi have transferred much of their information and findings to their counterparts at the Army Research Laboratory (ARL), including the effects of degradation of yarn by moisture. These parameters have been used in material response models for a woven composite in a code used at ARL for impact events. The research groups are currently interacting to identify mechanisms for loss in ballistic performance of composites after exposure to heat and moisture. ARL will transfer experience in using Digital Image Correlation as a diagnostic technique in upcoming impact tests at Berkeley.

Electromagnetic Fabrics

Ballistic fabrics can fail to impede a projectile with a sharp point or jagged edges, even if the fibers in the fabric do not break. A sharp projectile tip or jagged piece of shrapnel can push the fibers aside and penetrate between them if the force of impact is strong enough. Farhat's and Zohdi's research groups are investigating a potential solution to this problem: using



an electromagnetic (EM) field to reduce the force of the impact.

Electromagnetically-sensitive fabrics represent a potentially significant enhancement over traditional ballistic shielding fabrics. EM forces on the fabric can cause a projectile to rotate, making it less likely to penetrate the fabric. Induced projectile tumbling has been attempted for many years by non-EM means. However, no prior research has explored the deformation of EM-sensitive fabric shielding.

High-performance computing algorithms are under development, capable of treating the type of unique physics involving multiphysical contact, transient current flow through a fabric network, EM fabric deformation (and rupture) and electromagnetically-induced thermodynamic (Joule) heating. Laboratory experiments conducted at Berkeley will calibrate the HPC models, and the massively parallel computational models will be developed and executed at the computational facilities at Stanford. ARL researchers are currently collaborating with Zohdi on implementing a failure response model for composites, and with Zohdi and Farhat to manufacture EM-sensitive panels for ballistic testing at ARL.

Impact on Soft Materials

Very little is known about the specific mechanisms by which soft materials (including human tissue) fail after ballistic impact. Constructing computer simulations of materials undergoing shock or ballistic



Pneumatic gun, breech, and barrel setup for laboratory validations of ballistic impact simulations. (Tarek Zohdi, The University of California, Berkeley)

impact is especially difficult because complex changes occur rapidly. Soft materials absorb and dissipate impact energy, deform, melt, and crack, in a few milliseconds. The number of computations required to construct a realistic simulation can tax the resources of even the best high performance computers. Thus, it is necessary to frame the problem and design the computational codes to use computing resources efficiently, without sacrificing accurate results.

Adrian Lew (Stanford University), Mark Potts (HPTi), and co-workers have delivered the first version of COMODIN++ (Spanish for “wild card”), a parallel code for nonlinear solid dynamics, to the ARL, and subsequent versions are under development. This code is capable of fully asynchronous time stepping, allowing calculations in rapidly deforming regions of the solid to be performed in fine detail while using a coarser resolution in more stable areas. The code has been scaled to more than 1000 processors and modeled nearly 700 million elements.

COMODIN++ is being developed to model the response of ballistic gelatin (a human tissue simulant) to stress, impact, heat, and shock, and to simulate the evolution of the cavity behind a projectile. The

Asynchronous time stepping adjusts the mesh size to fit the calculation. (Adrian Lew, Stanford University)

simulations are being extended to cover a period of 10–50 milliseconds after impact (a significant period of time in the world of computer simulations for this problem) and to calculate the amount of energy involved in damaging the material.


Current research focuses on increasing the number of processors for the modeling and simulation runs, investigating the effect of techniques for simplifying the calculations on the accuracy of the results, and making the models more realistic by adding features such as material failure and energy dissipation effects. Results of the modeling studies are compared with laboratory tests on Permagel performed in Zohdi's laboratory at Berkeley.

Recent efforts address the appearance of a cavity behind a bullet as a function of the material properties. A viscoelastic material model has been implemented with several time constants provided by ARL.

One of Lew's students, Raymond Ryckman, demonstrated to ARL researchers how the asynchronous time step algorithm was implemented. The researchers have expressed an interest in implementing this capability into the DoE codes used at ARL for large-scale simulations.

Biological Warfare Agent Release

Gianluca Iaccarino, Eric Shaqfeh and Mark Jacobson (Stanford University) use HPC to create realistic simulations of biological warfare agent (BWA) release scenarios. Recent efforts have focused on Oklahoma City, using conditions that replicate the July 2003 Joint Urban Atmospheric Dispersion Study. They are developing a computational framework for modeling the dispersion of aerosolized BWA particles, from a few nanometers to more than 100 micrometers across, in a turbulent air flow. They are coupling two modeling methods, one that operates over large areas (10–100 square km), and one that models building-scale details (30–50 cubic m). In order to increase the fidelity of the overall simulation, each modeling method passes data to the other. Simulations run over



Tracer concentration map for Oklahoma City simulation.
(Mark Z. Jacobson, Stanford University)

time scales that are long enough to provide useful results (~20 min). The models incorporate topological and photochemical effects, among many other factors.

Simulated vertical transport effects have been compared with black carbon profile measurements from a 2009 study conducted over the Pacific Ocean. A simulation using 68 layers, from 0 to 60 km in altitude, produced results that differed from the measured column loading by 1.4%, and produced a mean vertical profile similar to that of the observed data. This is a significant improvement over published results from 14 other global models that over-predicted black carbon concentrations by a factor of five and produced vertical profile slopes that were effectively vertical in the troposphere (an indicator of numerical diffusion).

The team has an ongoing interaction with researchers at the U.S. Army Edgewood Chemical and Biological Center and the Army Research Laboratory at Adelphi, Maryland. One ARL researcher stated that current modeling capabilities do a very poor job of modeling hazard propagation in an urban environment, and that the type of work being done by AHP-CRC is of crucial importance for analysis and operational type planning.


Blood Cells in Microfluidic Flow

Eric Shaqfeh, Eric Darve, and their students (Stanford University) have developed simulations to aid in understanding the mechanisms of flow for platelets and particles of essentially arbitrary shape circulating in the smallest blood vessels. These simulations include adsorption, an important factor for the first step in clot formation and trauma response in the microcirculation.

Shaqfeh and Darve's models take into account factors such as electrical forces, flexible or rigid particles of various shapes, Brownian (random) motion, and sedimentation effects. At present, no other computational simulation techniques exist that include all of these factors in the same package. Their simulation codes handle orientable objects in a flow with hydrodynamic interactions as well as complex microfluidic environments and particle shapes, deformable particle surfaces, and complex interactions between the solid and liquid phases. The model now accounts for adhesion between particles and the blood vessel wall (plaque or clot formation) and will in the future include adhesion among the particles themselves.

The computational models are being compared with laboratory-made particles of various shapes and sizes, flowing through straight or bifurcated channels, generated by Samir Mitragotri's group at the Institute for Collaborative Biotechnologies (ICB) at the University of California at Santa Barbara. Sumita Pennathur, also of ICB, is measuring the adsorption rates of these particles, and the simulations are also compared with this data.

Researchers at Walter Reed Army Institute of Research (WRAIR) are interested not only in adhesion properties, but also in examining the concentration distribu-



Cross-sectional profile of blood vessel. Red line: red blood cell concentration, 25% hematocrit. Green bars: platelet concentration. (Eric Shaqfeh, Stanford University)

tion of platelets in the microvesicles, particularly how the use of freeze-dried platelets affects concentration distributions. Shaqfeh and Darve have successfully replicated experimental observations of red blood cells concentrating toward the center of the blood vessel, forcing the platelets closer to the vessel walls, an effect that is more pronounced at higher concentrations of red blood cells (higher hematocrit).

Materials Research

Microstructural Defect Modeling

Dislocations (defects) can cause metal parts to crack and break, and they cause degradation in semiconductor infrared (IR), radio-frequency (RF) and micro-electro-mechanical system (MEMS) devices that are essential for the modernization of the Army. At the micro-scale at which these devices work, materials may behave very differently than they do on a bulk scale.

Wei Cai is working with research associate Sylvie Aubry and their graduate students (Stanford University) to develop metal and semiconductor microstructure modeling capabilities using dislocation dynamics (DD), which tracks defect motion through a crystal lattice. They are comparing their results with those obtained using molecular dynamics (MD) modeling, which characterizes the behavior of atoms and small ensembles of atoms, with the intent of bridging these two methods and the length scales to which they apply.


Cai and his collaborators at Lawrence Livermore National Laboratory have developed a publicly available, massively parallel dislocation dynamics simulator called ParaDiS, with the intention of overcoming many of the problems associated with conventional DD programs. By using thousands of processors simultaneously, ParaDiS has, for the first time, successfully captured the strain hardening behavior of a $10\mu\text{m}^3$ representative volume in a bulk metal. The program runs routinely on 100–1000 processors, and it has been demonstrated on the 132,000 processors of the BlueGene/L supercomputer.

Under the AHPCRC program, Cai is developing numerical algorithms and computer programs (imple-


mented in ParaDiS) that allows DD simulation of thin films and micro-cylinders. An efficient image stress (surface effects) algorithm has been implemented in ParaDiS for thin film and cylinder geometries. The thin film algorithm is also working in parallel.

Impact studies on metal parts, including vehicle panels and projectile tips, have revealed a microscale phenomenon called adiabatic shear banding (ASB), providing important insights into why metallic parts bend and break during impact. Shear bands form in metal when the stress and deformation are localized into a small area, and they act as sites for future failures. Microcompression experiments are being conducted at ARL/WMRD to find the origin of this size effect, which is still under debate. Cai's group is coordinating DD simulations for micro-pillar deformations for comparison with results from the ARL group.

Cai's group has performed DD simulations in metal thin film to compare with the experimental measurements performed at ARL/SEDD, to assist in developing a multilayered MEMS. Thin film simulations are also used for examining the dislocation evolution in semiconductor thin films in IR and RF devices (*one example is shown below*). Cai's group provides training and technical support for ARL researchers who wish to use the ParaDiS program.



Dislocation structure in gallium nitride (semiconductor) thin film. (Wei Cai, Stanford University)



Graphene sheet (framework) with chemisorbed hydrogen atoms (green). (Evan Reed, Stanford University)

Graphene-Based Electrical Devices

The discovery of a practical manufacturing process for single-layer graphene opens the potential for fabrication of 2D devices, and may lead to novel electronics applications. Before this potential can be translated into practical application, however, unwanted electronic effects, introduced by myriad chemical impurities and disorder, must be mitigated.

The electrical, mechanical, optical, and other properties of graphene electronic and sensing devices are reconfigurable—this opens up possibilities for even more applications. Chemisorbing one hydrogen atom onto each carbon atom (i.e., chemically bonding hydrogen atoms to the carbon surface) produces graphane, a semiconductor-like material with an electronic bandgap. Electronic structure calculations show that graphane boundaries can be fabricated to produce 2D patterns on a graphene sheet. Such carefully controlled configurations of hydrogen are likely to be very challenging to make in real devices. Some modeling work has been reported on mechanical properties of graphene partially covered with hydrogen (*illustrated above*) using empirical potentials. However, little attention has been focused on the thermodynamically favorable arrangements of hydrogen that would be generated by a practical deposition and annealing approach and their impact on electronic properties.

The Army Research Laboratory approached the AHPCRC Consortium about formulating a project that could provide theoretical insight on these potential applications of graphene, which the consortium

was able to implement quickly. Evan Reed's group (Stanford University) is performing quantum atomistic simulation studies of hydrogen adsorption on graphene for electronics and other applications. They are determining thermodynamically stable states of hydrogen adsorbed on graphene edges and defects and exploring the chemical changes associated with high temperature processing and annealing. They are studying hydrogen deposition on graphene as a method for control of the electronic and other properties. Using quantum approaches and molecular dynamics simulations, they are determining the thermodynamically favorable arrangements of hydrogen around a graphene edge and graphene imperfections (e.g., vacancies). They are characterizing the impact of hydrogen adsorption on the undesirable electronic effects that accompany disorder and defects, as a function of the fraction of hydrogen coverage on the graphene sheet.

In conjunction with these studies, Reed's group is applying their unique expertise in nanoscale piezoelectric and mechanical phenomena to study the potential for modification of graphene's electronic properties through mechanical effects. The controlled adsorption of hydrogen, lithium, and other dopants may lead to exciting new classes of 2D devices that seamlessly integrate electronic and mechanical effects. To address these questions, they are combining quantum-based computational tools, including density-functional theory (DFT) and the self-consistent charge density-functional tight-binding method (SCC-DFTB).

Graphene-based fabrication methods capitalize on carbon, a benign, plentiful resource, and could produce highly miniaturized, versatile electronic devices. Quantum mechanical model development using HPC simulations would not only advance the field of graphene devices, but provide new computational

capabilities for application in other areas. The calculations to date have shown that doping in graphene is electrostatic, which was insight not anticipated by ARL researchers. This opens the door to the development of novel devices, which may out-perform conventional semiconductor-based devices for certain applications.


The All-Electron Battery

Electronic devices can save lives: sensors can detect people through walls, gunshot detectors locate snipers, satellite phones do not require local cellular networks, and enhanced night vision devices remove an adversary's element of surprise. However, each new device adds weight and requires power, for transporting the devices themselves and for generators to recharge the batteries.

Fritz Prinz and co-workers (Stanford University) use high performance computing in their search for materials to construct the all-electron battery (AEB), a new type of device that may deliver both high power density and high energy density. AEBs show potential for efficient energy storage, a lifetime similar to present-day capacitors, no catastrophic failure modes, fast charging and discharging, safe operation. Research on the AEB produced two patent applications in 2010. The group is currently fabricating and testing a proof-of-concept device, evaluating materials for each component of the device, and testing the scalability of the device by increasing the size and adding more layers. The AEB stores energy through charge separation. The charge carriers are moving electrons, which are lighter, and therefore faster than the moving ions typical of most batteries. In the AEB, quantum dot inclusions are embedded in the dielectric structure between two electrodes of a capacitor. Electrons can tunnel through the dielectric between the electrodes and the inclusions, thereby increasing the charge storage density relative to a conventional capacitor.

Developing practical AEB devices requires an understanding of the charge transfer and storage mechanisms involved. Prinz and his group are designing simulations and experiments to test their hypotheses, and to evaluate the effects of quantum dot size and

Quantum dots separated by a dielectric layer: the conceptual basis of the all-electron battery. (Fritz Prinz, Stanford University)




material, as well as dielectric combinations, on charge storage. The group has used quantum calculations to screen viable architectures and materials, in preparation for generating a detailed design of the AEB from the nanoscale up. For validation, the predictions are tested against laboratory measurements.

Army researchers are watching this basic research project with great interest. The concept not only has broad implications for producing new energy storage devices, but it has already provided a new method for producing quantum dots, which are used in other electronic device applications.

Stream Programming

Parallel programming is an intrinsic part of high performance computing (HPC). Whether a programmer is adapting existing software or building new capabilities, codes must be designed to run accurately, reliably, and efficiently on systems that may contain tens to thousands of processors working cooperatively. Parallel programming is not merely a problem of dividing computational tasks among processors. In fact, the most difficult part of parallelism is often moving data to where it is needed, when it is needed. High performance computers and clusters have not reached the state of standardization that allows a programmer to write code that runs equally well on most machines. This is especially true today as the HPC world undergoes a revolution in architecture with the development of heterogeneous platforms (those containing more than one class of processors) and multi-core platforms. To write a parallel program that achieves the best performance on any specific system, a programmer must understand the characteristics of that system, for example the memory architecture, and design the code accordingly. Code that works especially well on one architecture may not achieve the same level of performance on a system with a different size or structure. Conversely, programs written to be highly portable may not perform optimally on any system.

Professors Alex Aiken, William Dally, and Patrick Hanrahan (Stanford University) are leading a group that recently delivered the first version of the Sequoia programming language to the ARL. This language pro-



Hierarchical memory systems are typical of high performance computing environments. (Alex Aiken, Stanford University)

vides Army researchers with the ability to port parallel programs to many types of computing systems and architectures without sacrificing performance. Sequoia allows programmers to write code that is functionally correct on any system, then tune the performance to the characteristics of a specific system. Sequoia syntax is an extension of the C++ programming language, but Sequoia introduces language constructs that produce a programming model that is very different from C++. The Sequoia language makes it easier to develop a parallel program that is “aware” of the memory hierarchy configuration in the machine on which it is running. Computations are localized to specific memory locations, and the language mechanisms describe communications among these locations.

A complete Sequoia programming system has been implemented and released to ARL. The system includes a compiler and runtime systems that deliver efficient performance for both Cell processors and distributed memory clusters.

ARL(CISD) researchers are interested in Sequoia because it fills a gap—unlike other approaches, it actually targets the complex memory systems found in today’s evolving microprocessors. This will enable Army computer scientists to develop high performance codes by targeting a variety of hybrid binary computing systems in real operational scenarios. This programming system can also be used on tactical or deployed HPC platforms, delivering real-time intelligence to the warfighter.

★

AHPCRC Publications and Presentations June–December 2010

A complete list of publications and presentations is available at <http://www.ahpcrc.org/publications.html>

Project 1–1: Multifield Simulations of Accelerated Environmental Degradation of Fabric, Composite, and Metallic Shields and Structures

- Dynamics of clusters of charged particulates in electromagnetic fields. Zohdi, T. I. *The International Journal of Numerical Methods in Engineering*. Published online 24 August 2010, doi:10.1002/nme.3007
- Electromagnetically-induced deformation of functionalized fabric. Zohdi, T. I. *The Journal of Elasticity* (in press).
- Simulation of coupled microscale multiphysical fields in particulate-doped dielectrics with staggered adaptive FDTD. Zohdi, T. I. *Computer Methods in Applied Mechanics and Engineering*, 199(49–52), 3250–3269, 2010. doi:10.1016/j.cma.2010.06.032.
- Joule-heating field phase-amplification in particulate-doped dielectrics. Zohdi, T. I. *The International Journal of Engineering Science*, 49(1), 30–40, 2011. doi:10.1016/j.ijengsci.2010.06.021.
- Multi-Scale Modeling and Large-Scale Transient Simulation of Ballistic Fabric and Fabric-Resin Composites. Powell, D., Farhat, C., Zohdi, T. 9th World Congress on Computational Mechanics, Sydney, Australia, July 2010.
- Modeling and simulation of multiphysical processes in particulate media. Zohdi, T. University of Colorado, Boulder, Department of Mechanical Engineering. Invited lecture (colloquium), May 2010.

Project 1–2: Simulation of Ballistic Gel Penetration

- Parameterization of planar curves immersed in triangulations with application to finite elements. Rangarajan R., Lew A. submitted.
- Optimal convergence of a discontinuous-Galerkin-based immersed boundary method. Lew A., Negri M. submitted.
- Explicit asynchronous contact algorithm for elastic rigid body interaction. Ryckman R., Lew A. submitted.
- Stability and convergence proofs for a discontinuous-Galerkin-based extended finite element method for fracture mechanics. Shen Y., Lew A. *Computer Methods in Applied Mechanics and Engineering*, in press.
- An optimally convergent discontinuous-Galerkin-based extended finite element method for fracture mechanics. Shen Y., Lew A. *International Journal for Numerical Methods in Engineering*, 82:6, 716–755, 2010.
- An adaptive stabilization strategy for enhanced strain methods in nonlinear elasticity. TenEyck A., Lew A. *International Journal for Numerical Methods in Engineering*, 81:11, 1387–1416, 2010.

Project 1–3: Multidisciplinary Parametric Modeling and Lift/Drag Quantification and Optimization

- Hybrid optimization schemes for wing modeling of micro-aerial vehicles. Velazquez, L., Argaez, M., Culbreth*, M., Sanchez, R.*, Ramirez, C.*, Hernandez IV, M.* User Group Conference Proceedings, *IEEE-CS Journal*, Schaumburg, IL, June 2010.

Project 1–4: Flapping and Twisting Aeroelastic Wings for Propulsion

- Effects of mass ratio to flexible flapping-wing propulsion. M. Xu, M. Wei, T. Yang, Y. Lee, and T. D. Burton. *Bulletin of the American Physical Society*, Vol. 55, No. 16, Long Beach, CA, 2010.
- A fully-coupled approach to simulate three-dimensional flexible flapping wings. T. Yang, and M. Wei. *Bulletin of the American Physical Society*, Vol. 55, No. 16, Long Beach, CA, 2010.
- Effect of gust on flow patterns around a robotic hummingbird wing. E. N. Marquez, H. Evans, R. Alarcon, G. Whitehouse and B.J. Balakumar. *Bulletin of the American Physical Society*, Vol. 55, No. 16, Long Beach, CA, 2010.
- Lift, drag and flow-field measurements around a single-degree-of-freedom toy ornithopter. R. Alarcon, B.J. Balakumar, and J. Allen. *Bulletin of the American Physical Society*, Vol. 55, No. 16, Long Beach, CA, 2010.
- A global approach for reduced-order models of flapping flexible wings. M. Wei, T. Yang. AIAA paper 2010–5085, Chicago, IL, 2010.
- Numerical Study of Flexible Flapping Wing Propulsion. T. Yang, M. Wei. *AIAA Journal*, Vol. 48, No. 12, pp. 2909–2915, 2010.
- Robust and Provably Second-Order Explicit-Explicit and Implicit-Explicit Staggered Time-Integrators for Highly Nonlinear Fluid-Structure Interaction Problems. C. Farhat, A. Rallu, K. Wang, T. Belytschko. *International Journal for Numerical Methods in Engineering* (in press).
- Total Energy Conservation in ALE Schemes for Compressible Flows. Dervieux, C. Farhat, B. Koobus, M. Vazquez. *European Journal of Computational Mechanics* (in press).
- Nonlinear Structural response in flexible flapping wings with different density ratio. Xu, M., Wei, M., Yang, T., and Burton, T. D. Submitted to AIAA ASM, Orlando, FL, 2011.
- Numerical Study of Flexible Flapping Wing Propulsion. Yang, M. Wei, H. Zhao. AIAA paper 2010-0553, Orlando, FL, 2010.
- Computational Analysis of Hovering Hummingbird Flight. Z. Liang, H. Dong, M. Wei. AIAA paper 2010-0555, Orlando, FL, 2010.
- Optimal Flight of Rufous Hummingbirds in Hover: An Experimental Investigation. H. Bocanegra Evans, J. J. Allen, and B. J. Balakumar. AIAA paper 2010-1028, Orlando, FL, 2010.

Project 1–5: Numerical Simulation of Flapping Flows

- Grid and Time Step Requirements to Accurately & Efficiently Resolve Flow Around a Rigid Flapping Airfoil using OVERFLOW. Lefell, J. and Pulliam, T. Presented at the 49th AIAA Aerospace Sciences Meeting in Orlando, FL, January 4–7, 2011.

Project 1–6: The All-Electron Battery: Quantum Mechanics of Energy Storage in Electron Cavities

- Quantum dot ultracapacitor and electron battery. Holme, Timothy P., Prinz; Friedrich B. United States Patent Application 20100183919. July 22, 2010.
- All -electron battery having area-enhanced electrodes. Holme; Timothy P., Prinz; Friedrich B., Usui, Takane. United States Patent Application 20100255381. October 7, 2010.

Project 1–7: Advanced Optimization Algorithms and Software

- A Regularized Active-set Method for Sparse Convex Quadratic Programming. C. M. Maes. PhD thesis, Stanford University, November 2010.
- 40 Years of Linear Algebra and Optimization at Stanford. M. A. Saunders, Mathematics and Systems Biology seminar, University of Iceland, Reykjavik, Iceland, October 4, 2010.
- QPBLUR: A regularized active-set method for sparse convex quadratic programming. M. A. Saunders, keynote speaker (with C. M. Maes). OPTEC Workshop on Large-Scale Convex Quadratic Programming—Algorithms, Software, and Applications, Katholieke Universiteit Leuven, Belgium, Nov 25–26, 2010.
- An algorithm for nonlinear optimization problems with binary variables. Walter Murray and Kien-Ming Ng. *J. Computational Optimization and Applications*, 47:2, 257–288 (2010).
- LSMR: An iterative algorithm for sparse least-squares problems. Michael Saunders, plenary speaker. Second International Conference on Numerical Linear Algebra and Optimisation, University of Birmingham, UK, Sep 13–15, 2010.
also: LSMR: An iterative algorithm for sparse least-squares problems. M. A. Saunders (with D. Fong). Mathematics seminar, Delft Institute of Applied Mathematics, Delft, The Netherlands, Nov 29, 2010 and Mathematics seminar, Applied Analysis and Computational Science (AACS), University of Twente, Enschede, The Netherlands, Dec 2, 2010.
- MINRES-QLP: a Krylov subspace method for indefinite or singular symmetric systems. S.-C. T. Choi, C. C. Paige, and M. A. Saunders. *SIAM J. Sci. Comp.* (submitted March 2010), 26 pp.
- Presentation on sparse least-squares problems. M. Saunders, D. Fong. Eleventh Copper Mountain Conference on Iterative Methods, April 2010.

Project 2–1: Dispersion of Biowarfare Agents in Attack Zones

- Development and application to Oklahoma City of a new mass, energy, vorticity, and potential enstrophy conserving scheme for 3-D nonhydrostatic atmospheric flows with complex boundaries. Ketefian, G.S., and M.Z. Jacobson. American Geophysical Union Fall Meeting, San Francisco, California, Dec. 13–17, 2010.
- Jacobson, M.Z. Numerical Solution to Drop Coalescence/Breakup With a Volume-Conserving, Positive-Definite, and Unconditionally-Stable Scheme. *J. Atmos. Sci.*, in press, doi :10.1175/2010JAS3605.1, 2011.
- Ketefian, G.S., and M.Z. Jacobson. A piecewise-linear boundary scheme for the shallow water equations that conserves mass, energy, vorticity, and potential enstrophy. *J. Comp. Phys.*, in press, 2011.
- The global-through-urban nested 3-D simulation of air pollution with a 13,600-reaction photochemical mechanism. Jacobson, M.Z., Ginnebaugh, D. L. *J. Geophys. Res.*, 115, D14304, 13 pp., doi:10.1029/2009JD013289, 2010. (www.stanford.edu/group/efmh/jacobson/3Dgas-photochem.html).
- Modeling Normal Reynolds Stress Anisotropy for use with Algebraic Scalar Flux Closures. Philips, S., Iaccarino, G., in preparation.
- A numerical study of scalar dispersion downstream of a wall-mounted cube using direct simulations and algebraic flux models. Rossi, R., Philips, D., Iaccarino G. I. *J. Heat Fluid Flow*, accepted 2010.

Project 2–2: Micro- and Nano-fluidic Simulations for Biowarfare Agent Sensing and Blood Additive Development

- On the physical mechanism of platelet margination in the microvasculature. Zhao, H. and Shaqfeh E.S.G. *Physical Review Letters* (submitted August 2010).
- The dynamics of a vesicle in simple shear flow. Zhao, H. and Shaqfeh E.S.G. *Journal of Fluid Mechanics* (revision submitted October 2010).
- On the Irreversible Adsorption and Taylor Dispersion of Particles in Channel Flows of General Cross Section. Fitzgibbon, S. and Shaqfeh, E.S.G. *Phys. Fluids* (submitted March 2010)

continued on page 28

continued from page 27

Project 2-4: Protein Structure Prediction for Virus Particles

- Constraint Logic Programming in Evolutionary Biology. B. Chisham, E. Pontelli, T. Son, B. Wright. *Theory and Practice of Logic Programming* (Submitted).
- Constraint based fragment assembly. A. Dal Palu, A. Dovier, E. Pontelli. International Joint Conference on Artificial Intelligence (IJCAI), (Submitted).
- Constraint-based Protein Fragment Assembly. A. Dal Palu, A. Dovier, F. Fogolari, E. Pontelli. *Theory and Practice of Logic Programming*, 10(4-6):709-724, 2010.
- Computing approximate solutions of the protein structure determination problem using global constraints on discrete crystal lattices. A. Dal Palu, A. Dovier, E. Pontelli. *International Journal on Data Mining and Bioinformatics*, 4(1):1-20, 2010.
- An investigation in parallel execution of answer set programs on distributed memory platforms: task sharing and dynamic scheduling. E. Pontelli, H. Le, T. Son. *Computer Languages, Systems & Structures*, 36(2):158-202, 2010.
- CDAO-STORE: A New Vision for Data Integration. B. Chisham, T. Le, E. Pontelli, T. Son, B. Wright. *Nature Precedings*, doi:10.1038/npre.2010.4586.1, 2010.
- A New Vision for Data Integration in Computational Biology. B. Chisham. Presentation, iEvoIO Workshop, Portland, July 2010.
- Protein Fragments Assembly in CLP. E. Pontelli, A. Dovier, A. Dal Palu, F. Fogolari. International Conference on Logic Programming, Edinburgh, Scotland, July 2010. [BEST PAPER AWARD]
- CLP-based Protein Fragment Assembly. A. Dovier, A. Dal Palu, F. Fogolari, E. Pontelli. *Theory and Practice of Logic Programming*, 10(4-6):709-724, 2010.
- An empirical study of constraint logic programming and answer set programming solutions of combinatorial problems. A. Dovier, A. Formisano, E. Pontelli. *Journal of Experimental and Theoretical Artificial Intelligence*, 21(2):79-121, 2009.
- Structure prediction for the helical skeletons detected from the low resolution protein density map. Al Nasr, K., Sun, W., He, Jing. *BMC Bioinformatics*, vol. 11, 2010.
- Enumeration of all geometrically constrained assignments of the secondary structures using a graph for the protein structure prediction. Al Nasr, K., Ranjan, D., He, J. *International Conference in Bioinformatics and Systems Biology*, vol. 3, 2010.
- Computing Approximate Solutions of the Protein Structure Determination Problem using Global Constraints on Discrete Crystal Lattices. A. Dovier, A. Dal Palu, E. Pontelli. *International Journal of Data Mining and Bioinformatics*, 4(1):1-20, 2010.

Project 2-5: Nanoscale Dislocation Dynamics in Crystals

- Dislocation Junctions and Jogs in Free Standing Thin Films. Seokwoo Lee, Sylvie Aubry, William D. Nix and Wei Cai. *Modelling and Simulation in Materials Science and Engineering*, 19, 025002 (2011).
- The stability of Lomer-Cottrell Jogs in Nano-Pillars. Christopher R. Weinberger and Wei Cai. *Scripta Materialia*, 64, 529 (2011).
- Equilibrium Shape of Dislocation Shear Loops in Anisotropic alpha-Fe. Sylvie Aubry, Steven P. Fitzgerald, Sergei L. Dudarev and Wei Cai. Submitted to *Modelling and Simulation in Materials Science and Engineering*.
- The Stability of Lomer-Cottrell Jogs in Nanopillars. Christopher R. Weinberger, Wei Cai. Submitted to *Scripta Materialia*, 2010.
- Plasticity of metal wires in torsion: molecular dynamics and dislocation dynamics simulations. Christopher R. Weinberger, Wei Cai. *Journal of Mechanics and Physics of Solids*, 58, 1011 (2010).
- Orientation dependent plasticity in metal nanowires under torsion: twist boundary formation and Eshelby twist. Christopher R. Weinberger, Wei Cai. *Nano Letters*, 10, 130142 (2010).

Project 3-1: Information Aggregation and Diffusion Under Mobility

- Interactive Analysis and Simulation of VANETs Using MOWINE. Ian Downes, Branislav Kusy, Omprakash Gnawali, and Leonidas Guibas. *Proceedings of the IEEE Vehicular Networking Conference (VNC 2010)*, December 2010.
- Collaborative Image Annotation Using Image Webs. Zixuan Wang, Omprakash Gnawali, Kyle Heath, and Leonidas Guibas. Presentation AO-03, *Proceedings of the 27th Army Science Conference (ASC 2010)*, November 2010.
- END: A Topology-Aware Metric for Sensor Networks. Daniele Puccinelli, Omprakash Gnawali, SunHee Yoon, Silvia Giodano, Leonidas Guibas. Poster in ACM Conference on Embedded Networked Sensor Systems (SenSys 2010), November 2010.
- Interactive Analysis and Simulation of VANETs Using MOWINE. Ian Downes. Presentation at IEEE Vehicular Networking Conference, December 2010.
- A Case for Evaluating Sensor Network Protocols Concurrently. Omprakash Gnawali (presenter), Leonidas Guibas, Philip Levis. *Proceedings of the Fifth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, September 2010.
- Data Stashing: Energy-Efficient Information Delivery to Mobile Sinks through Trajectory Prediction. HyungJune Lee, Martin Wicke, Branislav Kusy, Omprakash Gnawali, Leonidas Guibas. *Proc. of ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, April 2010.

- Fingerprinting Mobile Users in Wireless Sensor Networks with Network Flux Information. Mo Li, Xiaoye Jiang, Branislav Kusy, Leonidas Guibas. 30th International Conference on Distributed Computing Systems (ICDSC), June 2010.
- Image Webs: Computing and Exploiting Connectivity in Image Collections. K. Heath, N. Gelfand, M. Ovsjanikov, M. Aanjaneya, L. J. Guibas. 23rd IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2010.
- Connected Dominating Sets on Dynamic Geometric Graphs. Leonidas Guibas, Nikola Milosavljevic, Arik Motzkin. 22nd Canadian Conference on Computational Geometry (CCCG), July 2010.
- END: A Topology-Aware Metric for Sensor Networks. Daniele Puccinelli, Omprakash Gnawali, SunHee Yoon, Silvia Giordano, Leonidas Guibas. *Proc. of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, November 2010 (poster).

Project 3–2: Scalable Design Methods for Topology Aware Networks

- Subgraph Sparsification and Nearly Optimal Ultrasparsifiers. Kolla, Y. Makarychev, A. Saberi, S. Teng. *Proceedings, 42nd ACM Symposium on Theory of Computing (STOC 2010)*. Available at: <http://www.stanford.edu/~saberi/sparsifier.pdf>

Project 3–3: Secure Sensor Data Dissemination and Aggregation

- Jamming Dust: A Low-Power Distributed Jammer Network. H. Huang, N. Ahmed, S. Pulluru. Poster BP-03, 27th Army Science Conference, Orlando FL, November 29–December 2, 2010.
- On the Low-Power Distributed Jammer Network. H. Huang, N. Ahmed, and S. Pulluru, submitted for publication.
- Achieving Optimal Tradeoff between Efficiency and Security in Sensor Data Aggregation. H. Huang, V. Kodali, and Y. Katuru, submitted for publication.

Project 3–4: Robust Wireless Communications in Complex Environments

- Analytical Function for Multiple Distance Measures in a Mixed Wireless Network. Abdulaye Traore. Masters Thesis, December 2010.
- Modeling and Managing QoS in Mixed Wireless Networks using the Power Performance Measure. Astatke Y., Dean R. Globecom 2010.
- Mixed Network Clustering with Multiple Ground Stations and Node Preference. Traore, O., Gwanvoma, S., Dean, R. ITC 2010.
- Mixed Networks Interference Management with Multi-Distortion Measures. Traore, A., Dean, R. ITC 2010.
- QoS Performance Management in Mixed Wireless Networks. Astatke A., Dean, R. ITC2010.
- Three student Summer reports are available for the work accomplished for Summer 2010.

Project 3–5: Mobile Brain–Machine Interface for Integrated Information–Social/Cognitive Network Operations

- Novel beamformers for multiple correlated brain source localization and reconstruction. Hung V. Dang, Kwong T. Ng, and James K. Kroger (in review). 2011 International Conference on Acoustics, Speech and Signal Processing.
- Novel vector beamformers for EEG source imaging. Hung V. Dang, Kwong T. Ng, and James K. Kroger (in review). 2011 IEEE International Symposium on Biomedical Imaging.
- FDEHMT: A Finite Difference Electromagnetic Head Modeling Toolbox. H.V. Dang and K. T. Ng. Biomedical Engineering Society 2010 Annual Meeting, Austin, TX, October 2010.

Project 4–1: Stream Programming for High Performance Computing

- Programming the Memory Hierarchy Revisited: Supporting Irregular Parallelism in Sequoia. M. Bauer, J. Clark, E. Schkufza, A. Aiken. Accepted to the Symposium on Principles and Practice of Parallel Programming 2011.

Project 4–2: Massive Scale Data Analysis on the Flexible Architecture Research Machine (FARM)

- Accelerating CUDA graph algorithms at maximum warp. S. Hong, S. Kim, T. Oguntebi, K. Olukotun. *Proceedings of the 16th ACM SIGPLAN Annual Symposium on Principles and Practices of Parallel Programming (PPOPP)*, 2011.
- Hardware Acceleration of Transactional Memory on Commodity Systems. J. Casper, T. Oguntebi, S. Hong, N. Bronson, C. Kozyrakis, K. Olukotun. 16th International Conference on Architectural Support for Programming Languages and Operating Systems, 2011.
- FARM: A prototyping environment for tightly-coupled, heterogeneous architectures. T. Oguntebi, S. Hong, J. Casper, N. Bronson, C. Kozyrakis, and K. Olukotun. In *Proceeding of the 18th Annual Symposium on Field-Programmable Custom Computing Machines (FCCM '10)*, 2010.

Project 4–3: Specifying Computer Systems for Field-Deployable and On-Board Systems of Multicore Processors

- FAIRIO: An Algorithm for I/O Performance Differentiation Based on Bottleneck Analysis. S. Arunagiri, Y. Kwok, S. Seelam, P. Teller, and R. Portillo. Submitted to the 2011 IEEE International Parallel & Distributed Processing Symposium (IPDPS 2010), Anchorage, Alaska, May 16–20, 2011.

continued on page 30

continued from page 29

- Power versus Performance Tradeoffs of GPU-accelerated Backprojection-based Synthetic Aperture Radar Image Processing. Portillo, R., S. Arunagiri, P. Teller, L. H. Nguyen, S. J. Park, S. J., D. R. Shires, and J. C. Deroba. To appear in *Proceedings of the Modeling and Simulation for Defense Systems and Applications VI Conference*, part of the SPIE Defense, Security, and Sensing Conference, Orlando, FL, April 25–29, 2011.
- Stereo Matching: Performance Study of Two Global Area-Based Algorithms. Arunagiri, S., V. Barraza, P. Teller, J. C. Deroba, D. R. Shires, L. H. Nguyen, and S. J. Park. To appear in *Proceedings of the Radar Sensor Technology XV Conference*, part of the SPIE Defense, Security, and Sensing Conference, Orlando, FL, April 25–29, 2011.
- Power vs. Performance Evaluation of Synthetic Aperture Radar Image-Formation Algorithms and Implementations for Embedded HEC Environments (Ongoing Study). Portillo, R., S. Arunagiri, P. Teller. Technical Report, UTEP-CS-10-47, Department of Computer Science, The University of Texas at El Paso, El Paso, TX, October 2010. www.cs.utep.edu/vladik/2010/tr10-48.pdf
- Embedded High-end Computing in Mobile Tactical Radar Systems—Power versus Execution-Time Performance Tradeoffs. Portillo, R., S. Arunagiri, and P. Teller. UTEP Research Booth poster presented at SC10, The 23rd International Conference for High Performance Computing, Networking, Storage and Analysis, New Orleans, LA, November 13–19, 2010.
- FAIRIO: An Algorithm for Differentiated I/O Performance. Arunagiri S., Y. Kwok, S. Seelam, P. Teller, and R. Portillo. UTEP Research Booth poster presented at SC10, The 23rd International Conference for High Performance Computing, Networking, Storage and Analysis, New Orleans, LA, November 13–19, 2010.
- GPGPU Programming Approach Productivity Comparisons—CUDA vs. OpenCL vs. PGI. Kwok, Y., J. McCartney, S. Arunagiri, and P. Teller. UTEP Research Booth poster presented at SC10, The 23rd International Conference for High Performance Computing, Networking, Storage and Analysis, New Orleans, LA, November 13–19, 2010.
- GPGPU Programming Approach Productivity Comparisons—CUDA vs. OpenCL vs. PGI. Kwok Y. AHPCRC Research Booth presentation at SC10, The 23rd International Conference for High Performance Computing, Networking, Storage and Analysis, New Orleans, LA, November 13–19, 2010.
- GPUs in Mobile Tactical Systems: Power versus Image Quality Tradeoffs. Portillo, R. AHPCRC Research Booth presentation at SC10, The 23rd International Conference for High Performance Computing, Networking, Storage and Analysis, New Orleans, LA, November 13–19, 2010.
- Synthetic Aperture Radar (SAR) Image Formation (IF) Power vs. Performance Study Phase I – Preliminary Analysis. R. Portillo, S. Arunagiri, P. Teller. Technical report, in preparation.
- OpenCL, CUDA, PGI compiler: Performance Studies Using Simple Kernels. Y. Kwok, J. L. McCartney, S. Arunagiri, P. Teller. Technical report, in preparation.
- Stereo Matching: Comparative Performance Studies of Graph Cut vs. Simulated Annealing. S. Arunagiri, V. J. Barraza, P. Teller. Technical report, in preparation.
- On the Use of Shareable Resource Signatures and Hardware Thread Priorities to improve throughput of (SMT) Processors. M. R. Meswani, P. J. Teller, S. Arunagiri. Submitted to 2010 IEEE International Symposium on Workload Characterization (IISWC-2010), Atlanta, GA, December 2–4, 2010.
- Preparing Students to Meet Tomorrow's Challenges in Education. P. Teller. Keynote Presentation, DoDHPCMP JEOM Research Workshop, ARL-Aberdeen Proving Grounds, June 29, 2010.
- Preparation for STEM Professional Careers in the Computational Sciences in Academia. P. Teller. Plenary Presentation, DoDHPCMP JEOM Research Workshop, ARL-Aberdeen Proving Grounds, June 30, 2010.
- Extending the Monte Carlo Modeling Technique: Statistical Performance Models of the Niagara 2 Processor. J. Cook. Paper presented at the International Conference on Parallel Processing, September 2010.

Project 4–6: Hybrid Optimization Schemes for Parameter Estimation Problems

- Symbolic Dynamics for Localization of the Subcortical Structures during Deep Brain Stimulation Surgery for Parkinson's Disease. Paper accepted: North American Fuzzy Information Processing Society (NAFIPS) Conference, El Paso, December 2010.
 - A Note on the Use of Optimal Control on a Discrete Time Model of Influenza Dynamics. Paper accepted: Mathematical Biosciences and Engineering, December 2010.
 - An algorithm for constrained l_1 minimization problems and applications. Poster presentation: Sixth Blackwell-Tapia Conference, Columbus, Ohio. November 2010.
- also: International Conference on Applied Mathematics and Informatics, San Andres island, Colombia November 2010.
- Hybrid optimization for parameter estimation problems. Demonstration at the AHPCRC booth: The International Conference for High Performance Computing (SC10). New Orleans, LA, November 2010.
 - Convex optimization in digital image processing. Presentation: 8th Joint UTEP/NMSU Workshop on Mathematics, Computer Science and Computational Sciences, The University of Texas at El Paso. El Paso, Texas. November 2010.

- Optimal control applied to a discrete influenza model. Book article and Invited Presentation, published in *Proceedings of the XXXVI International ORAHS Conference*, pp. 13–27, July 2010.
also: 8th Joint UTEP/NMSU Workshop on Mathematics, Computer Science and Computational Sciences, The University of Texas at El Paso, El Paso, Texas, November 2010,
Geoepidemiology workshop, University of New Mexico, Albuquerque, New Mexico, October 2010.
- Hybrid optimization schemes for wing modeling of micro-aerial vehicles. Velazquez, L., Argaez, M., Culbreth*, M., Sanchez, R.*, Ramirez, C.*, Hernandez IV, M.* Invited paper, User Group Conference Proceedings, *IEEE-CS Journal*, Schaumburg, IL, June 2010.
also: Presentation: International Conference on Applied Mathematics and Informatics, San Andres island, Colombia November 2010.
- A comparison of wavelet-based schemes for parameter estimation. Hernandez IV, M*, Velazquez, L., Argaez, M. Invited and published paper, User Group Conference Proceedings, *IEEE-CS Journal*, Schaumburg, IL, June 2010.
- A path following method for Large-scale and dense l1-underdetermined problems and its applications to compressed sensing. Submitted to *Mathematical Programming Computation Journal*.
- A Hybrid Algorithm for Global Optimization: Wing Modeling of Micro-Aerial Vehicles. Invited presentation, EURO Conference, Lisbon, Portugal, July 2010.
- A note on the use of optimal control on a discrete time model of influenza dynamics. Book article, to be published in *Math. Biosciences and Engineering*.
- Hybrid optimization schemes for wing modeling of micro-aerial vehicle. Pan American Workshop in Applied & Computational Mathematics, Choroní, Venezuela. June 2010.
- A path following method for Large-scale and dense l1-underdetermined problems and its applications to compressed sensing. Pan American Workshop in Applied & Computational Mathematics, Choroní, Venezuela. June 2010.
- A hybrid optimization scheme for parameter estimation problems. 6th Annual Minority Serving Institutions Research Partnerships Consortium Conference, Baltimore, MD, April 14–17, 2010.

Project 4–7: Evaluating Heterogeneous High Performance Computing for Use in Field-Deployable Systems

- A Statistical Performance Model of the Opteron Processor. Paper presented at the workshop on Performance Modeling, Benchmarking, and Simulation of High Performance Computing Systems, held in conjunction with Supercomputing 2010 (SC10). This work was partially supported by a prior AHPCRC-funded project.
- SAR backprojection implementation. Soumik Banerjee, Tomasz Tuzel. Technology demonstration, AHPCRC exhibit, SC10.
- Extending the Monte Carlo Processor Modeling Technique: Statistical Performance Models of the Niagara 2 Processor. W. Alkohani, J. Cook, R. Srinivasan. *Proceedings of the International Conference on Parallel Processing (ICPP)*, Sept. 2010.
- Extending the Monte Carlo Modeling Technique to Superscalar, Out-of-Order Architectures: The Opteron Performance Model. J.M. Cook, J.E. Cook, W. Alkohani. Submitted to International Symposium on Performance Analysis of Systems and Software (ISPASS), Sept. 2010.
- GPU Power Measurement and Modeling: A Survey of Techniques. T. Tuzel, J. Cook. New Mexico State University Technical Report, TR 10 – XXX.
- FPGA Power Measurement, Modeling, and Optimization: A Survey of Techniques and Tools. M. Guzman, S. Chacon, J. Cook. New Mexico State University Technical Report, in preparation.

★ ★ ★

AHPCRC Consortium Members



Stanford University
High Performance Technologies, Inc.
Morgan State University
New Mexico State University at Las Cruces
University of Texas at El Paso
NASA Ames Research Center



Army High Performance Computing Research Center

c/o High Performance Technologies, Inc.
11955 Freedom Drive, Suite 1100
Reston, VA 20190-5673
ahpcrc@hpti.com
703-682-5368
<http://www.ahpcrc.org>

INSIDE THIS ISSUE

AHPCRC News	2
Education and Outreach	3
Flapping Wing MAVs	5
Materials Modeling	10
HPC On Board	13
The Researchers	18
Project Briefs	19
Publications, Presentations	26



TECHNOLOGY DRIVEN. WARFIGHTER FOCUSED.

This work was made possible through funding provided by the U.S. Army Research Laboratory (ARL) through High Performance Technologies, Inc. (HPTi) under contract No. W911NF-07-2-0027, and by computer and software support provided by the ARL DSRC.