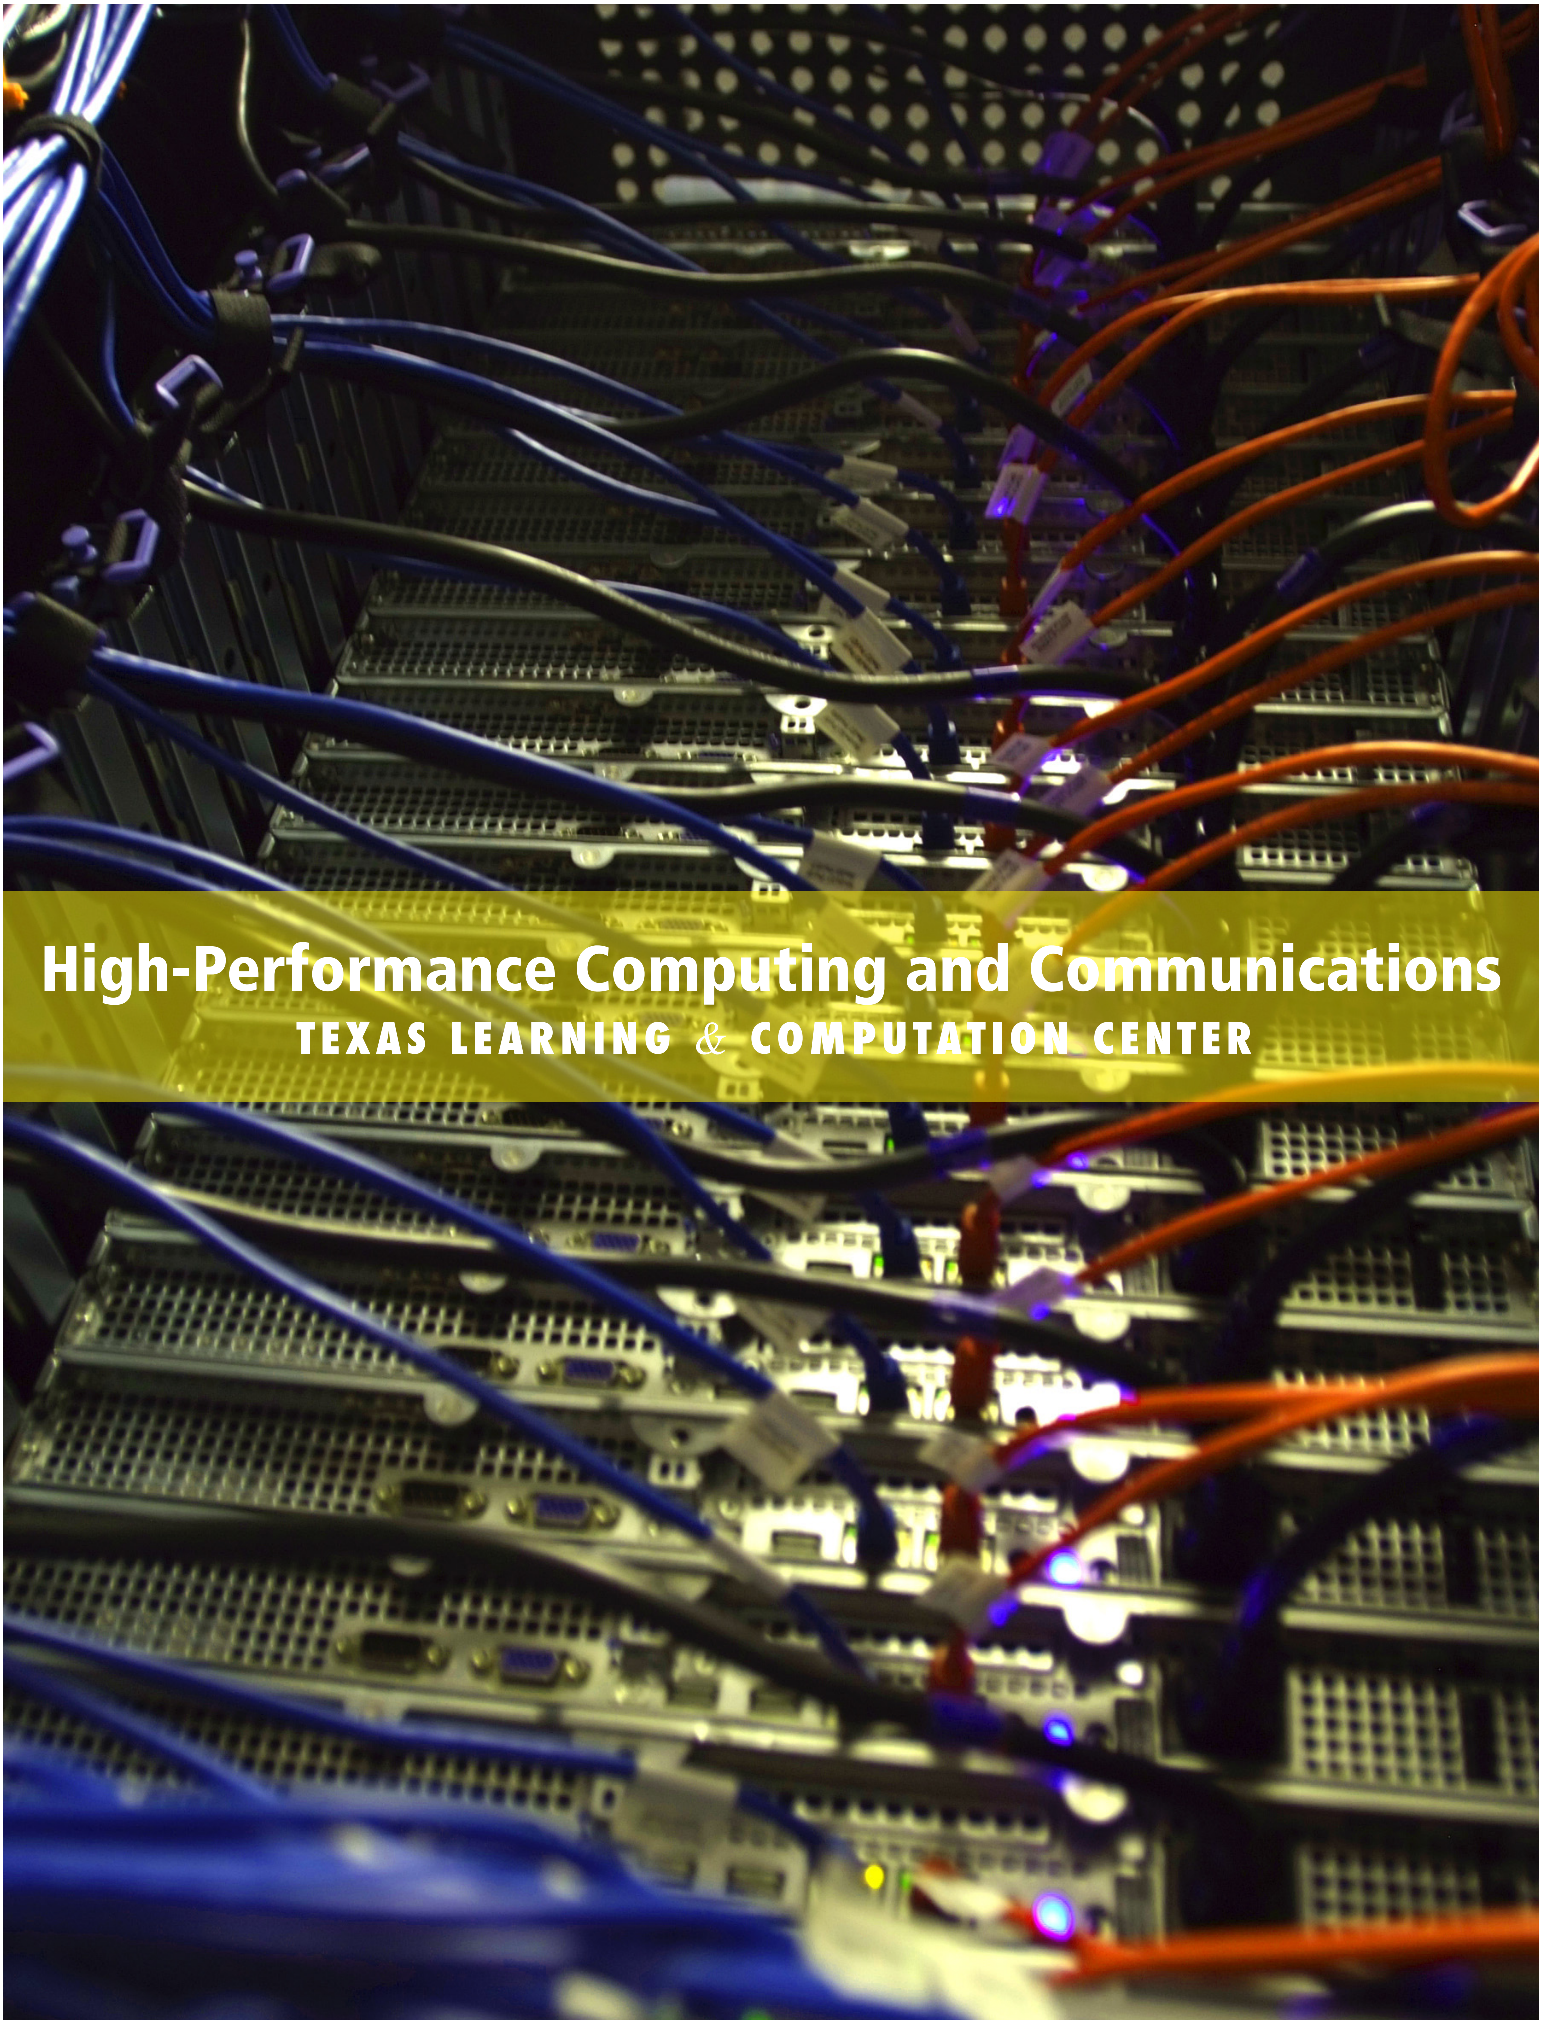

High-Performance Computing and Communications

TEXAS LEARNING & COMPUTATION CENTER

High-Performance Computing and Communications

TLC² affiliated faculty doing research in the fields of high-performance computing and communication

Robert Azencott
Professor of Mathematics

Edgar Bering
Professor of Physics and Electrical and Computer Engineering

Eric R Bittner
John & Rebecca Moores Professor of Chemistry

Barbara Chapman
Professor of Computer Science

Yuriy Fofanov
Associate Professor of Computer Science and Biology and Biochemistry

Stephen Huang
Professor of Computer Science

Lennart Johnsson
Director of TLC²,
Hugh Roy and Lillie Cranz Cullen Distinguished Professor of Computer Science, Mathematics, and Electrical & Computer Engineering

Ioannis A. Kakadiaris
Eckhard Pfeiffer Professor of Computer Science and Electrical and Computer Engineering

Donald Kouri
Hugh Roy and Lillie Cranz Cullen Distinguished Professor of Chemistry, Mathematics, Mechanical Engineering, and Physics

Ramanan Krishnamoorti
Dow Chair Professor and Chair of

Emmanouil Papadakis
Associate Professor of Mathematics

Ioannis Pavlidis
Eckhard Pfeiffer Professor of Computer Science

The Texas Learning & Computation Center (TLC²) is an interdisciplinary center at the University of Houston whose mission is to foster and support collaborative interdisciplinary research, education, and service. TLC² supports and enhances the operations of several centers, labs, and projects throughout the University, providing seed funding, high-performance computing, storage, networking and visualization facilities, high-performance networking, technical project management, research administration, and event and outreach support.

Research conducted by TLC² affiliated centers, labs, and projects covers a broad range of issues from hardware design to applications in areas ranging from life sciences, environment, energy, biology, chemistry, and physics and includes methodologies, software tools, and algorithms for high-performance computing data management, visualization, and communication. Many of the projects are multidisciplinary and involve faculty and students from different departments, centers and labs.

ADVANCED COMPUTING RESEARCH LABORATORY

Dr. S. Lennart Johnsson

The Advanced Computing Research Laboratory (ACRL) finds new and innovative ways to design and harness High-Performance Computing (HPC) resources for scientific and engineering applications. Multi-core and heterogeneous architectures using various forms of acceleration presents new challenges in high-performance software design and portability since the software must adapt to architectures as well as applications in order to operate efficiently on many platforms.

Currently, researchers are collaborating with industry and the Partnership for Advanced Computing in Europe (PRACE) in developing energy efficient HPC systems and methodologies, software tools and libraries for energy-efficient use of high-performance computing platforms. The ACRL's novel approach to the design and development of high-performance software is an integral component in collaborations with researchers, nationally and internationally, and places them at the forefront of Advanced Computing Research for tera- and peta-scale systems.

BIOMEDICAL IMAGING LAB

Dr. George Zouridakis

Large time series processing is required in many modern signal processing applications, such as functional brain connectivity analysis. Today's brain mapping machines are capable of acquiring neurophysiological electroencephalographic (EEG) signals from hundreds of locations on the scalp with sub-millisecond resolution. The most accurate methods to obtain a functional brain connectivity network rely on the solution of a large system of multivariate autoregressive models representing the simultaneous interaction of all nodes.

Biomedical Imaging Lab researchers have created a scalable OpenMP code to estimate Granger causality measures in EEG recordings from up to 256 scalp locations. Currently, lab members are developing an MPI version of the Granger causality algorithm that uses a master-worker paradigm, hoping to exploit Single Instruction Multiple Data (SIMD)-style parallelism in addition to multithreading and resulting in considerable overall speedup of the Granger connectivity algorithm.

CENTER FOR BIOMEDICAL AND ENVIRONMENTAL GENOMICS

Dr. Yuriy Fofanov

Researchers at the Center have developed efficient algorithms and data structures for genome analysis that concurrently evaluate not only exact matches of subsequences of nucleotides but also one, two, three, or more nucleotide mismatches in the subsequence. The number of subsequences of length n increases as 4n making the computational complexity grow very rapidly with the length of the subsequence and the number of mismatches considered.

The technique allows for efficient determination of statistical properties of mutations critical for evolution. It also enables the design of robust, ultraspecific probes, which lead them to propose the first-ever inexpensive, reliable, high-throughput technology to estimate the sizes of and total genomic diversity of microbial communities in any type of environmental and clinical samples. These technologies will advance scientists' understanding of the dynamic changes in microbial communities in response to such external stimuli as medical treatment, disease, and host-immune response.

COMPUTATIONAL BIOIMAGING LAB

Dr. Ioannis Kakadiaris

Researchers in the Computational Biomedicine Lab have developed highly efficient multi-modal algorithms for 3-D phase recognition and biomedical imaging applications, for instance for vulnerable plaque detection, and for biometrics applications. URxD, the Lab's face recognition system, was the top performer in the international Face Recognition Vendor Test organized by NIST in 2006. Research also include biomechanical models for soft body parts, such as the female breast, and physics-based simulation for post-mastectomy breast reconstructive surgery planning, allowing patients to communicate to surgeons preferences about the desired shape after reconstruction and for surgeons to predict (prior to surgery) the shape of a breast for a given breast tissue volume, breast implant size, shape, and placement. In this work, Distributed Approximating Functionals (DAFs) have been coupled with deformable models to obtain very efficient multiresolution deformable models.

In another effort, a theoretical framework and computational tools have been developed for non-intrusive human motion estimation from a monocular video camera for the tele-operation of ROBONAUT, an anthropomorphic robot developed at NASA JSC capable of dexterous, human-like maneuvers. Recently, work has been initiated on adapting the algorithms to GPUs and a small GPU accelerated cluster has been acquired.

COMPUTATIONAL PHYSICS

Dr. Lawrence Pinsky

One exciting opportunity enabled by the heavy ion radiation transport code FLUKA is improved cancer treatment through proton therapy. UH researchers have modified the Monte-Carlo code to directly calculate dose-equivalents as a standard scoring option for proton therapy, leading to lower dosages with better efficiency. The code is in use at Massachusetts General Hospital in Boston, and collaboration with the world-renowned University of Texas M. D. Anderson Cancer Center is under way. Other applications of the FLUKA tools include target design, calorimetry, detector design, and neutrino physics. Proton treatment also allows for PET (Positron Emission Tomography) imaging during treatment, potentially leading to real-time imaging and treatment adjustment.

B. Montgomery Pettitt
Hugh Roy and Lillie Cranz Cullen Distinguished Professor of Chemistry, Computer Science, Physics, and Biology and Biochemistry

Lawrence Pinsky
Professor and Chair of Physics

Shishir Shah
Assistant Professor of Computer Science

Pradeep Sharma
Bill D. Cook Associate Professor of Mechanical Engineering

Jaspal Subhlok
Professor of Computer Science

Ricardo Vilalta
Assistant Professor of Computer Science

Arthur Weglein
Hugh Roy and Lillie Cranz Cullen Distinguished Professor of Physics and Earth and Atmospheric Sciences

Richard Willson
Professor of Chemical Engineering

Rong Zheng
Assistant Professor of Computer Science and Electrical and Computer Engineering

David C. Zimmerman
Professor of Mechanical Engineering

George Zouridakis
Professor of Computer Science, Electrical and Computer Engineering, and Engineering Technology

Another HEP application requiring substantial computation, storage and network capabilities is ALICE, A Large Ion Collider Experiment, one of four experiments conducted at the Large Hadron Collider (LHC), In 2006, through TLC[2], UH was the second-largest contributor to the LHC grid data challenge, and has been a global Top-15 contributor during the last three years, contributing 78 percent of computer cycles and storage of the US ALICE project.

HIGH PERFORMANCE COMPUTING AND TOOLS

Dr. Barbara Chapman

Current research in the HPCTools group has focused on support for OpenMP. OpenUH is a freely available, robust compiler framework that generates native code for Pentium, Opteron and Itanium-based systems and emits heavily optimized source code that makes calls to its custom runtime library for other platforms. A specific feature of OpenUH is its ability to interact with several popular performance analysis tools in order to simplify the task of tuning OpenMP codes (as well as MPI codes or programs that use both of these programming interfaces). The tools currently include TAU, KOJAK, PerfSuite and VAMPIR. In collaboration with the NASA Ames Research Center the OpenUH compiler framework has been used to successfully improve the performance of production codes, and the usefulness of experimental language extensions for OpenMP demonstrated.

Through the participation in the DoE SciDAC Center for Programming Models for Scalable Parallel Computing led by Argonne National Lab, the group has researched language features that enhance the usefulness of OpenMP for scalable parallel computing, where many threads may be utilized to execute a program, or where OpenMP may be combined with other models. The group is also exploring ways in which the OpenUH compiler can be adapted to generate code that takes into account the resource sharing of multi-core platforms and the limited memory per thread utilizing the OpenMP 3.0 tasking features.

HIGH PERFORMANCE SYSTEMS LAB

Dr. Jaspal Subhlok

HPS researchers address some of the key issues in effective use of dynamic distributed computing environments. Skeletons have been developed for performance prediction in unreliable environments. A skeleton is a short-running program whose performance in any scenario reflects the performance of the application it represents. Skeletons can be employed to quickly estimate the performance of a large application under a new and unpredictable environment. Another initiative is Volpex that seeks to develop a framework for parallel execution on volatile nodes, such as desktop computers that are idle most of the time but can become unavailable suddenly and without notice.

Specific research directions are Volpex MPI, an MPI implementation customized for robust execution; Volpex Dataspace, focused on execution of communicating parallel programs; and Volpex Simulation for per¬formance prediction of parallel applications by creating a virtual model of the real environment. Volpex has established a campus testbed, the Virtual Campus Supercomputer Center consisting of campus PCs, computing lab machines, desktop and laptops belonging to faculty, staff, and students, and home PCs belonging to alumni.

INSTITUTE FOR DIGITAL INFORMATICS AND ANALYSIS

Dr. Donald Kouri

IDIA researchers have developed a new approach to functional approximations called the Distributed Approximating Functionals (DAFs) and use them to solve a variety of both linear and nonlinear partial differential equations. Current work includes applications in digital signal processing, imaging, data compression, pattern recognition, neural networks, and wavelet theory. For instance, DAFs have been used to develop new effective "filters" for denoising data. IDIA researchers are also working on a new generalization of the Heisenberg uncertainty principle that includes multi-resolution analysis.

This generalization is important for nonlinear optics (e.g., preparation of optimal coherent light pulses), Bose-Einstein condensation and confinement of material particles, quantum computing, and many other areas. This work has led to the development of models for generalized Gaussian and coherent quantum states. These promise to be important both theoretically and experimentally. Other efforts include time-independent wave-packet formulations of reactive scattering with different approximations applied in separate dynamical regions.

INSTITUTE OF FLUID DYNAMICS AND TURBULENCE LAB

Dr. Fazle Hussain

Understanding turbulent flows and developing techniques for control thereof have been the focus of the IFDT since its inception. The institute uses large-scale simulations in combination with experiments for validation and verification. The techniques developed apply to a large range of applications such as jet engine and aircraft design for energy efficiency but also for noise reduction.

IFDT experimental facilities include an anechoic chamber suitable for instability studies in a quiet environment and study of jet noise sources. The Holographic Particle Velocimetry, a three-dimensional flow measurement technology developed at IFDT has applications in many fields, including medical sciences where it is used for the study of blood flow, for instance in the development of the artificial heart.

INSTITUTE FOR MOLECULAR DESIGN

Dr. B. Montgomery Pettitt

The most widely used method for calculating the thermodynamic effects of solvation fields between large molecules in liquids is based on solving the Poisson Boltzmann (PB) equation by a multi-grid PDE method. The IMD's new theory uses more accurate liquid state methodology combining integral equation (IE) methods with multilevel solvers to calculate the liquid state distributions by representing the correlations as functions of the inter-particle separation and mutual angles which provide remarkable accuracy. Our numerical methods compete effectively with PB-PDE solvers, particularly with full three-dimensional (3D) dependence.

The large size of the data structure lends itself to parallelization, but the implementation of the necessary three-dimensional Fourier transforms, often split over different memory spaces, is non-trivial. With large molecular species, accelerating the approach to solution is necessary, with General Minimized Residual (GMRES) and Modified Direct Inversion of Iterative Subspace (MDIIS) methods being most commonly used. The overall structural details of the solvent contact with the hydrophobic surfaces are consistent between the different 3D IE methods and in reasonable agreement with the Molecular Dynamics simulation results of the same model. The 3D IEs are able to qualitatively, but not quantitatively, reproduce the solvent effects on the potential of mean force.

INTRUSION DETECTION RESEARCH GROUP

Dr. Stephen Huang

Computers and networks are subject to an increasing number of attacks. Intrusion detection is the process of identifying and responding to malicious activity targeted at computing and networking resources. The Intrusion Detection Research Group works in the following areas in intrusion detection and other areas of computer security: stepping-stone, where the goal is to prevent attackers from using computing resources to attack other systems; connection traceback, whose goal is to trace the connection back to the originating host; and information leaking, the goal of which is to quantify the information leaking so that it can be compared.

MISSION-ORIENTED SEISMIC RESEARCH PROGRAM

Dr. Arthur Weglein

A pressing challenge in petroleum seismology is the inability to locate and define hydrocarbon targets beneath complex media, e.g., salt, basalt, and karsted sediment. M-OSRP is pioneering and developing a direct response to this challenge. Objectives include the removal of free surface and internal multiples and depth imaging and inverting primaries, without the traditional need for subsurface information above the target, including a velocity model.

M-OSRP researchers have developed the most effective and comprehensive set of algorithms for removing multiples able to accommodate complex heterogeneous ill-defined media while requiring absolutely no subsurface information. Current goals include accurate depth imaging and inverting primaries beneath a complex overburden without knowing or determining any property of the medium above the target. M-OSRP represents and provides a fundamentally new and effective response to the pressing challenges of exploration seismology. M-OSRP has significant collaboration and interaction with its petroleum industry sponsors.

PATTERN ANALYSIS LABORATORY

Dr. Ricardo Vilalta

The massive amounts of data being collected in many disciplines require new approaches for effective knowledge extraction. PAL researchers are developing new data mining techniques that can search for domain specific patterns efficiently, with a current focus on data from physical and astronomical experiments, such as data from the CMS and ALICE experiments at the Large Hadron Collider, and from telescopes and satellites. One effort seeks to create methodologies to automate the process of signal enhancement in identifying jet streams in particle physics by multivariate classification techniques. In another effort methods from the fields of digital terrain analysis, data mining, and computer vision are employed and enhanced to produce algorithms for automatic identification and characterization of landforms on Mars to map the topography and better understand the climate of Mars.

Other efforts seeks to extend our knowledge of how humans learn by studying brain function, behavior and detect cognitive impairment through continuous non-invasive monitoring of human physiology, analyze facial expressions and the underlying cognitive state, and to improve biometrics-based security. The research includes methods for creating robust databases for millions of entries with indexing and advanced search capabilities, and design, implementation and test of algorithms for finding domain specific patterns in databases.

THEORETICAL BIOLOGICAL AND SOFTMATTER PHYSICS

Dr. Margaret Cheung

The group uses extensive simulations with molecular dynamics codes such as NAMD and Gromacs to study the effects of macro-molecular crowding on protein-folding dynamics and protein shape. In one computational study, an off-lattice model for an apoflavodoxin protein and hard-sphere particles for Ficoll 70 showed that, in the presence of 25 percent volume occupancy of spheres, native flavodoxin is thermally stabilized and the free energy shifts to favor compact structures in both native and denatured states. The native-state compaction originates in stronger interactions between the helices and the central beta-sheet, as well as less fraying in the terminal helices. This study was the first to demonstrate that macromolecular crowding has structural effects on the folded ensemble. Protein folding dynamics is also being studied in nano-sized confinements with interacting wall-proteins.

Another area of research is hydrophobic interactions in hydrocarbons and surfactants at multiple scales ranging from a nano-sized pore in an aquifer to a meter-sized reservoir. For this research multi-grained molecular simulations are developed to properly handle the large range of length and time scales involved.

WIRELESS SYSTEM RESEARCH GROUP

Dr. Rong Zheng

The Wireless System Research group (WiSeR) focuses on advancing the fundamental understanding and development of efficient robust solutions for wireless data networks. Projects span core system building blocks, data and network management tools and services, mathematical principles, and application of wireless networks to scientific, industrial, and health-care disciplines. One project seeks to develop theoretical models and algorithms for robust resource management in Industrial, Scientific and Medical (ISM) bands to minimize the outage and/or disruption of service levels under varying resource availability in 802.11 like networks.

Another project for Structure Health Monitoring (SHM) seeks to develop a scalable framework and efficient algorithms for cooperative active wireless sensing systems with a focus on coverage models for active sensing under realistic domain-specific constraints and energy efficiency in cooperative active sensing and processing. Other projects address theoretical as well as system aspects of designing efficient data dissemination and aggregation protocols for multihop wireless networks, and localization using indoor wireless networks.