



**NATIONAL SCIENCE FOUNDATION**  
Division of Computer and Network Systems  
CISE Computing Research Infrastructure  
**CRI WORKSHOP**



**SNOWBIRD UTAH**  
**JUNE 23-25, 2006**





**NATIONAL SCIENCE FOUNDATION**  
**Division of Computer and Network Systems**  
**CISE Computing Research Infrastructure**  
**CRI-WORKSHOP**

**PROGRAM ORGANIZATION**

Rita Rodriguez, CISE, NSF  
Steve Mahaney, CISE, NSF  
Malek Adjouadi, FIU

**LOGISTICS**

Dana Neill  
Computing Research Association

**INVITED SPEAKERS**

Peter Freeman, Assistant Director CISE, NSF  
Deborah Crawford, Deputy Assistant Director, CISE, NSF

**PANELS**

**CISE and the Global Implications:**

Larry Peterson, Princeton; David Novick, UTEP; Patrick Crowley, Washington U. St. Louis;  
Ram Dantu, U. of North Texas; Debby Cheng, Michigan State U.; and Bryant York, Portland  
State U.

**The GENI Program**

Guru Parulkar, GENI Program Director, CISE, NSF

**The Industry/University Cooperative Research Centers (I/UCRCs)**

Alex Schwarzkopf, Program Director, NSF, Division of Eng Education & Centers  
Sayfe Kiaei, Arizona State University, Director of Connection One and WINTech.

**ORGANIZATION OF PROCEEDINGS**

Malek Adjouadi with thanks to the FIU-CATE  
students Magno Guillen and Javier Delgado.

**PHOTOS ON THE COVER PAGE**

Courtesy of the Cliff Lodge  
Snowbird Ski & Summer Resort

**WITH THANKS**

To all of the investigators, presenters and  
participants of this NSF Workshop.

SNOWBIRD UTAH  
JUNE 23-25, 2006





**NATIONAL SCIENCE FOUNDATION**  
**Division of Computer and Network Systems**  
**CISE Computing Research Infrastructure**  
**WORKSHOP - AGENDA**

**Friday: June 23**

Registration: 4:00 PM – 6:00PM  
Reception: 6:00 PM  
Dinner: 7:00 PM

**Saturday: June 24**

7:30 - 8:00 Breakfast

8:00 - 8:20 Welcome  
Steve Mahaney, CISE/CNS Program Director, NSF  
Rita Rodriguez, CISE/CNS Program Director, NSF  
Malek Adjouadi, Florida International University

8:20 – 10:20 Session I: First-Year Awards

Sabharwal, Ashutosh	Rice University
Shieber, Stuart	Harvard University
Callan, Jamie	Carnegie Mellon University
Manikopoulos, Constantine	Foundation @ NJIT
Chandra, Namas	Florida A&M University
Raju, G.V.S.	Univ. Texas San Antonio

10:20 - 10:40 Break

10:40 - 11:10 Invited Presentation  
The GENI Program  
Parulkar, Guru CISE/GENI Program Director, NSF

11:10 - 11:40 Invited Presentation  
The Industry/University Cooperative Research Centers (I/UCRCs)  
Schwarzkopf, Alex, Program Director, NSF, Division of Eng Education & Centers  
Kiaei, Sayfe, Arizona State University, Director of Connection One and WINTech.

11:40-12:20 Session II: First-Year Awards Cont.

Pustejovsky, James	Brandeis University
Turner, Jonathon	Washington U at St Louis

12:20 - 12:30 Break

- 12:30 - 1:30 Lunch
- 1:30 – 2:10 Invited Session: CISE and the Global Implications  
 Panel: Larry Peterson, Princeton; David Novick, UTEP; Patrick Crowley, Washington U. St. Louis; Ram Dantu, U. of North Texas, Betty Cheng, Michigan State U.; and Bryant York, Portland State U.
- 2:10 – 3:10 Session III: First-Year Awards Cont.  
 Henderson, Thomas U Washington  
 Wawrzynek, John U of Cal Berkeley  
 Schulzrinne, Henning Columbia University
- 3:10 – 3:30 Break
- 3:30 – 4:30 Session IV: First-Year Awards Cont.  
 Massey, Daniel Colorado State U  
 Hwu, Wen-mei U of Ill Urbana-Champaign  
 Dantu, Ram University of North Texas
- 4:30 – 5:10 Session V: Completing Awards  
 Peterson, Larry Princeton University  
 Hauck, Scott U of Washington
- 5:10 - 5:30 CISE - Overview  
 Peter Freeman, Assistant Director CISE, NSF
- 5:30 – 6:00 Break and Poster Preparation
- 6:00 – 7:00 Poster Session
- 7:00 – 7:30 Break
- 7:30 Dinner

## Sunday: June 25

- 8:00 - 8:30 Breakfast
- 8:30 – 9:50 Session VI: Completing Awards Cont.  
 Psarris, Kleanthis U of Texas at San Antonio  
 Griffioen, James U of Kentucky Res Fdn  
 Zhang, Aidong SUNY Buffalo  
 Nava, Patricia U of Texas El Paso

9:50 – 10:20 Break

10:20 – 11:40 Session VII: Planning Awards

Note: These are planning grants awarded in August 2005. Those grants awarded in March 2006 were not included this year for presentations

Becerra-F., Irma	Florida International University
Argamon, Shlomo	Illinois Inst of Tech
Dua, Sumeet	Louisiana Tech University
Nance, Kara	U of Alaska Fairbanks

11:40-12:00 Closing Remarks and Planning for Next Year Meeting

Deborah Crawford, Deputy Assistant Director, CISE, NSF  
 Steve Mahaney, CISE/CNS Program Director, NSF  
 Rita Rodriguez, CISE/CNS Program Director, NSF

## TABLE OF CONTENTS

<b>Arizona State University</b> RI: An Interdisciplinary Research Environment for Motion Analysis _____	<b>1</b>
<b>Boston University</b> RI: Sensorium: Research Infrastructure for Managing Spatio-Temporal Objects in Video Sensor Networks _____	<b>6</b>
<b>Brandeis University</b> in collaboration with U. of Colorado, U. of Pennsylvania and U. of Pittsburgh CRI: Towards a Comprehensive Linguistic Annotation of Language _____	<b>12</b>
<b>California Institute of Technology</b> RI: Wide-Area-Network in a Laboratory _____	<b>17</b>
<b>City College of New York - CUNY</b> CRI: Center for Perceptual Robotics, Intelligent Sensors and Machines at CCNY _____	<b>22</b>
<b>Colorado State University</b> in collaboration with Michigan State University CRI: Collaborative Project: Repository for Model Driven Development (ReMoDD) _____	<b>27</b>
<b>Colorado State University</b> in collaboration with University of Memphis, University of Oregon, and UCLA CRI: Collaborative Research: Building the Next-Generation Global Routing Monitoring System _____	<b>32</b>
<b>Columbia University</b> in collaboration with Lucent Bell Labs, Penn State University and University of California at Santa Barbara CRI: WORKIT: A Universal Wireless Open Research KIT _____	<b>37</b>
<b>Florida A &amp; M University-Florida State University</b> CRI: A Linux Cluster for Research in Computational Nanotechnology and Education in Parallel Computing _____	<b>42</b>
<b>Florida International University</b> in collaboration with Emory University CRI: Enabling Globally-Distributed Electronic Collaboration (GloDEC) for Expertise Studies and Human and Social Dynamics _____	<b>47</b>
<b>Florida International University</b> in collaboration with Florida Memorial University MII: Infrastructure for Research and Training in Database Management for Web-based Geospatial Data Visualization with Applications to Aviation _____	<b>52</b>

<b>Florida International University</b>	
MII: Hardware-Software Integration for the Design of Real-Time Prototypes Merging Assistive Technologies to Neuroscience _____	<b>57</b>
<b>Harvard University</b>	
CRI: Infrastructure for Multi-Agent Decision-Making Research _____	<b>62</b>
<b>Howard University</b>	
MII: Research Infrastructure Improvement _____	<b>67</b>
<b>Hunter College, City University of New York</b>	
CRI: Planning SWAMI – A Model of Support for a Women and Minority PhD Pipeline _____	<b>72</b>
<b>Indiana University at Bloomington</b>	
CRI: A Research Infrastructure for Collaborative High-Performance Grid Applications _____	<b>77</b>
<b>Louisiana Tech. University</b>	
CRI: Planning a Heterogeneous Data Repository for Computing Research _____	<b>82</b>
<b>New Mexico State University</b>	
CRI: Infrastructure for Networked Sensor Information Technology _____	<b>87</b>
MII: Computing Support for the Next Generation Declarative Programming Systems _____	<b>92</b>
MII: Frameworks for the Development of Efficient and Scalable Knowledge-based Systems _____	<b>97</b>
<b>North Carolina A&amp;T State University</b>	
CRI: Advanced Sensor Network Applications For Environmental Monitoring Systems _____	<b>102</b>
<b>Ohio Sate University</b>	
RI: High-End Computing and Networking Research Testbed for Next Generation Data Driven, Interactive Applications _____	<b>107</b>
<b>Pennsylvania State University</b>	
RI: I <sup>3</sup> C: An Infrastructure for Innovation in Information Computing _____	<b>112</b>
<b>Princeton University</b>	
RI: Pervasive Computing: Applications and Systems _____	<b>117</b>
CRI: PlanetLab: A Community Resource Development Program _____	<b>122</b>
<b>Rice University</b>	
CRI: Wireless Open-Access Research Platform (WARP): A Scalable and Extensible Testbed for High-Performance Wireless Systems _____	<b>127</b>
<b>Rutgers University</b> in collaboration with Illinois Institute of Technology	
CRI: Community Resources for Research in Automated Authorship Attribution _____	<b>132</b>

<b>State University of New York at Buffalo</b>	
RI: MultiStore: A Research Infrastructure for Management, Analysis and Visualization of Large-Scale Multidimensional Data Sets _____	137
<b>Texas A&amp;M University – Corpus Christi</b>	
MII: Improving the Pipeline in Applied Computer Science _____	142
<b>University of California, Berkeley</b>	
RI: UC Berkeley Wireless Research Infrastructure _____	147
CRI: RAMP: Research Accelerator for Multiple Processors - A Community Vision for a Shared Experimental Parallel HW/SW Platform _____	152
<b>University of California, Los Angeles</b>	
CRI: EMSTAR: A Community Resource For Heterogeneous Embedded Sensor Network Development _____	157
<b>University of California, San Diego</b>	
CRI: FWGrid: A Research Infrastructure with Fast Wireless, Wired, Compute, and Data Infrastructure for Next Generation Systems and Applications _____	162
<b>University of Chicago</b> in collaboration with Toyota Technological Institute at Chicago	
CRI: Standard ML Software Infrastructure _____	167
<b>University of Connecticut</b>	
CRI: Computing Infrastructure for the UConn Health-Grid Initiatives _____	172
<b>University of Illinois, Urbana-Champaign</b>	
CRI: A Configurable, Application-Aware, High-Performance Platform for Trustworthy Computing _____	177
<b>University of Kentucky, Lexington</b>	
RI: The Metaverse: A Laboratory for Digital Media Networks _____	182
<b>University of Maryland</b>	
RI: High-Performance and Visualization Cluster for Research in Coupled Computational Steering and Visualization for Large-Scale Applications _____	187
<b>University of North Texas</b> in collaboration with Purdue, UC Davis, and Columbia Univ.	
CRI: A Testbed for Research and Development of Secure IP Multimedia Communication Services _____	192
<b>University of Puerto Rico at Mayaguez</b>	
MII: WALSAIP: wide Area large scale automated information processing _____	197
<b>University of Texas at Austin</b>	
RI: Mastodon: A Large-Memory, High-Throughput Simulation Infrastructure _____	202

<b>University of Texas at El Paso</b>	
MII: UTEP Distributed Computing Lab: Creation of a Pipeline for Minority Participation in Research _____	<b>207</b>
<b>University of Texas at San Antonio</b>	
CISE MII: Research Experience for Minority Students in High Performance Computing and Communications _____	<b>212</b>
CRI: Experimental Research in High Performance Computing and Wireless Networks _____	<b>217</b>
<b>University of Southern California</b> in collaboration with University of California Berkeley	
CRI: DETER Cyber Community of Researchers (DECCOR) _____	<b>222</b>
<b>University of Washington</b> in collaboration with ICSI Center for Internet Research and Georgia Institute of Technology.	
CRI: Developing the Next-Generation Open-Source Network Simulator (ns-3) _____	<b>227</b>
<b>Vassar College</b> in collaboration with Princeton University and International Computer Science Institute	
CRI: An Open Linguistic Infrastructure for American English _____	<b>232</b>
Author Index _____	<b>237</b>

# RI: An Interdisciplinary Research Environment for Motion Analysis

0403428, 09/2004-08/2009

Arizona State University

Gang Qian, Arts, Media and Engineering Program and Dept. of Electrical Engineering  
gang.qian@asu.edu

Thanassis Rikakis	Arts Media and Engineering Program	t.rikakis@asu.edu
Sethuraman Panchanathan	Computer Science and Engineering Dept.	panch@asu.edu
Jiping He	The Harrington Bioengineering Dept.	jiping.he@asu.edu
Michael McBeath	Dept. of Psychology	mmcbeath@asu.edu

Project Website: <http://ame.asu.edu/irema>

## 1. Abstract

Motion analysis is a complex problem due to the 3D nature of the human body; the infinite possibilities of human movements; variability of movement execution between different people; continuously adaptive learning through feedback from and interactions with the environment; and the inherent multiple levels of movement structure in terms of time, space and energy. This makes it unrealistic for one single discipline to address all aspects. To address this issue Arizona State University (ASU) has founded the Interdisciplinary Research Environment for Motion Analysis (IREMA) initiative that integrates researchers from eleven disciplines to create a holistic model for motion analysis research and education. Key research areas of IREMA include multimodal sensing and extraction of low, middle and high level features of movement. Key application areas include rehabilitation, K-12 education, everyday systems and culture and arts. This research infrastructure grant has allowed us to overcome significant infrastructure constraints, achieve significant progress in motion analysis and core computer and information science and engineering areas, advance our application areas and embed their results in the community and create a new model of interdisciplinary research and education with broad impact.

## 2. Introduction/Background

The complex and multifaceted nature of motion analysis demands research initiatives that cross disciplines. Such initiatives allow for the transfer of solutions among disciplines, efficient distribution of resources, decreased doubling of effort and emergence of new hybrid research models. The IREMA project at ASU exemplifies this model of collaborative research and education.

IREMA is part of the Arts, Media and Engineering Program (AME) at ASU. AME, focusing on the computation and communication of experiences, brings together 57 faculty members, more than 40 students, and other researchers from a wide range of disciplines including Computer Science, Bioengineering, Mechanical and Aerospace Engineering, Electrical Engineering, Dance, Music, Theatre, Visual Arts, Psychology, Anthropology, Sociology, Architecture and Design, and Education. This interdisciplinary collaboration has been formalized into a shared curriculum (in the form of graduate concentrations) across all participating Departments. In 2002, AME researchers from all participating disciplines came together to form the basis for IREMA and together offer a unique and powerful palette of perspectives to address existing research questions in motion capture and analysis, create new approaches and tools, and identify new areas of inquiry.

IREMA approaches the complex task of motion analysis through several levels of interrelated networks of experts, locations, infrastructures, and applications. Our network of experts comprises all disciplines participating in AME. Our network of locations is comprised of the Motion Analysis (MA) Lab, acting as a central hub, and a number of connected, satellite labs: Motor Control and Rehabilitation Lab, Center for Cognitive Ubiquitous Computing, Psychology Cognition Lab, Robotics and Structures Lab, and Digital Conducting Lab. It allows for an exchange of resources where ideas, algorithms and software/hardware equipment are being shared between different laboratories.

### 3. Goals, Objectives, Targeted Activities

IREMA conducts cutting-edge movement sensing and analysis research integrated into applications in the areas of health, education, everyday systems, culture and arts.

Our goals center on holistic multimodal movement analysis at all levels of movement structure (low, middle and high), with the following objectives:

- To create integrated multimodal sensing infrastructure (marker-based motion capture, video camera array, pressure sensing, audio tracking, physiological sensing) capable of synchronized data capture;
- To develop robust algorithms to extract, track and recognize low, middle and high level movement features at multiple temporal and spatial scales. These features range from joint angles, to poses and gestures, people/object location, shape and effort qualities and activity patterns in group movement.
- To train and develop low cost, unobtrusive movement sensing and analysis systems;
- To customize the sensing and analysis systems developed for various applications.

The major IREMA *Infrastructure* development includes the design and implementation of a large area pressure sensitive floor, the establishment of a multi-camera video capture system, the installation of a marker-based motion capture system, related physiological sensing equipment and high performance computing and data storage facilities, and the integration of all these systems.

The movement analysis *Research* of IREMA has six pilot projects: *Extraction and Recognition of Middle and Low Level Features of Movements*, *Vision-based Motion Capture Using Domain Knowledge*, *Archival of 3D Motion Data*, *Real-time Interactive Media Flow Architectures*, *Pressure Sensitive Floor*, and *Kinematics of Reach to Grasp Defined by Object Properties*.

The above research efforts and results are being integrated into six *Application* area projects, including: *Movement Based Interactive Environments*, *Rehabilitation and Biofeedback*, *K-12 Experiential Education*, *Modeling of Human and Robotic Heuristics for Projectile Interception*, *Structure Discovery for Dance*, and *Digital Conducting*.

### 4. Infrastructure and Facilities

The following multimodal movement sensing and analysis infrastructure is being developed by IREMA:

*Marker-based Motion Capture*: we are equipping IREMA with a 24-camera marker-based motion capture system. A 16-camera Eagle system from Motion Analysis Corporation was installed in Fall of 2004. Eight more cameras will be added in Fall of 2006.

*Video Camera Array*: an array of 24 high quality video cameras is being established at IREMA for video-based movement analysis. Three Dragonfly2 high speed (60 fps at 640x480) and two Dragonfly2 high resolution (30 fps at 1024x768) cameras have been installed; four Dragonfly Express high speed (200 fps at 640x480) are being added.

*Pressure Sensing*: we are developing a large area (22.8'-by-12'), high resolution (6 sensors/in<sup>2</sup>) pressure sensitive floor, which will cover the central area of the MA lab of IREMA. Pressure readings from each sensor are represented by one byte, i.e. 0 to 255. The entire floor will be finished by August 2006.

*Physiological Sensing*: two physiological sensing equipments are being deployed, including a Telemyo 2400 R 16 Channel Wireless EMG system used in the *Rehabilitation and Feedback* project and a COSMED K4b<sup>2</sup> telemetry metabolic system.

*Microphone Array*: an 8-microphone array has been developed in the *K-12 Experiential Education* project to conduct audio-based tracking and people activity inference such as conversation.

*Computing and Storage*: 24 high performance workstations and two servers are to be deployed at IREMA, mainly for the implementation of vision-based motion capture system. 4 high performance dual-CPU Dell Precision workstations have been purchased and more will be added in the summer. IREMA will also be supported by large capacity (~21 terabytes) high throughput (160-200 MB/s) storage equipment. An XserveRAID (~7 terabytes) together with SAN client software is already installed. A similar system will be added in Fall 2006.

### 5. Project Description

*Multimodal movement sensing and analysis*

Our goal is to develop an integrated multimodal sensing (marker-based motion capture, video camera array, pressure sensing, audio tracking, physiological sensing) and analysis environment. This environment will be able to

concurrently capture a wide range of human movements (from everyday walking to highly structured dance movement) in multiple scales of space (from group dynamics in large areas to fine muscle activities and weight distribution of individuals), and time (from a simple gesture of a stroke patient to complex movements of large group over a period of 30 minutes). We are working on the development of a robust pose recognition system using labeled or unlabeled markers or simply videos. We are creating robust multiple (with/without markers) people/object tracking systems using multiple cameras. We are developing computationally efficient framework using hidden Markov model and dynamic programming alignment for gesture recognition, which is also applicable to other general statistical pattern recognition problems. Movement analysis is also being tackled through the fusion of multiple sensing modalities, such as the pressure sensitive floor and marker-based system, pressure sensitive floor and video camera array, and video camera and microphone arrays. A major outcome of this project will be a low-cost, non-obtrusive system for human movement sensing and activity analysis. Education, rehabilitation, everyday systems culture and arts are the major application areas of this project.

#### *Rehabilitation and Biofeedback*

We have been working on developing an experiential biofeedback system for stroke patients. The system integrates task dependent physical therapy and cognitive stimuli within an interactive, multimodal environment. This biofeedback system allows for more efficient rehabilitation. It decreases duration, enhances results, improves performance evaluation accuracy, and increases patient engagement and enjoyment. It also opens the possibility for remote, daily home-based rehabilitation. The key motion analysis challenges are real-time low-cost, detailed tracking and analysis of arm movement, patient customization of motion analysis engine, and real-time evaluation of patient performance.

#### *K-12 Experiential Education*

Our vision is to develop a user centered, hybrid physical-digital learning environment (SMALLab), promoting collaborative, embodied learning through many sensory modalities and freeing children from the keyboard, mouse and screen. Such learning environments increase understanding of complex concepts. It prepares children for a productive and rich experience in a hybrid physical-digital world. Our specific research goals for SMALLab include: extraction of movement activity patterns of children in the learning environment; inference of engagement and learning from such movement activity and the correlation of movement activity to feedback; user customization of motion analysis engine; development of dynamic learner profiles; creation of adaptive feedback; design of learning modules; development of individual assessment and design of hybrid system evaluation strategies. We emphasize experiential learning through movement and sound as these are primary interaction modes for children. We are developing different learning modules for different desired outcomes. Extensive outreach activities are being realized in collaboration with local schools.

#### *Culture and Interactive Arts*

This project is focused on the development of interactive environments for movement based performance. Based on the results from multimodal sensing and analysis project, we are developing new pieces for dance and interactive multimedia. Poses, dynamic gestures, shape and effort qualities, phrase recognition and classification and patterns of activity in groups of movers, are being used as communication cues between the performers and the responsive environment. The novelty of this research is that it allows communication between the user and computing (sensing and modeling) at the level of semantics arising out of syntax and form. The advances made by this project in semantic level movement-based human computer interaction can be used in facilitating new intuitive human computer interfaces.



Figure 1: A subject interacting with the multimodal biofeedback system for rehabilitation



Figure 2. Students and Instructor engage in SMALLab



Figure 3: A snapshot of a movement-based interactive dance piece developed at IREMA.

## 6. Indicators of Success and Major Accomplishments

The following summaries of accomplishments include references (e.g. [P1]) to the Publication section.

### *Multimodal movement sensing and analysis*

We have developed robust and efficient algorithms for middle-level human movement analysis such as reliable tracking of people [P4], autonomous human skeleton modeling [P2], kinematic recovery [P6,P8,A4], pose [P1,P5,P9] and dynamic gesture recognition, segmentation and phrasing [P3,S2] using marker-based or video-based systems. We have also conducted research on the quantitative analysis of higher level features of human movement through novel approach such as computational extraction of Laban movement qualities. Our research also resulted in efficient theoretical computational models for statistical pattern recognition [P7, S2].

### *Rehabilitation and Biofeedback*

In the past two years of the grant period, we have laid down critical foundations for the development of experiential rehabilitation systems through real-time sensing, online performance evaluation and multimodal feedback [P12, S1]. We have made significant progress in the areas of analysis, feedback and validation metrics. A full version of an interactive biofeedback system for stroke patients is now available in IREMA. A large longitudinal trial (~one year) with a group of patients from collaborating hospitals is under way.

### *K-12 Experiential Education*

Four core learning modules have already been implemented: Sonic Gesture Discovery, Video Harvesting, Sound Poem, and Spring Sling. Each module has its own customized activities, learning objectives, relevant educational theories, technological framework, evaluation strategies and connections to STEM curricula. We have undertaken workshop sessions and a six-week pilot study with middle school students through after-school sessions in our lab at IREMA. We have fostered numerous partnerships with local school districts and museums for on-site installation that support our proposed evaluation activities. We are pursuing second-generation innovations in our learning technologies based on computational frameworks for multimodal sensing, multimodal data fusion and activity analysis, interaction context modeling, and multimodal feedback, and computational assessment strategies based on the emergent syntax of interaction. Our project was awarded as a finalist for the prestigious Peoria Prize in Creativity ([www.peoriaprize.org](http://www.peoriaprize.org)), lending it tremendous exposure amongst the arts community at large and also industry professionals. A few scientific papers have resulted from the work of this project and have been published [P10], accepted [A5] or pending for acceptance [S4].

### *Culture and Interactive Arts*

Three movement-based interactive dance pieces have been created and successfully presented nationally and internationally with significant media attention. Attendance in presentations has exceeded 2000 people. Several articles have been published about this project in major news publications internationally (e.g. New York Times, New Scientist, Dance Magazine). We are presenting a lecture/demonstration at the upcoming Congress on Research Dance (CORD) conference. CORD is the premiere international dance research organization. We have been invited to present our work at the Monaco Dance Forum in December 2006. The Monaco Dance Forum is a premiere international modern dance and technology venue and receives a wide range of European press interested in technology and education. A number of papers [P1, P4, P5, P9, A3, S3] have been produced as well.

## 7. Publications/Patents

Through the IREMA project, 31 scientific papers have been archived in various journals and conference proceedings. Seven more are in press and eight in submission. The lists below include representative publications showing the accomplishments mentioned above as well as those from other IREMA projects. A full list of related publications can be found at the project website (<http://ame.asu.edu/irema>).

### Published

- [P1] G. Qian, F. Guo, T. Ingalls, L. Olson, J. James, T. Rikakis, "A Gesture-Driven Multimodal Interactive Dance System", Proceedings of IEEE International Conference on Multimedia & Expo, p. 1579, vol. 3, (2004).
- [P2] S. Rajko, G. Qian, "Autonomous Real-Time Model Building for Optical Motion Capture", Proceedings of the IEEE International Conference on Image Processing, p. 1284, vol. III, (2005).
- [P3] S. Rajko and G. Qian, "A Hybrid HMM/DPA Adaptive Gesture Recognition Method", Proceedings of International Symposium on Visual Computing, p. 227, (2005).
- [P4] D. Whiteley, G. Qian, T. Rikakis, J. James, T. Ingalls, S. Wong, L. Olson, "Movement Analysis for Interactive Dance Using Motion Capture Data, Part I: Real-Time Tracking of Multiple People from Unlabelled Markers", Proceedings of British Machine Vision Conference, (2005).
- [P5] Y. Wang, G. Qian and T. Rikakis, "Robust Pause Detection Using 3D Motion Capture Data For Interactive Dance", Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, p.

- 1173, vol. 2, (2005).
- [P6] F. Guo and G. Qian, "Singularity Detection and Consistent 3D Arm Tracking Using Monocular Videos", Proceedings of International Conference on Image Analysis and Recognition, p. 844, (2005).
- [P7] F. Guo and G. Qian, "Sample-Efficiency-Optimized Auxiliary Particle Filter", Proceedings of IEEE Workshop on Statistical Signal Processing, p. 393, (2005).
- [P8] F. Guo and G. Qian, "Robust Contour Line Extraction Using Context", Proceedings of International Conference on Acoustics, Speech, and Signal Processing, (2006).
- [P9] F. Guo and G. Qian, "Dance Posture Recognition Using Wide-baseline Orthogonal Stereo Cameras", Proceedings of IEEE International Conference on Automatic Face and Gesture Recognition, p. 481, (2006).
- [P10] S. Benton, A. Spanias, K. Tu, H. Thornburg, G. Qian, and T. Rikakis, "Adaptive Beamforming and Particle Filtering", Proceedings of The Third IASTED International Conference on Signal Processing, Pattern Recognition and Applications, (2006).
- [P11] P. Srinivasan, D. Birchfield, G. Qian and A. Kidane, "Design of a Pressure Sensitive Floor for Multimodal Sensing", Proceedings of the International Conference on Information Visualization, p. 41, (2005).
- [P12] H. Huang, T. Ingalls, L. Olson, K. Ganley, T. Rikakis, J. He, "Interactive Multimodal Biofeedback for Task-Oriented Neural Rehabilitation", Proceeding of 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, p. 2547, (2005).
- [P13] L. Peng, K. Candan, K. Ryu, K. Chatha, H. Sundaram, "ARIA: An Adaptive and Programmable Media-flow Architecture for Interactive Arts", Proceedings of ACM Multimedia Conference, (2004).

## 8. Outreach Programs and Supported Students

We have undertaken several diverse forms of outreach to inform the general public about the research and applications being carried out under the auspices of this project.

### *Culture and Interactive Arts*

The works of this project have been presented internationally, attended by over 2000 people and reviewed in popular press. In conjunction with these performances we developed a very detailed website (<http://ame.asu.edu/motione/>) geared to a general audience. The presentations of these three works were shot in High Definition Video in collaboration with the local PBS station (KAET).

### *Biofeedback for Rehabilitation*

We have developed collaborative relationships with Mayo Clinic and Good Samaritan Hospital for the implementation and evaluation of the project and the dissemination of the results. We are currently planning to demonstrate the system in both ACM Multimedia 2006 and Neuroscience Conference 2006, both of which are premier conferences in their respective fields.

### *K-12 Experiential Education*

We have undertaken workshop sessions and a six-week pilot study with middle school students through *Herberger College for Kids*, an after-school program supported by ASU. We are starting another installation with *Herrera Middle School*, a public school in Tempe, AZ with a primarily minority (Latino) and underserved population.

### *Skysong project*

ASU in collaboration with City of Scottsdale is developing the Skysong project (a technology park that will combine research, commercial, retail and culture activity). IREMA, as a participant of the project, will be making its results and activities available for hands-on experience to the general population. *Supported Students*

IREMA has provided access to high end infrastructure and a rich network of interdisciplinary collaborators to 40 graduate students and a few undergraduates from Computer Science, Electrical Engineering, Psychology, Mechanical Engineering, Biomechanics, Kinesiology, Music, Dance, and Visual Art. All 40 students share common equipment, space and courses. Diverse methodologies, insights and approaches from disparate disciplines are being shared for research and education purposes. Tangible results from this increased exchange can be evidenced by the growing list of IREMA publications that include interdisciplinary teams as authors and the diversity of the student participants. In addition, the central infrastructure has facilitated our efforts to provide cross-disciplinary mentorship to students.

## 9. Future and Strategic Directions

The development and implementation of low-cost, non-obtrusive movement sensing and analysis systems will be the major thrust of IREMA. Such systems will contribute to the next generation of human centric computing. They will allow movement-based human-computer communication at the level of meaning and experience. These systems and our related application activities will have significant impact in health, education, everyday living, and culture and arts.

## **RI: Sensorium: Research Infrastructure for Managing Spatio-Temporal Objects in Video Sensor Networks**

*Azer Bestavros (PI/PD), Margrit Betke, John Byers, Mark Crovella, Gene Itkis, Assaf Kfoury, George Kollios, Ibrahim Matta, Leo Reyzin, Stan Sclaroff, Rich West, and Hongwei Xi*

<http://www.cs.bu.edu/groups/sensorium>

Computer Science Department  
Boston University

1. **Abstract:** This project, developing an open sensor network research infrastructure (Sensorium), aims to catalyze fundamental advances in image and video computing, network protocols, and resource management to deal with unique spatio-temporal constraints of sensor networks. A Sensorium is composed of a sensor network of video cameras spanning several rooms, networked processing units, and a terabyte database, managed together to satisfy queries using those generated by mobile users within this environment. The infrastructure enables the following research projects: (a) Modeling, interpretation, and prediction of human motion in video streams at multiple scales in space/time and at multiple layers of detail; (b) Development of efficient location management, routing, transport, and content distribution protocols for multi-resolution/scale streaming sensory data networks; (c) Characterization of traffic and access patterns in mobile sensory networks; (d) Instrumentation of embedded real-time operating systems to enable coordinated resource management and the development of middleware services for the management of active sensor networks; (e) Indexing and mining of large spatio-temporal non-textual sensory datasets, with a particular emphasis on mining of human motions and activities; (f) Enhancing code safety for embedded systems through the use of type systems and run-time support, with emphasis on flow-oriented programming; (g) Development of algorithms and protocols for supporting security and trust, and for protecting the confidentiality and integrity of data in video sensor networks and repositories.
2. **Introduction:** Recent trends suggest that inexpensive networked video sensing elements will be pervasively deployed in our environment. Already, they are embedded in devices such as computers and mobile phones and they are mounted in public spaces such as malls and airports. In many ways, this trend could prove beneficial to society, in that information collected by sensors could be shared for the better good. Harnessing the power of these emergent sensory environments will hinge on our ability to build applications capable of gathering, interpreting and storing data from distributed sensors and to provide scalable mechanisms for managing the networks and systems resources that these applications consume. We aim to address these challenges within a research infrastructure composed of a sensor network of video cameras spanning several rooms, networked processing units and a terabyte database which are managed together to satisfy user queries including those generated by mobile users within this environment.
3. **Goals, Objectives, and Targeted Activities:** Sensors in the Sensorium infrastructure are fixed video cameras with pan/tilt/zoom capability. At any instant in time, we refer to the finite spatial extent visible to a sensor as that sensor's *view volume*; objects within that volume are recorded at a resolution which can vary due to pixel and frame rate sampling. For a group of video cameras, the union of view volumes define a *domain volume* for a monitored room or public space. Each video camera is attached to PCs for processing and indexing sensory information. A user with her mobile unit can ask the Sensorium to perform high-level, automated sensing tasks, such as locating and tracking a human moving through the Sensorium. Queries can be sent either through a wireless access point or through another user within communication range (as in *ad-hoc* networks). In response to user queries, or in cooperation with other sensors, cameras perform basic operations.




Figure -1: Main Subsystems of the Sensorium: Functional View (left) and Topological View (right).

such as adjusting their resolution or altering their view volumes. Responding to a single such query is already a challenging research problem from a networking and computer vision standpoint, but we are interested in an even broader challenge: Designing a scalable Sensorium system architecture capable of handling concurrent queries of this form which will in general place competing demands on systems resources.

In contrast to current sensory systems which are conceived (and built) as *special purpose* systems with custom-developed architectures, we envision sensoria built using commodity hardware and freely available software, opening up much wider access to sensory systems to the computer science research community. Our desire to develop sensory commodity system for a broad class of environments also highlights some of the current impediments to building such a system. Among these, we believe that (1) multi-resolution encoding of sensory data, and (2) cognizance of spatio-temporal and resource constraints at every level of a system's architecture will be critical core technologies. These two threads are common to all Sensorium research projects.

4. **Infrastructure and Facilities:** Figure -1 shows functional and topological views of the various subsystems that comprise the Sensorium infrastructure. In addition to hardware and software instrumentation, our award is partially supporting a research engineer who is assisting in software development and system integration efforts.

**5. Description of Sample Projects Enabled Through the Infrastructure:** The Sensorium has inspired and is supporting either directly or indirectly over 30 projects. Space limitations prohibit us from describing them all. The following are descriptions of representative Sensorium projects, which should give the reader a sense of the scope of our efforts in various areas, as well as the collaborations that the Sensorium has catalyzed across areas. For details on these as well as other projects, we refer the reader to the web pages of the various research groups and projects affiliated with this effort, all of which are accessible from the Sensorium main webpage.

- + *Estimating and Recognizing 3D Articulated Motion via Uncalibrated Cameras (PI: Sclaroff):* The main goal of this effort is to develop new algorithms for 3D articulated structure and motion estimation, from one or more uncalibrated video streams. Articulated motion is exhibited by jointed structures like the human body and hands, as well as linkages more generally. In this project, 3D articulated structure and motion estimation algorithms are developed that can automatically initialize themselves, estimate multiple plausible interpretations along with their likelihood, and provide reliable performance over extended sequences.
- + *Video-based Interfaces For People With Disabilities (PI: Betke):* This project aims to develop human-computer interfaces for people with severe disabilities by detecting various body cues. These interfaces are intended to run on an average PC with video input from an inexpensive USB camera. The face is tracked using multi-scale template correlation. Symmetry between left and right eyes is exploited to detect if the computer user is looking at the camera, or off to the left or

right side. The detected eye direction can then be used to control applications such as spelling programs or games. To test the system, we have developed a game 'BlockEscape' to gather quantitative results to evaluate the interface system with test subjects. Other HCI tools for people with disabilities are developed as part of this project. Examples include "BlinkLink", which automatically detects a user's eye blinks and accurately measures their durations, and the 'EyebrowClicker', which automatically detects when a user raises his or her eyebrows and then triggers a mouse click.

- + *Pattern Discovery in Signed Languages and Gestural Communication (PI: Sclaroff, Betke, and Kollios)*: The goal of this project is the development of algorithms for pattern analysis and discovery in signed languages and gestural communication. A key emphasis is on analysis of the co-occurrence, overlap, relative timing, frequency, and magnitude of linguistically significant movements of the hands, face, and upper body. Clustering and indexing algorithms are developed for identification of groups of similar gestural components from mixed discrete annotation event labels and sampled measurement data. This project involves an interdisciplinary collaboration with linguists at BU.
- + *Human Motion Mining (PIs: Sclaroff, Kollios, and Betke)*: This project develops and tests methods for indexing, retrieval, and data mining of human motion trajectories in video databases. Computer vision techniques are being devised for automatic extraction of human motion time series data from video. Data mining algorithms are being developed that can be used to discover clusters and other patterns in the extracted motion time-series data. We have developed a strategy for approximation of expensive distance measures used in indexing and nearest neighbor classification problems. Using Lipschitz embeddings, "vantage objects" are embedded in a Euclidean space in which similarities can be rapidly measured using a weighted Manhattan distance. Embedding construction is formulated as a machine learning task, where AdaBoost is used to combine many simple, 1D embeddings into a multidimensional embedding that preserves a significant amount of the proximity structure in the original space.
- + *Sketching Sensor Network Streams (PIs: Byers and Kollios)*: We consider problems associated with performing queries across sensor fields of large scale. We treat the sensor streams as a distributed database, albeit one in which approximations (or sketches) of the sensed data is stored. Our techniques are based on duplicate insensitive randomized algorithms that can be combined with robust routing algorithms in sensor networks, such as multi-path routing. The approximate representations that we consider save space, can be transmitted with smaller messaging and energy cost, and admit more flexibility when used in concert with in-network aggregation. Our methods apply from sensor applications ranging from mote-based systems to sensoria to wide-area network measurement.
- + *Scalable Real-time Delivery of Streaming Media and Sensory Data in Ad-Hoc Networks (PIs: Bestavros and Matta)*: Ad-hoc overlay networks are increasingly used for sharing static bulk content but their promise for scaling the delivery of on-demand, real-time content is yet to be tapped. Motivated by many Sensorium applications, we are investigating how ad-hoc overlay networks could be used for efficient on-demand distribution of real-time content (e.g., popular streaming media, or field measurements) to a large number of clients. In addition to leveraging the local caches at clients in the network, our approaches leverage the mobility of individual nodes to improve the timeliness and query recall rates in such systems.
- + *Sensorium Security through Fault-Tolerant Cryptography (PIs: Itkis and Reyzin)*: Traditional fault-tolerance focuses on dealing with (accidental or malicious) errors in computation. In the area of security and cryptography this focus must be extended to include as a main topic security compromises, such as exposures of confidential information, obtaining unauthorized privileges, falsifying records, repudiating contracts, etc. This is particularly important for large heterogeneous systems, such as the Sensorium, where components are geographically dispersed and may be administered by different entities with variable physical and electronic security devices. The efforts of this project are directed towards developing techniques that anticipate key compromise, detect and/or gracefully recover from it.
- + *Utilizing Biometrics and Strong Authentication Techniques in Sensorium-like Settings (PI: Reyzin)*: Humans are quite limited in their ability to generate and reliably recall the long random strings required for traditional cryptographic protocols, and hence resort to using low-security devices such as passwords. To that end, there have been proposals to use biometric information collected in Sensorium-like environments for access control, for example. Since biometric information is not precisely reproduced from one reading to the next, existing approaches require storage of the original reading in order to check if it is close to the current reading. This introduces a set of challenging security problems that we address in this project. For example, while the database of biometric data is sensitive private information that must be secured, in order for the system to be reliable and reasonably fast, the database must be always available, through replication. These two requirements—security and availability—are in conflict, particularly for systems such as the Sensorium that need to scale to large

numbers of users and access points.

- + *User-Level Sandboxing (PI: West)*: This work is motivated by the desire to allow **Sensorium** applications to customize and extend the system for their specific needs. It aims to develop safe and efficient methods for user-level extensibility of commercial off-the-shelf (COTS) systems. As with micro-kernels, user-level services and extensions can be deployed without polluting the kernel address space with potentially unsafe application-specific code. User-level sandboxing provides a clean separation of core system functionality from higher-level abstractions. Likewise, user-level code can leverage libraries and system calls, it can be rapidly prototyped without causing system failure, and it can be written without knowledge of kernel internals.
- + *Compositional Specification of Networking Systems using iBench (PIs: Bestavros, Kfoury, and Matta)*: The science of network service composition has clearly emerged as one of the grand themes driving many of our research questions in the networking field today. This driving force stems from the rise of sophisticated applications and new networking paradigms. By “service composition” we mean that the performance and correctness properties local to the various constituent components of a service can be readily composed into global (end-to-end) properties without re-analyzing any of the constituent components in isolation, or as part of the whole composite service. In this project we focus on services built upon an overlay management architecture, and we use control theory and Network Calculus as example theories from which we lift up compositional specifications that allow us to verify parallel and sequential compositions of QoS-constrained flows with various networking elements (e.g., shapers and schedulers) and protocols (e.g., TCP).
- + *snBench: Sensor Network programming workBench (PIs: Bestavros and Kfoury)*: The snBench project aims to design and implement the programming and run-time infrastructure necessary for developers to specify and deploy truly distributed applications over a heterogeneous network of Sensing Elements (SEs) and of Computing Elements (CEs). snBench offers the means via which applications may (1) locate, monitor, and query SEs and CEs for services they support, and (2) initiate, control, or otherwise use such SE and CE services. In addition to supporting these functionalities, the snBench provides basic run-time resource management services for QoS support, including real-time scheduling and admission control functionalities. One of the salient features of snBench is extensibility. To that end, snBench provides a type-disciplined framework for defining and adding new services—both atomic services at SEs and compositional services at CEs.

**6. Indicators of Success and Major Accomplishments:** The **Sensorium** infrastructure has had a significant impact on the research, education, and outreach missions of the department, and continues to be instrumental in the development of the professional careers of its faculty members and students. Space limitations prohibit us from describing the various indicators of success and major accomplishments leveraged by the **Sensorium**, so far. The following tidbits should provide the reader with a sense of the significance of these achievements.

- + *Major Technical Findings*: The specific technical findings are too numerous to enumerate in this report. Perhaps it suffices to note that over the last two years alone, **Sensorium**-enabled research has been recognized by four best paper awards, and that publications by PIs and senior personnel have garnered over 3,000 citations.
- + *New Collaborations and Funding*: Many of the projects pursued under the **Sensorium** umbrella involve collaborations across disciplines within CS (e.g., Vision & Networking, Networking & Programming Languages, Vision & Operating Systems, Vision & Databases, Cryptography & Databases, Real-Time Systems & HCI, etc.) Already these collaborations have resulted in generating new sources of research funding – totalling over \$10M in the last four years – including some highly prestigious grants to PIs and senior personnel (e.g., three ITR awards, three CyberTrust awards, two CAREER awards, and two DoED GAANN awards).
- + *Inter-Disciplinary Impact*: In addition to intra-disciplinary collaborations, the **Sensorium** is contributing to other non-CS disciplines. Examples include medical technologies for assisted living and support of people with disabilities, American Sign Language interpretation research, environmental monitoring, among others. The **Sensorium** was also a catalyst for a wide collaboration on campus on graduate training and education in Device Networks.
- + *Contributions to Curriculum Development*: The **Sensorium** (and the synergy that it created across various groups) has also resulted in specific outcomes (to date) related to curriculum development. The following are two specific examples: (1) Throughout 2003/04, PIs Bestavros, Byers, Matta, and Kollios ran a year-long seminar entitled “Sensor Networks: The Dream, The Reality, and The Science in Between”, which resulted in publications, including papers earning “best paper awards” at top conferences, and PhD research projects. This is the first time this format of a year-long, team-taught seminar was attempted and

the general consensus is that it was a successful experiment that will be repeated in future years. (2) The code base of the snBench project was used in regular CS courses – namely, Betke’s CS-440 AI class and Kfoury’s CS-511 Software Engineering class – to allow graduate and undergraduate students to “experiment” with the **Sensorium** in their class projects. The projects pursued in CS-511 centered around extending the snBench code base itself!

+ *Tangible Impact on Cyber Infrastructure*: Our CHAIN methodology, developed as part of the iBench project uncovered an important “bug” in the HTTP protocol. Specifically, we were able to identify specific canonical arrangements of HTTP proxies and settings that are prone to deadlocks. Arguably, no protocol is more pervasively deployed and in use by millions of users at any point in time than the HTTP protocol. Nevertheless, this bug had remained under the radar of network researchers until publications of our results. Other outcomes from our iBench project were also instrumental in identifying unsafe behaviors in Web cache consistency protocols, in informed packet scheduling for streaming applications, and in compositional QoS systems (now being evaluated at BBN).

- + *Technology Transfer*: (1) One of the interface systems developed by Betke and her students has been commercialized successfully; it is now used in many early intervention facilities, homes, schools, and hospitals in Massachusetts, Connecticut, New Jersey, England, and Northern Ireland. (2) Several of the results by John Byers in his project on flexible strategies for content distribution are currently deployed as part of a commercialized technology by Digital Fountain Inc. (3) The work by Crovella and his student Anukool Lakhina on network anomaly detection (patent pending) is the catalyst for a very promising start-up founded by Dr. Lakhina. (4) The eTCP work by Bestavros and Matta (patent pending) is currently licensed by a west-coast startup for use in VoIP and gaming applications. (5) The work by Sclaroff and his students spurred industrial interest (and funding) from Honda, Mitsubishi Research, and Charles River Analytics.

+ *Bridging Research and Development*: The **Sensorium** was instrumental in initiating the Sensor Network Consortium at BU (see <http://www.bu.edu/systems/industry/consortium> for details). SNC, a collaborative effort to foster research in sensor networks and facilitate the interaction among academia and industry as well as between industry participants, provides an interesting mix of companies: start-ups, system integrators, users and a variety of small and large corporations whose strategic business opportunities relate to this emerging technology. In addition to Boston University, faculty from many academic institutions are involved in the consortium activities, including faculty members at Harvard, Umass Amherst, and Yale University. Current industrial members include BP International Ltd., Echelon Corporation, Ember Corporation, Honeywell International, INETCO Systems Limited, L-3 Communications - Henschel and Homeland Security Divisions, Millennial Net, Inc., Mitre Corporation, Radianse Inc., SAP, Sensicast Systems, and Textron Systems.

- + *Broader Impact on Society*: Results from our various **Sensorium** projects are proving to be crucial for automated analysis, classification, and recognition of human movement in open **Sensorium**-like settings and applications that are valuable to society, including homeland security; video-based analysis of human biomechanics for occupational safety, as well as dance and sports training; archive management and analysis for news, entertainment, and sports video; and video

database management for non-intrusive monitoring of the motion patterns of handicapped, infirm, or elderly people to detect decline, danger, and to alert caregivers when needed.

**7. Publications, Patents and Products**: To date, projects leveraged by the **Sensorium** infrastructure have resulted in over 150 refereed journal and conference proceedings, over a dozen software artifacts and data sets, and 8 patent applications and invention disclosures. This body of publications is clearly too large to enumerate in this report (details are available through the **Sensorium**’s publication page).

To underscore the quality of our publication record, it suffices to say that these publications appear regularly in the absolute best venues in each of the fields of study covered by the **Sensorium** project. Example venues in which our papers appear regularly include Sigmod, VLDB, ICDE, and TKDE for database research; CVPR and PAMI for vision research; Crypto and Eurocrypt for security research; Sigcomm, Infocom, ICNP and ToN for networking research; RTSS, RTAS, TPDS, and ToC for real-time systems research; PoPL and PLDI for software systems research.

In addition to peer-refereed papers in journals and conferences, software artifacts and data sets from our research projects have been used by many other research groups. A recent example would be the 3D

articulated structure and motion estimation algorithms developed by Sclaroff and his students, which were implemented in Matlab and C++ (with associated training and test data sets, example demos, and documentation), and which were shared with colleagues at Ohio State U., U of Palermo and U of Florence. Another example is the TRAFFIC web-based tool for verification of safe network compositions from the iBench project, which has catalyzed in some early collaborations with BBN.

**8. Outreach Programs and Supported Students:** Outreach efforts are likely to bear fruits only if they are consistent, not over-reaching, and only if they do not insist on local (as opposed to global) successes. In particular, we view our outreach efforts to under-represented groups as successful even they result in members of such groups joining or succeeding at other institutions. Along these lines, our efforts are coordinated at multiple levels involving the individual PIs, the Department's administration, the College, and the University, all of which are spelled out in our annual reports. For brevity, we mention two recent examples targeting two very different constituents. (1) Every Spring semester, PIs Sclaroff and Betke host an afternoon visit by various Girl Scout troops. In these visits, girls are actively engaged in a series of interactive demonstrations of computer vision tracking of human motion, digital imaging methods, computer interfaces for severely disabled people, and virtual reality. This is followed by a lively discussion (led by some of our department's female CS PhD students) of CS as a potential field of study and its potential to change society. (2) Azer Bestavros is spearheading an initiative to bring members of minority groups to attend major networking conferences. To that end, he secured funding from NSF and industrial partners for travel awards to junior faculty members from under-represented groups to attend IEEE ICNP 2005. 10 minority faculty members and 4 minority graduate students benefitted from this program. The idea for this initiative, which he proposed first at the 2004 NSF RI/MII PI meeting, is that while PI meetings for programs like the RI and MII help faculty in minority institutions compare notes and build bridges with main research institutions, these meetings tend to bring researchers from very diverse subfields, making it hard for a representative from a minority institution to find more than one or two other faculty members with similar research interests to "network" with. Alternately, to foster better interactions, faculty members from minority institutions must be encouraged (and supported) to travel and attend *Top* CS conferences in their chosen areas of research.

Over 110 students (58 already graduated) worked on various **Sensorium** projects and used the NSF-funded research infrastructure in their research training. Of these 110 students, 24 are female (13 already graduated) and 3 were members of under-represented groups (2 already graduated). Of these 110 students, 49 students are in our PhD program (13 already obtained their PhDs), 38 students are in our MA program (28 already obtained their MAs), and the remaining 23 students are in our undergraduate program (20 already obtained their BAs).

In terms of placement, PhD graduates who have participated in **Sensorium** research have secured positions in academia, including faculty members at U of Wisconsin, NCSU, and CWRU, in research labs, including Microsoft, Google, BBN, Amazon, and in a number of startups, including ones co-founded by the students themselves. Many of our BA and MA graduates who have participated in **Sensorium** research have moved on to PhD programs, many at some of the best programs in the nation, including MIT, CMU, UCLA, among others.

**9. Future and Strategic Directions:** Our **Sensorium** initiative has focused primarily on issues of scalability, extensibility, and programmability of pervasive video sensor networks, focusing on multi-resolution spatio-temporal representation and on resource management issues. Having made significant strides along these lines, it has become evident that issues of security, privacy, and safety of such sensor-rich, open environments are paramount. In that regard, we have started on a number of projects along these lines – some of which were mentioned earlier in this report. Tackling these issues constitute what we perceive to be the next major "challenge" for **Sensorium**-like environments. We are very well poised to steer our research into this direction, leveraged by the establishment of a new center for Reliable Information Systems and Cyber Security (RISCS) research and education at Boston University, and by continued infrastructure support from NSF.

# CRI: Towards a Comprehensive Linguistic Annotation of Language

NSF-CRI-0551615

## Year 1

Brandeis University, NYU, U. of Colorado, U. of Pennsylvania, U. of Pittsburgh

James Pustejovsky (PI), Adam Meyers (co-PI), Janyce Wiebe (co-PI)

Martha Palmer (co-PI), Mitchell Marcus (co-PI)

[jamesp@cs.brandeis.edu](mailto:jamesp@cs.brandeis.edu),

[meyers@cs.nyu.edu](mailto:meyers@cs.nyu.edu),

[wiebe@cs.pitt.edu](mailto:wiebe@cs.pitt.edu),

[martha.palmer@colorado.edu](mailto:martha.palmer@colorado.edu)

[mitch@linc.cis.upenn.edu](mailto:mitch@linc.cis.upenn.edu)

URL: [www.unified-annotation.org](http://www.unified-annotation.org)

## 1. Abstract

Many recent annotation efforts for English have focused on pieces of the larger problem of semantic annotation, rather than producing a single unified representation. This project investigates the issues involved in merging several of these efforts into a unified linguistic structure: PropBank, NomBank, the Discourse Treebank, TimeBank, and opinion annotation. We discuss resolving overlapping and conflicting annotation as well as how the various annotation schemes can reinforce each other to produce a representation that is greater than the sum of its parts.

## 2. Introduction/Background

The creation of the Penn Treebank (Marcus et al, 1993) and the word sense-annotated SEMCOR (Fellbaum, 1997) have shown how even limited amounts of annotated data can result in major improvements in complex natural language understanding systems. These annotated corpora have led to high-level improvements for parsing and word sense disambiguation (WSD), on the same scale as previously occurred for Part of Speech tagging by the annotation of the Brown corpus and, more recently, the British National Corpus (BNC) (Burnard, 2000). However, the creation of semantically annotated corpora has lagged dramatically behind the creation of other linguistic resources: in part due to the perceived cost, in part due to an assumed lack of theoretical agreement on basic semantic judgments, in part, finally, due to the understandable unwillingness of research groups to get involved in such an undertaking. As a result, the need for such resources has become urgent.

Many recent annotation efforts for English have focused on pieces of the larger problem of semantic annotation, rather than producing a single unified representation like Head-driven Phrase Structure Grammar (Pollard and Sag 1994) or the Prague Dependency Tecto-grammatical Representation (Hajicova & Kucerova, 2002). PropBank (Palmer et al, 2005) annotates predicate argument structure anchored by verbs. NomBank (Meyers, et. al., 2004a) annotates predicate argument structure anchored by nouns. TimeBank (Pustejovsky et al, 2003) annotates the temporal features of propositions and the temporal relations between propositions. The Penn Discourse Treebank (Miltsakaki et al 2004a/b) treats discourse connectives as predicates and the sentences being joined as arguments. Researchers at Essex were responsible for the coreference markup scheme developed in MATE (Poesio et al, 1999; Poesio, 2004a) and have annotated corpora using this scheme including a subset of the Penn Treebank (Poesio and Vieira, 1998), and the GNOME corpus (Poesio, 2004b). The Pittsburgh Opinion annotation project marks features of “private states” (opinions, evaluations, sentiments, etc.) as they occur in text. This project will create a *Unified Linguistic Annotation (ULA)* by merging annotation of examples using the schemata from these efforts.

We describe below existing independent annotation efforts, each one of which is focused on a specific aspect of the semantic representation task: semantic role labeling, discourse relations, temporal relations, etc. They have reached a level of maturity that warrants a concerted attempt to merge them into a single, unified representation, ULA. There are several technical and theoretical issues that will need to be resolved in order to bring these different layers together seamlessly. Most of these approaches have annotated the same type of data, Wall Street Journal text, so it is also important to demonstrate that the annotation can be extended to other genres such as spoken language. The demonstration of success for the extensions would be the training of accurate statistical semantic taggers.

**PropBank:** The Penn Proposition Bank focuses on the argument structure of verbs, and provides a corpus annotated with semantic roles, including participants traditionally viewed as arguments and adjuncts.

**NomBank:** The NYU NomBank project can be considered part of the larger PropBank effort and is designed to provide argument structure for instances of about 5000 common nouns in the Penn Treebank II corpus (Meyers, et al., 2004a).

**Discourse Treebank:** The Penn Discourse Treebank (PDTB) (Miltsakaki et al 2004a/b) is based on the idea that discourse connectives are predicates with associated argument structure (for details see (Miltsakaki et al 2004a, Miltsakaki et al 2004b).

**TimeBank:** The Brandeis TimeBank corpus, funded by ARDA, focuses on the annotation of all major aspects in natural language text associated with temporal and event information (Day, et al, 2003, Pustejovsky, et al, 2004).

**Opinion Annotation:** The Pittsburgh Opinion annotation project funded by ARDA, focuses on the annotation of opinions, evaluations, emotions, sentiments, and other private states in text.

### 3. Goals, Objectives, Targeted Activities

The main goal of this project is to create community resources including: a semantic representation of the sentences in a corpus which integrates several current annotation schemata. This annotated corpus should open the door to corpus-based techniques for processing text to produce a similar structure, thus removing some of the obstacles faced by NLP and linguistic research in semantic interpretation. Specifically, we will create a *Unified Linguistic Annotation (ULA)* and a larger (550K words) and more balanced corpus than any semantically annotated corpus currently available, in which most basic types of semantic information are annotated according to high-quality schemes and using the state-of-the-art annotation technology. Crucially, all individual annotations will be kept separate in order to make it easy to produce alternative annotations of a specific type of semantic information (e.g., word senses, or anaphora) without need to modify the annotation at the other levels.

We now outline five specific subgoals of this research. First, we will design a meta-specification framework that will both (a) allow individual annotations to cohabit with one another (consistency), and (b) allow specification components from different annotation schemas to communicate with one another, in order to refer to merged information (integration). Our second aim is to advance semantic annotation technology and our understanding of semantic judgments in two ways. First of all, we will test (semi)-automatic annotation techniques for the types of semantic annotation that can already be done reliably. To this end, we will study annotator agreement on types of semantic annotation that have been shown to be useful in a variety of tasks, but for which we do not yet have reliable annotation guidelines. Annotating several different types of semantic information will enable progress on techniques for one type of semantic processing without the need to wait for the development of high-performing systems for other aspects of semantic interpretation –e.g., Gold Standard sense tags can enable the development of anaphoric resolvers relying on disambiguated word senses before accurate automatic taggers are available—in the same way that the Penn Treebank has allowed progress on anaphora resolution while waiting for the development of high-performance parsers. Our third aim follows logically from the first two, namely the production of a stable and language-independent methodology for the process of unified linguistic annotation, complete with widely accessible tools.

The fourth aim seeks to validate the generality and robustness of the ULA by incorporating additional annotation schemes, new genres, and additional languages. This will be done by solicitation associated with a series of workshops organized and funded through this research. Finally, given the standards-oriented and community-focused nature of the proposal, our fifth aim is to actively promote the dissemination of the techniques embodied in the ULA, as well as the resulting annotated corpora throughout the community, for use and further evaluation.

There is a growing community consensus that general annotation, relying on linguistic cues, and in particular lexical cues, will produce an enduring resource that is useful, replicable and portable. Our objective is not to prove the truth or correctness of any particular theoretical approach, but rather to provide a useful level of representation. We believe that a shallow dependency structure, i.e., predicate-argument structure, for verbs, participial modifiers, nominalizations, and adjectives, provides a feasible level of annotation that would be of great benefit. This annotation, coupled with word senses, minimal co-reference links, event identifiers, and discourse and temporal relations, could provide the foundation for a major advance in our ability to automatically extract salient

relationships from text. This will in turn facilitate breakthroughs in message understanding, machine translation, fact retrieval, and information retrieval.

## 4. Infrastructure and Facilities

The necessary web infrastructure will be created to enhance collaboration between all sites involved in the project. Brandeis University will host a server to allow access to a common repository of annotated corpora and annotation tools, and keep code and data under CVS version control. In addition, there will be a project web site with password-protected subsites, discussion forums, and mailing lists.

## 5. Project Description

### 5.1 Merging Annotations

The 1M words of Wall Street Journal text known as the Penn Treebank has already had several additional types of annotation added, over and above the original part-of-speech tagging and phrase structure annotation. However, these layers are not necessarily complete, and there are gaps with respect to phenomena and lexical items. During this project, annotators at various sites will focus on completing the various layers of annotation for a selected 310K word subset of the Wall Street Journal corpus. The annotation schemas will include: sense tagging, discourse connectives and selected temporal relations. This data set will comprise the original test corpus for merging the different layers in a Unified Linguistic Annotation. A major focus will be the investigation of the use of eventuality ID's for both verbal predicate argument structures and nominal predicate argument structures to allow discourse relations and temporal relations straightforward references to eventualities as arguments.

Our first step is to define a single format for a merged representation of PropBank and NomBank, effectively the core predicate argument structures and the referents for the arguments. The goal is to provide a general architecture where each type of annotation is viewed as an independent analysis that contributes new information to the whole -- all redundancies and conflicts between annotation being initially ignored. This will necessitate the definition of guidelines for converting various types of annotation (offset annotation, sgml markup, etc.) into the chosen merged format. One possibility would be to convert each annotation into features and values to be added to a larger feature structure. The resulting feature structure would be a combination of stand alone and offset annotation -- the feature structure would include actual words and features from the text as well as special features that point to the actual text (character offsets) and, perhaps, syntactic trees (offsets along the lines of PropBank/NomBank). Alternative global annotation schemes that deserve consideration include annotation graphs (Cieri & Bird, 2001), and MATE (Carletta, et. al., 1999). We will then take a sample of the multiply annotated WSJ text and experiment with automatic techniques for producing the merged representation. Each of these automatic results will have to be carefully examined for conflicts, overlaps and duplications, as discussed below. There are many areas in which the boundaries between these annotations have not been clearly defined, such as the treatment of support constructions and light verbs. Determining the most effective approach to handling these phenomena will require detailed discussion of alternative treatments for several actual examples.

Once we have reconciled our predicates and their argument references, we are then faced with the task of elucidating the relationships between them, i.e., merging with TimeBank and the Discourse TreeBank which annotate temporal and discourse relations between eventualities. A prerequisite for defining relationships between eventualities is identifying them. Not all verbal and nominal argument structures actually instantiate events. Thus, at one level the event identification task is fairly straightforward, involving the association of eventuality variables or IDs with the PropBank I predicate argument structures, with a certain amount of hand correction. For example, phrases that contain an aspectual verb combined with another verb, such as, *they are continuing the search*, or *they are beginning the search*, should only get one ID with an associated temporal predicate instead of two. In addition, only nominalizations that denote events, such as *The real estate development ran way over budget* should get IDs, rather than all noun predicates.

In spite of the seemingly simple definition of event IDs, their impact is enormous, since they provide the key to merging with TimeBank and the Discourse TreeBank. The event IDs allow us to link together the individual predicate argument structures the way beads are strung on a string. They provide a structure for the semantic representation that is independent of the syntactic structure, giving an additional flexibility to our ability to match sentences that are semantically similar but have quite different syntactic structures.

Once we have satisfactorily combined the existing semantic annotations, the next step is to consider richer layers of annotation. One of the themes of our workshops will be exploring new schemas for annotating additional linguistic phenomena, and incorporating them into our unified framework. The Opinion Annotation described above will be our first candidate for inclusion, and exemplifies the style of richer annotation we will be aiming at (Wiebe et al, 2005). Richer discourse structure annotation that moves us closer to Marcu's Rhetorical structure (Marcu, 2000) will be an important goal (e.g., Calhoun 2005). We will also be interested in extending the annotation to additional syntactic categories such as adjectives, prepositions, quantifiers, and comparatives to provide a more complete representation of the entities and events in each sentence. The workshops will provide us with a forum for discussing advantages and disadvantages of various proposals, and coming to a consensus on prioritization.

### *5.2 Training of Automatic Statistical Taggers*

The intended goal of the annotation is to provide training data for automatic taggers. There is a well-known synergy between the development of the annotation schemas and the training process, since training taggers can also provide a test of the consistency and quality of the annotated data. Once sufficient amounts of data are available, automatic taggers will be trained. Their accuracy should be sufficient for them to then be used as pre-processors. This can improve the consistency of additional data as well as speeding up the annotation process, thus providing more data for training improved taggers. The taggers include semantic role labels, senses, temporal expressions, events, and subordination relationships.

### *5.3 Stabilizing the Methodology*

A key prerequisite for the dissemination of this technology will be the production of a stable, language independent methodology based on guidelines and universally accessible tools. The guidelines need to explicate the details of each individual annotation process as well as interactions between the different types of annotation. For instance, the guidelines for the Proposition Bank outline a process that begins with creating a Frame File for each individual verb in the corpus to be annotated. The Frames Files provide invaluable additional direction for the annotation of the individual verbs over and above the general annotation guidelines. The NomBank annotation in turn begins by referencing the verb Frame Files for associated nominalizations whenever possible. The same approach has been used successfully for Chinese probanking and nombanking. However, it needs to be tested on at least two or three additional disparate language families before it can be considered a stable part of the comprehensive linguistic annotation process.

## **6. Indicators of Success and Major Accomplishments**

We are currently in the first three months of our research.

## **7. Publications/Patents**

We are currently in the first three months of our research.

## **8. Outreach Programs and Supported Students**

By gathering together the experts on various annotation layers to give training seminars to workshop attendees we will also be providing an environment for in-depth discussion of annotation merging. The student exchanges will also be dual purpose, in that they will provide more in-depth training as well as cross-fertilization of ideas that can benefit the host institution as much as the visiting student.

## **9. Future and Strategic Directions**

We are currently in the first three months of our research.

## **References**

Burnard, Lou (2000) User Reference Guide for the British National Corpus. Technical report, Oxford University Computing Services.

Calhoun, S., M. Nissim, M. Steedman and J. Brenier (2005) A Framework for Annotating Information Structure in Discourse. In Proc. Workshop Frontiers in Corpus Annotation II: Pie in the Sky. ACL 2005, Ann Arbor, Mich.

- Carletta, J. and A. Isard (1999), The MATE Annotation Workbench: User Requirements. In Proc. of the ACL Workshop: Towards Standards and Tools for Discourse Tagging. Univ. of Maryland, 11-17.
- Cieri, C. and S. Bird (2001), Annotation Graphs and Servers and Multi-Modal Resources: Infrastructure for Interdisciplinary Education, Research and Development Proc. of the ACL Workshop on Sharing Tools and Resources for Research and Education, 23-30.
- Day, D., L. Ferro, R. Gaizauskas, P. Hanks, M. Lazo, J. Pustejovsky, R. Saurí, A. See, A. Setzer, and B. Sundheim (2003), The TIMEBANK Corpus. *Corpus Linguistics*.
- Fellbaum, C. (1997), *WordNet: An Electronic Lexical Database*, MIT Press..
- Hajivcova, E. and I. Kuvcerov'a (2002). Argument/Valency Structure in PropBank, LCS Database and Prague Dependency Treebank: A Comparative Pilot Study. In the *Proc. of the Third International Conference on Language Resources and Evaluation (LREC 2002)*, 846—851.
- Marcu, Daniel (2000). The Rhetorical Parsing of Unrestricted Texts: A Surface-Based Approach. *Computational Linguistics*, 26 (3), pages 395-448.
- Marcus, M., B. Santorini, and M. Marcinkiewicz (1993), Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*, 19:313—330.
- Meyers, A., R. Reeves, C. Macleod, R. Szekely, V. Zielinska, B. Young, and R. Grishman (2004a), The NomBank Project: An Interim Report, *Proc. of the HLT-EACL Workshop on Frontiers in Corpus Annotation*.
- Miltsakaki, E., R. Prasad, A. Joshi and B. Webber. (2004a), The Penn Discourse Treebank. In *Proc. of the 4th International Conference on Language Resources and Evaluation (LREC 2004)*.
- Miltsakaki, E., R. Prasad, A. Joshi and B. Webber. (2004b), Annotation of Discourse Connectives and Their Arguments, in *Proc. of the HLT-NAACL Workshop on Frontiers in Corpus Annotation*.
- Palmer, M., D. Gildea, P. Kingsbury (2005), The Proposition Bank: A Corpus Annotated with Semantic Roles, *Computational Linguistics Journal*, 31:1.
- Palmer, M., H. Dang, C. Fellbaum, (to appear, 2006) Making Fine-grained and Coarse-grained sense distinctions, both manually and automatically, *Journal of Natural Language Engineering*
- Poesio, M. (2004a), The MATE/GNOME. Scheme for Anaphoric Annotation, Revisited, *Proc. of SIGDIAL*.
- Poesio, M. (2004b), Discourse Annotation and Semantic Annotation in the GNOME Corpus, *Proc. of the ACL Workshop on Discourse Annotation*.
- Poesio, M. and R. Vieira (1998), A corpus-based investigation of definite description use. *Computational Linguistics*, 24(2).
- Poesio, M., F. Bruneseaux, and L. Romary (1999), The MATE meta-scheme for coreference in dialogues in multiple language, *Proc. of the ACL Workshop on Standards for Discourse Tagging*.
- Pollard, C. and I. A. Sag (1994), *Head-driven phrase structure grammar*. University of Chicago Press.
- Pustejovsky, J., B. Ingria, R. Saurí, J. Castaño, J. Littman, R. Gaizauskas, A. Setzer, G. Katz, and I. Mani (2003), The Specification Language TimeML. In I. Mani, J. Pustejovsky, and R. Gaizauskas, editors, *The Language of Time: A Reader*. Oxford University Press.
- Pustejovsky, J., R. Saurí, J. Castaño, D. R. Radev, R. Gaizauskas, A. Setzer, B. Sundheim and G. Katz (2004), Representing Temporal and Event Knowledge for QA Systems. In Mark T. Maybury (ed.), *New Directions in Question Answering*, MIT Press.
- Wiebe, J., Wilson, T., Cardie, C. (2005) Annotating expressions of opinions and emotions. *Language Resources and Evaluation*, volume 39, issue 2-3, pp. 165-210.

# RI: Wide-Area-Network in a Laboratory

Proposal # EIA-0303620; Project Year

California Institute of Technology

Steven H. Low, John Doyle, Harvey B. Newman, Demetri Psaltis

slow@caltech.edu, doyle@cds.caltech.edu, newman@hep.caltech.edu, psaltis@sunoptics.caltech.edu

<http://wil.cs.caltech.edu>

## Abstract

WAN-in-Lab is a hardware testbed for the design, development, testing and evaluation of TCP protocols. It uses real carrier-class networking hardware to avoid the artifacts introduced by network emulation, while being localized to allow detailed measurement of network performance.

## Introduction/Background

It has recently become apparent that very high speed wide area networks (WANs) are pushing the limits of existing protocols, such as the transmission control protocol (TCP). This has led to many proposed replacements for TCP [1], [2], [3], [4], [5], [6], [7], [8], which must be evaluated, optimized and eventually developed into the next generation of TCP. Although initial design and testing can be performed using mathematical modeling and software simulation, there is ultimately a need to implement the selected algorithms in real networks. This project seeks to develop a test facility called WAN-in-Lab, consisting of a complete WAN in a laboratory environment, which is to be available for use by the global networking research community. Developing protocols requires many forms of testing. Existing tools are listed below, in decreasing order of abstraction. Except recirculating loops, all of these play a role in developing congestion control protocols. However, there is a big gap between emulation and production networks, which WAN-in-Lab seeks to fill.

**a) Mathematical modeling** explores fundamental limits of algorithms. It is the only tool that can explore infinite classes of networks rather than specific instances, but analysis requires simplification of actual protocols, ignoring important implementation issues.

**b) Simulation** is the most inexpensive way to test new protocols, and is a useful first step. However, network simulation is hundreds of times slower than real-time, and slower for high-speed networks. Moreover, it does not allow actual prototype deployments to be tested and optimized.

**c) Emulation** of large delays and link impairments using Dummynet also works well for links below 1 Gbit/s. A successful example of this approach is EmuLab [9]. Unfortunately, software emulation introduces artifacts into the traffic, which become increasingly severe at higher link rates.

**d) Recirculating optical loops** can also emulate long physical links with high delay. This approach is suitable for studying physical-layer effects on single packets, but cannot delay a continuous stream of packets, and hence is unsuitable for flow control experiments.

**e) Production networks** are the ultimate testing ground for new protocols, allowing “black-box” evaluations through limited end-to-end measurements. This is suitable for tests which will not disrupt other traffic, but not for testing many failure modes, such as heavily overloaded networks, or equipment failures.

## Goals, Objectives, Targeted Activities

The aim of WAN-in-Lab is to provide a realistic but controlled environment, which enables detailed monitoring of all aspects of protocol operation. This will allow an integrated approach, where theory development, implementation, experiments, and deployment inform and influence each other intimately. WAN-in-Lab is an open resource, intended for use by the entire networking community.

The hardware in WAN-in-Lab is generic and can be used for many aspects of protocol development. However, our goal is to continue to develop a higher layer software infrastructure targeting performance benchmarking of different TCP variants. This will involve developing a measurement capability which can identify the queuing and burstiness of traffic, and correlate these measurements with the internal state of the protocols, such as window sizes and timeout thresholds.

## Infrastructure and Facilities

As this is an infrastructure project, the distinction between infrastructure and outcomes is blurred. This section describes the current status of the WAN-in-Lab hardware.

WAN-in-Lab is centered around an SDH/Sonet optical backbone with four OC48 routers, with wireless and gigabit Ethernet access networks. Delay is provided by 24 spools of 100 km of single-mode fiber.

To increase the delay, data traverses a spool sixteen times. A router's 2.5 Gbit/s OC48 stream is multiplexed onto a 10 Gbit/s OC192 stream, which is amplified and returned. The returning stream is multiplexed onto a second OC192 slot, and the process is repeated for four timeslots and four wavelength division multiplexing (WDM) wavelengths. This yields a round trip time of 28.8ms from a single path of four spools.

Currently, WAN-in-Lab has four Cisco routers, each with a single OC48 line card. This allows two links. The fiber spools allow an aggregate delay of 63ms, adjustable in steps of 1 ms. More complex topologies can be created with short gigabit Ethernet (GbE) links between the routers and the 20 high-speed servers.

Some of these servers are configured as Dummynets, while others are software routers to test active queue management (AQM) protocols.

## Project Description

The WAN-in-Lab project consists of the implementation of the infrastructure described in the previous section, and the use of it for the development, testing and evaluation of TCP-like protocols. The initial phase, assembling the core hardware, is largely complete. The three remaining phases are (i) completing the software infrastructure to allow end-users to conduct experiments with minimal impact on one another and with minimal knowledge of the specific hardware (ii) enhancing the features of the network, by improving the monitoring ability, external connectivity and internal connectivity to expand the range of topologies that can be studied (iii) testing the hardware and software by performing actual protocol research; this will be in the form of a benchmark suite to determine strengths and weaknesses of proposed TCP protocols, with the aim of selecting a replacement for the current TCP NewReno.

## Indicators of Success and Major Accomplishments

The basic WAN-in-Lab hardware is fully installed and operational, apart from three of the 24 optical links. It is now waiting on the software infrastructure to become fully operational. In January 2006, a remote reconfigurability capability was added using a 144×144 port optical switch. This will allow topology changes between experiments or within an experiment. Rerouting is particularly important for delay-based protocols such as FAST. More vitally, this capability also allows the delay between routers to be adjusted from 0ms to 63ms (round trip), in increments of 1ms. Using IP loopbacks, it is possible to achieve delays up to 125ms, corresponding to trans-Atlantic distances. TCP protocols behave very differently depending on the round trip delays, and so it is very important to be able to test and evaluate them in this wide range of conditions.

The software infrastructure is also taking shape. Servers can be manually allocated to projects by means of an interactive web interface, and the interconnectivity can be controlled using an intuitive web interface to a gigabit Ethernet switch. Virtualization of important portions of the filesystem allows the servers to be configured separately for each project, protecting users from inadvertent reconfiguration by other users and allowing servers to be reconfigured easily for different projects.

The software infrastructure now allows the bandwidth, buffer size and delay of each OC48 link to be set without needing to detailed knowledge of router configuration.

## Publications/Patents

Since the purpose of this project is to develop a research infrastructure, there is a lag between the project itself and the resulting publications. To date there has been one publication arising from this work,

Martin Suchara, Ryan Witt, Bartek Wydrowski, “TCP MaxNet – Implementation and Experiments on the WAN in Lab”, in *Proc. IEEE Int Conf On Networks (ICON)* 2006.

and one other conference presentation

George S. Lee, Lachlan L. H. Andrew, David X. Wei, Bartek P. Wydrowski, Cheng Jin, John Doyle, Steven H. Low, Harvey B. Newman, “A WAN-in-Lab for Protocol Development”, TERENA Networking Conference, 2006

In addition, there are several works nearing completion based on results obtained using the WAN-in-Lab, either as main results or as benchmarks for comparison:

David X. Wei and Pei Cao, “NS-2 TCP-Linux: A TCP implementation in NS-2 with Congestion Control Algorithms from Linux Kernels”

Joon-Young Choi, Hyungmo Khoo, David x. Wei, Jin S. Lee and Steven H. Low, “Global Exponential Stability of FAST TCP”

David X. Wei, Pei Cao and Steven H. Low, “Time for a TCP Benchmark Suite?”

Ao Tang, David Wei, Steven H. Low and Mung Chiang, “Bandwidth Allocation for Networks with Heterogeneous Congestion Control Protocols. Submitted to ICNP.

Lachlan L. H. Andrew, Martin Suchara, Ryan Witt and Bartek P. Wydrowski, "MaxNet: From theory to implementation".

as well as many projects in early stages.

## Outreach Programs and Supported Students

Several groups outside Caltech have plans to use the WAN-in-Lab infrastructure. Arrangements are being made with groups of researchers at the University of Pittsburgh and the University of Melbourne, Australia to make use of this facility. This is in addition to the work already carried out by researchers at Stanford University and the Pohang University of Science and Technology, Korea.

The WAN-in-Lab project is providing the basis for George Lee's masters dissertation under Steven Low's supervision. Equipment within WAN-in-Lab is also supporting Jerome White's masters thesis within Caltech.

## Future and Strategic Directions

Both the hardware and software of WAN-in-Lab are still evolving. The management software in particular requires work. As an increasing number of independent projects start using WAN-in-Lab, there is increased need for automated resource allocation, and time-sharing on finer timescales, such as days or hours rather than weeks. This may involve the adoption of the eWAN framework developed at INRIA [10], France and Caltech's MonALISA framework [11]. Going forward, our own research using WAN-in-lab will be into benchmarking the many proposed TCP variants. A general purpose framework will be developed to allow a wide range of measurements to be made. Where possible, this will use user-level code to remain as independent as possible of the kernel-level implementation of the TCP variant. This will also require a standardized set of traffic generation tools, which will need to keep sufficient statistics. Updates to the software infrastructure will be driven primarily by this research.

Earlier plans to support a richer set of topologies through the acquisition of additional line cards have been delayed by policy changes within one of the corporate sponsors of WAN-in-Lab, but that remains a long term goal. In the mean time, creative ways are being devised to emulate more complex virtual topologies using the existing infrastructure.

Another strategic goal is to connect WAN-in-lab to the Ultralight 10 Gbit/s experimental/production network [12]. This will allow the monitoring facilities of WAN-in-Lab to study real-world traffic, allow studies of incremental deployment of protocols, and increase the range of topologies available to WAN-in-Lab.

Arrangements are also being made to attach measurement equipment (DAG cards) to the routers, to allow the detailed measurement of queue sizes and the burstiness of data.

## References

- [1] C. Jin, D. X. Wei, S. H. Low, G. Buhrmaster, J. Bunn, D. H. Choe, R. L. A. Cottrell, J. C. Doyle, W. Feng, O. Martin, H. Newman, F. Paganini, S. Ravot, and S. Singh. FAST TCP: From theory to experiments. *IEEE Network*, 19(1):4-11, 2005.
- [2] C. Jin, D. X. Wei, and S. H. Low. TCP FAST: Motivation, architecture, algorithms, performance. In *Proceedings of IEEE Infocom*, March 2004. <http://netlab.caltech.edu>.
- [3] S. Floyd. HighSpeed TCP for large congestion windows. Internet draft draft-floyd-tcp-highspeed-02.txt, work in progress, <http://www.icir.org/floyd/hstcp.html>, February 2003.

- [4] R. N. Shorten and D. J. Leith. H-TCP: TCP for high-speed and long-distance networks. in *Proc. PFLDnet*, Argonne, 2004. <http://hamilton.ie/net/htcp3.pdf>.
- [5] T. Kelly. Scalable TCP: Improving performance in highspeed wide area networks. *Computer Communication Review*, 32(2), April 2003.  
<http://www-lce.eng.cam.ac.uk/~ctk21/scalable/>.
- [6] L. Xu, K. Harfoush, and I. Rhee. Binary increase congestion control for fast long distance networks. In *Proc. INFOCOM*, March 2004.
- [7] D. Katabi, M. Handley, and C. Rohrs. Congestion control for highbandwidth delay product networks. In *Proc. Sigcomm*, Aug. 2002.
- [8] B. Wydrowski, L. L. H. Andrew, and M. Zukerman, “MaxNet: A congestion control architecture for scalable networks,” *IEEE Commun. Lett.*, 7:511–513, Oct. 2003.
- [9] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb and A. Joglekar, “An Integrated Experimental Environment for Distributed Systems and Networks”, In *Proc. Fifth Symp. Operating Systems Design and Implementation*. Dec. 2002, pp. 255–270.  
<http://www.emulab.net/>.
- [10] <http://www.ens-lyon.fr/LIP/RESO/software/EWAN/>.
- [11] MonALISA web page, <http://monalisa.cern.ch/MONALISA/>.
- [12] Ultralight: An ultrascale information system for data intensive research.  
<http://www.ultralight.org>.

# Shaping the PRISM Center

## CRI: Center for Perceptual Robotics, Intelligent Sensors and Machines at CCNY

CNS-0551598, 1<sup>st</sup> Year

The City College of New York, CUNY

Jizhong Xiao, Zhigang Zhu, Myung Lee, George Wolberg

[jjiao@ccny.cuny.edu](mailto:jjiao@ccny.cuny.edu), [zhu@cs.ccny.cuny.edu](mailto:zhu@cs.ccny.cuny.edu), [lee@ccny.cuny.edu](mailto:lee@ccny.cuny.edu), [wolberg@cs.ccny.cuny.edu](mailto:wolberg@cs.ccny.cuny.edu)

URL: [http://www-cs.engr.ccny.cuny.edu/~zhu/Research/PRISM\\_Center.html](http://www-cs.engr.ccny.cuny.edu/~zhu/Research/PRISM_Center.html)

## Abstract

Supported by the newly awarded NSF CRI grant and significant amount of matching funds from the City College of New York (CCNY) and other external research grants, the PIs and involving faculty members are making collective efforts to establish a Center for Perceptual Robotics, Intelligent Sensors and Machines (PRISM) with the aim to make the CCNY a center of excellence in cutting-edge research, as well as a national urban model for minority education in robotics and related fields. This report highlights the recent progress and achievements in shaping the PRISM center.

## Introduction/Background

The multidisciplinary nature of intelligent robotics research requires collaborations among researchers with expertise in robotics, artificial intelligence, computer vision, wireless communication, advanced sensors and mechatronics. Robotics research and education at CCNY is currently scattered in several departments (EE, CS and ME). The goal of the newly awarded NSF CRI project (Grant No. CNS-0551598) is to establish a *Center for Perceptual Robotics, Intelligent Sensors and Machines* (PRISM) to consolidate the research activities and further enhance the collaboration among faculty members around a common research theme.

CCNY is one of the most diversified campuses in the U.S., with about 1/3 Hispanic and 1/3 Black students. One of the major challenges at CCNY is student retention. The funding for the establishment of the PRISM center contributes to the attraction and retention of minority students in our IT programs in the most profound way. This is achieved through integration of research with coherent educational and outreach activities and through early recruiting and retention via affinity mentoring at all levels of education.

Although the CRI grant was only a fraction of our originally requested budget, the CCNY provided strong institutional support with the total amount of \$212,424 in the forms of greatly reduced indirect costs, faculty release times, and the fund for PRISM lecture series.

## Objectives and Targeted Activities

The objectives of the project are to enhance comprehensive robotics research and education infrastructure and to group the researchers together to carry out nationally competitive research in robotics, computer vision, wireless communication, human-computer interface, and other related fields. The acronym “PRISM” reflects our desire to tackle technical challenges from various facets of real-world problems and to build the center upon three pillars: research, education, and outreach, with the aim to make the CCNY a center of excellence in cutting-edge research as well as a national urban model for minority education in robotics and related fields. Leveraging the previous research expertise of the team members in these related fields, the research theme of the PRISM center we have identified in the next few years is *Robotized Intelligent Sensor NETWORKs (RISE-NET) in 3D space*. Around this theme, we are conducting a broad spectrum of research, divided into four main research thrusts:

- *Distributed Perceptual Robotics* (robotized sensor networks, sensing, control, and coordination of robot teams in 3D space);

- *Multimodal Intelligent Sensors* (multimodal sensors for target recognition, subject tracking, and event understanding);
- *Pervasive Smart Machines* (human-robot interaction, and human-machine visual and speech interfaces in a large sensor network)
- *Reliable Communications* for RISE-NET (wireless sensor/ad hoc networks, standardization).

The education activities at the PRISM center include: introducing new courses with real research projects, reforming senior design courses, launching a student robotics club, and mentoring minority students in research projects. The PRISM center will further increase the cooperation with the NYC-LSAMP (Louis Stokes Alliance for Minority Participation) program, New York Academy of Sciences' Summer Research Training Program (SRTP), and NYC/NJ FIRST to carry out student nurturing activities and provide mentored research experience to minority students at CCNY.

## Infrastructure and Facilities

Due to the budget constraint, this NSF CRI grant has no money to purchase new equipments. However, since the PRISM center is built upon existing research laboratories directed by PRISM team members, the PRISM center has some basic facilities to conduct proposed research activities. We are going to apply NSF MRI grant to further enhance the research infrastructure. Following is the list of laboratories associated with the PRISM center.

- CCNY Robotics Lab (directed by Dr. Xiao, EE Dept.)
- CCNY Visual Computing Lab (directed by Dr. Zhu, CS Dept.)
- CCNY Advanced Wireless Networking Lab (directed by Dr. Lee, EE Dept.)
- CCNY Image Processing, Graphics, and Vision Lab (directed by Dr. Wolberg, CS Dept.)
- CCNY Speech and Natural Language Lab (directed by Dr. Levin, CS Dept.)
- CCNY Graphics Learning and Smart Sensors Lab (directed by Dr. Grossberg, CS Dept.)
- CCNY Multimedia Lab (directed by Dr. Uyar, EE Dept.)
- Center for Advanced Engineering Design and Development (directed by Dr. Sadegh, ME Dept.)

## Project Description

Supported by the PRISM infrastructure, the team members actively conduct research in the four thrusts under the NSF CRI grant. We also collaborated with leading research universities (e.g., UIUC, UMass, Columbia) and attracted many other research grants from the army research office (ARO), army research laboratory (ARL), air force research laboratory (AFRL), NSF MRI program, and industry. Some of the projects are listed below, which are part of the PRISM research thrusts.

### *Cooperative Wall-Climbing Robots (Xiao, Zhu, Sadegh, Grossberg, supported by ARO)*

We are developing a modular, reconfigurable, wall-climbing robotic system and to investigate intelligent control methods and vision algorithms to control and coordinate a team of such robots to perform various defense, security, and inspection missions. First, the modeling of individual wall-climber module, as well as the cooperative relationship of a group of such robots will be developed. Second, intelligent control methods will be designed to effectively coordinate the robot team to achieve a common goal (e.g., self-deployment, dynamic view planning, and object tracking, etc.) in a fault tolerant fashion. Third, based on our previous research in computer vision, new vision algorithms for scene understanding, object tracking and localization will be developed by taking advantage the climbing feature of multiple wall-climbers in 3-D environments.

### *Smart Brain for Miniature Robots (Xiao, Uyar, supported by NSF MRI program)*

The objective of this project is to develop a proof-of-concept, highly-adaptive computation module suitable for ultra-small robots for research in heterogeneous, reconfigurable teams of robotic agents. The embedded high-performance computation module based on System on Programmable Chip (SoPC) technology allows onboard sensor processing, advanced kinematic control, reliable communication, and real-time adaptation of the system to its environment enabling research in a wide variety of tasks in urban search and rescue, civilian and military surveillance, and planetary exploration.

### *Multimodal Surveillance (Zhu, Wolberg, Levin, with UIUC, supported by AFRL)*

In this project, we focus on the long-range (remote) personnel detection, either in open space or concealed in a cave or inside a room, from a distance of 100 meters up to kilometers. Particularly, a laser Doppler vibrometer (LDV) will be applied. Since a single sensor type is unlikely to yield a reliable level of personnel detection accuracy, we integrate the following novel sensors: video cameras, long range thermal IR cameras and laser Doppler vibrometers. Such an integrated system would have day and night operation. This system could also provide the feeds for advanced face and voice recognition systems.

### *Wireless Sensor Networks/Mesh Networks (Lee, Xiao, Umit, supported by Samsung, ARL)*

Dr. Lee leads the research efforts in wireless networking which include: conducting basic research on MAC/Routing/Transport, QoS, and cross-layer optimization; developing sensor network testbed and prototyping; engaging in standardization activities. We have proposed new routing protocols: adaptive robust tree (ART), meshed ART, and TDLS (Tree-guided Distributed Link State), and built testbed to verify the efficiency of the protocols in real environments. We have deployed 50 nodes which covers the 5<sup>th</sup> floor of CCNY engineering building. We made proposals to IEEE802.15.5 technical group, and ZigBee working group. We developed IEEE 802.15.4 NS-2 simulator which is included in the NS distribution. We also investigated the dynamic survivable resource pooling and its applications in distributed robotics.

### *3D and Motion Extraction from Large-Scale Scenes (Zhu, Wolberg, Xiao, with UIUC)*

We propose a dynamic pushbroom stereo mosaic approach for representing and extracting 3D structures and independent moving targets from urban 3D scenes. Our goal is to acquire panoramic mosaic maps with motion tracking information for 3D (moving) targets using a light aerial vehicle equipped with a video camera flying over an unknown area for urban surveillance. We propose a segmentation-based stereo matching approach with natural matching primitives to estimate the 3D structure of the scene, particularly the ground structures (e.g., roads) on which humans or vehicles move, and then to identify moving targets and to measure their 3D structures and movements. The stereo matching method is also being used for surface estimation for wall-climbing robots.

### *3D Virtualized Classrooms (Zhu, Levin, Grossberg, with UMass and CCNY Pathways)*

The project addresses fundamental research problems including automated intelligent capture of materials; intelligent compression; cross-media synchronization for search, retrieval, and media-component enhancement/substitution; archiving and distribution; and 3D, user-controlled, augmented presentation. This is a joint research between CCNY and UMass. We plan to use technology-supported, active, peer and cooperative learning in large, diverse introductory courses with a special emphasis on retention of women and underrepresented minorities at CCNY.

### *Speech and Natural Language Processing (Levin, with Columbia Univ., supported by AFRL)*

This research includes the following projects: (1). *Evaluation of the use of Latent Semantic Analysis (LSA) for unsupervised word sense disambiguation.* The hypothesis is that LSA can be used to compute context vectors for ambiguous words that can be clustered together – with each cluster corresponding to a different sense of the word. (2). *Optimization of dialog manager decisions.* This project is conducted within a newly established industrial collaboration with Telleureka, Inc. The goal of this project is to train a predictor that will be later integrated within the dialog manager. The function of the predictor is to detect in early stages of human-computer dialog, those dialogs that will not have a successful resolution. (3). *Automatic annotation of semantic categories using dialog data.* One of the most labor-intensive and expensive tasks in building spoken language understanding system (SLU) is annotation of utterances by semantic categories. In this project we are evaluating the use of logged dialog data for automatic probabilistic annotation of utterances.

## **Indicators of Success and Major Accomplishments**

Although this NSF CRI grant officially started in April 1, 2006, we have made some major accomplishments, thanks to the support of NSF MII planning project (CNS-0424539), and also due to our continuous efforts in the last several years.

### *Organization and Collaborations*

We are forming the PRISM Center, with Dr. Jizhong Xiao and Dr. Zhigang Zhu as co-directors, and both internal and external experts in our **Steering Committee and Advisory Committee**. We have organized weekly

seminars, where PIs present their research and have brainstorm on research issues of the PRISM center. We continue our **Lecture Series on Robotics, Computer Vision, and Human-Computer Interaction** by inviting experts from other universities to give talks. Collaborations between PRISM center and other institutions (e.g., UMass, UIUC, CUNY Hunter college, Columbia, AFRL) have been enhanced. Multiple joint research proposals are submitted and some of them are awarded.

#### *Research and Development*

PIs have performed joint work even before the establishment of the PRISM center. Supported by the PRISM infrastructure and other research grants, we have developed **3 generations** of wall-climbing robot prototypes which are named as City-Climbers. The video “City-Climber, a new generation of wall-climbing robots” (producers: Matthew Elliot, William Morris, Jizhong Xiao) is one of the three finalists for **the Best Video Award** at the Int. Conf. on Robotics and Automation (ICRA2006), May 15-19, Orlando, USA. We have **deployed 50 wireless sensor nodes** which cover the 5<sup>th</sup> floor of CCNY engineering building enabling us to conduct real experiments on wireless communication and distributed robotics. More than 20 publications and patents have been produced.

#### *Education and Outreach*

- **Minority Student Recruiting:** We have attracted three minority Ph.D. students to conduct research in PRISM Center. In collaboration with NYC-LSAMP program, we have identified several minority undergraduate students to conduct mentored research in the CCNY Robotics Lab and Visual Computing Lab under the direction of Dr. Xiao and Dr. Zhu during summer and in semesters.
- **Joint Capstone Projects for Seniors:** We have started offering joint senior design projects among multiple departments. For example, Dr. Xiao (EE) and Dr. Sadegh (ME) grouped EE and ME students in a team to design wall-climbing robots; Dr. Zhu (CS) and Prof. Charles B. Watkins (ME) are offering a joint capstone section for 4 ME seniors to design novel sensor devices for multimodal target detection.

## Publications/Patents

*The partial list of recent papers (19) and patents (3) related to the work in the PRISM Center is as follows:*

- [1]. Flavio Cabrera-Mora, Jizhong Xiao and Yi Sun, “Effects of Communication on Mobile Sensor Networks”, 2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Beijing, China, Oct. 2006, to appear.
- [2]. Yi Feng, Zhigang Zhu, Jizhong Xiao, “Heterogeneous Multi-robot Localization in Unknown 3D Space”, 2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Beijing, China, Oct. 2006, to appear.
- [3]. Narashiman Chakravarthy, Jizhong Xiao, “FPGA-based Control System for Miniature Wall-climbing Robots”, 2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Beijing, China, Oct. 2006, to appear.
- [4]. Matthew Elliot, William Morris, Jizhong Xiao, “City-Climber, a new generation of wall-climbing robots”, in Video Proceedings of 2006 IEEE International Conference on Robotics and Automation, May 15-19, 2006, Orlando, USA, **(one of the three finalist for the BEST VIDEO AWARD)**
- [5]. Jizhong Xiao, William Morris, Narashiman Chakravarthy, Angel Calle, “City-climber: a New Generation of Mobile Robot with Wall-climbing Capability”, Proc. of SPIE Vol. 6230, 62301B, 2006 SPIE Defense & Security Symposium, April 17-21, 2006, Orlando, Florida, USA.
- [6]. Abdelal, Hkelek, M. Uyar, and J. Zou; “Dynamic survivable resource pooling in FPGA-based distributed robotics system”, IEEE International Conference on Networking, Sensing and Control (ICNSC), pp. 1016-1021, Ft. Lauderdale, Florida, April 23-25, 2006.
- [7]. J. Xiao, Z. Zhu, “City Climbers at Work”, submitted to *IEEE Computer*, special issue on advances in. unmanned ground vehicles, December 2006.
- [8]. Z. Zhu, G. Wolberg, J. R. Layne, “Dynamic pushbroom stereo vision for surveillance and inspection”, a book chapter in *3D imaging for Safety and Security*, eds. A. Koschan, M. Pollefeys, and M. Abidi, Kluwer/Springer, 2006, in press. **(joint work of two labs at CCNY and AFRL)**
- [9]. Z. Zhu, A. Hanson, Mosaic-based 3D scene representation and rendering, Signal Processing: Image Communication, Special Issue on Interactive Representation of Still and Dynamic Scenes, Elsevier, 2006 (accepted). **(joint work of CCNY and UMass)**
- [10]. Z. Zhu, G. Xu, E. Riseman and A. Hanson, Fast Construction of Dynamic and Multi-Resolution 360-degree Panoramas from Video Sequences. Image & Vision Computing Journal, Volume 24, Issue 1, 1 January 2006, pp 13-26. **(joint work of CCNY and UMass)**
- [11]. Esther Levin, Mehrbod Sharifi and Jerry Ball "Word Sense Discrimination using LSA" submitted to the ISCLP 2006 conference. **(joint work of CCNY and AFRL)**
- [12]. Esther Levin and Rebecca Passonneau, ‘A WOz Variant with Contrastive Conditions’, in Proc Dialog-on-Dialog Workshop, INTERSPEECH 2006 — ISCLP, September 17-21, 2005, Pittsburg, PA **(joint work of CCNY and Columbia Univ.)**

- [13]. H. Tang, Z. Zhu, G. Wolberg and J. R. Layne, Dynamic 3D Urban Scene Modeling Using Multiple Pushbroom Mosaics, the *Third International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT 2006)*, University of North Carolina, Chapel Hill, USA, June 14-16, 2006, accepted.
- [14]. W. Li, H. Tang and Z. Zhu, Vision-Based Projection-Handwriting Integration in Classroom, *IEEE International Workshop on Projector-Camera Systems (PROCAMS'06)*, New York City, June 17, 2006 (in conjunction with IEEE CVPR 2006), accepted.
- [15]. Z. Zhu, H. Tang, "Content-Based Dynamic 3D Mosaics", *IEEE Workshop on Three-Dimensional Cinematography (3DCINE'06)*, June 22, New York City (in conjunction with CVPR), accepted
- [16]. W. Li, M. Liu, Z. Zhu and T. Huang, LDV Remote Voice Acquisition and Enhancement, *ICPR'06*, Hong Kong, China, August 2006, accepted (**joint work of CCNY and UIUC**)
- [17]. W. Li, H. Tang, C. McKittrick and Z. Zhu, "Classroom Multimedia Integration", accepted to *IEEE International Conference on Multimedia & Expo (ICME)*, Toronto, Canada, July 9-12 2006, poster presentation, accepted.
- [18]. W. Li, Z. Zhu and G. Wolberg, Remote Voice Acquisition in Multimodal Surveillance, accepted to *IEEE International Conference on Multimedia & Expo (ICME)*, Toronto, Canada, July 9-12 2006, oral presentation, acceptance rate 22%.
- [19]. Lingyun Liu, Ioannis Stamos, Gene Yu, George Wolberg, and Siavash Zokai, Multiview Geometry for Texture Mapping 2D Images Onto 3D Range Data, *CVPR 2006*, New York, NY: June 17-22, 2006 (**joint work of CCNY and CUNY Hunter College**)
- [20]. George Wolberg and Ali Sadegh, Automatic Page Turner with Turnstile Element, U.S. Patent No. 7,019,203. Issue Date: March 28, 2006.
- [21]. Ali Sadegh and George Wolberg, Automatic Page Turner with Belt Drive Element, U.S. Patent No. 6,935,058. Issue Date: August 30, 2005.
- [22]. Jizhong Xiao and Ali Sadegh, "Modular Wall Climbing Robots with Transition Capability", approved for disclosure by CUNY Intellectual Property Committee, patent filed in Feb. 2006.

## Outreach Programs and Supported Students

### *Outreach programs:*

- **CCNY Robotics Club:** The robotics club attracted many talented students doing projects at CCNY Robotics lab under the direction of Dr. Xiao. The club serves as the interface and major agent to carry out outreach activities, such as lab tours and project demonstrations for high-school students, mentor FIRST robotics competition teams, etc.
- **FIRST Robotics Competition:** Dr. Xiao and CCNY Robotics students have been mentoring high-school teams participating in the FIRST robotics competition since 2003 by organizing C programming workshops for the high-school students. Dr. Xiao served as the judge of the 2006 NYC/NJ FIRST competition .
- **Summer Intern Program for Undergrad and High School Students:** Dr. Xiao and Dr. Zhu are mentoring four undergrad students, and two high-school interns participating New York Academy of Sciences' SRTP program, doing research in robotics, computer vision and multimodal sensor fusion in the summer of 2006, one of them is minority student.

### *Supported Students*

The majority of the budget is used to support graduate students. Leveraging the NYC-LSAMP program, other research grants and scholarships, we maximized the possible number of students supported partially by the CRI grant.

- **PhD Students:** Angel Calle (EE, minority, by Xiao), Jorge Peche (EE, minority, by Xiao), Flavio Cabrera-mora (EE, co-mentored by Xiao and Lee), Edgardo Molina (CS, minority, co-mentored by Zhu and Wolberg), Yi Feng (CS, co-mentored by Zhu and Xiao), Weihong Li (CS, by Zhu), Hao Tang (CS, by Zhu), Zheng Chen (by Levin), Samratt Batth (EE, by Uyar)
- **Undergraduate Students:** William Morris (EE, by Xiao), Chukwuchem Orakwue (EE, by Xiao), Wai Khoo (CS, by Zhu), Nwanze Ononye (CS, by Zhu)
- **High School Students:** Jonathan Speiser (EE, by Xiao), Michael Chen (CS, by Zhu)

## Future and Strategic Directions

We believe we are in the right track to accomplish our goal -- making the CCNY a center of excellence in cutting-edge research as well as a national urban model for minority education in robotics, computer vision, wireless communication, HCI and other related fields. The PRISM team members will enhance the collaboration and jointly carry out the research thrusts within the RISE-Net theme of the PRISM Center.

# **CRI: Collaborative Project: Repository for Model Driven Development (ReMoDD)**

## **Planning Grant**

**CNS-0551622; Year 1 (Beginning 5/15/06)**

PIs: *Robert France*<sup>1</sup>, *James M. Bieman*<sup>1</sup>, *Betty H.C. Cheng*<sup>2</sup>

<sup>1</sup>Department of Computer Science  
Colorado State University  
{france,bieman}@cs.colostate.edu  
<http://www.cs.colostate.edu/~france>,  
<http://www.cs.colostate.edu/~bieman>

<sup>2</sup>Department of Computer Science and Engineering  
Michigan State University  
East Lansing, Michigan  
[chengb@cse.msu.edu](mailto:chengb@cse.msu.edu)  
<http://www.cse.msu.edu/~chengb>  
June 2006

## **Abstract**

This project will develop and deploy a repository-based community resource that supports research and education on model driven development (MDD). MDD research is developing techniques, methods, processes and tools that raise the level of abstraction at which large, complex software systems are conceived, analyzed, implemented and evolved. The Repository for MDD (ReMoDD) will contain artifacts that support MDD research and education. The ReMoDD platform will also provide interfaces and interchange mechanisms that will enable a variety of tools to retrieve artifacts from the repository and submit candidate artifacts to the repository. ReMoDD artifacts will include documented MDD case studies, examples of models reflecting good and bad modeling practices, reference models (including metamodels) that can be used as the basis for comparing and evaluating MDD techniques, generic models and transformation reflecting reusable modeling experience, descriptions of modeling techniques, practices and experiences, and modeling exercises and problems that can be used to develop classroom assignments and projects. The development of ReMoDD will be a collaborative effort involving teams from Colorado State University and Michigan State University. The team will also interact with other members of the MDD community, through a series of Repository Development Workshops, to collect and evaluate candidate ReMoDD artifacts. An advisory board comprising MDD researchers and industry representatives will provide the guidance needed to ensure that ReMoDD becomes a sustainable resource that significantly enhances MDD research productivity and improves the quality and relevance of educational materials.

## **1 Introduction**

The proposed project will develop and deploy a repository-based community resource that supports research and education on model driven development (MDD). MDD research targets the complex problem of developing software systems that play critical roles in organizations and society. Researchers in the MDD community are developing techniques, methods, processes and tools that allow developers to raise the level of abstraction at which they conceive, analyze, implement and evolve complex software systems. The Repository for Model Driven

Development (ReMoDD) will contain artifacts whose use can significantly improve MDD research productivity, improve industrial MDD productivity, and enhance the learning experience of MDD students. Artifacts will include detailed MDD case studies, examples of models reflecting good and poor modeling practices, semantic models for UML diagrams, reference models (including metamodels) that can be used as points against which MDD techniques are compared and evaluated, model and specification patterns, generic models and model transformations reflecting reusable modeling experience, descriptions of modeling practices and experience, and modeling exercises and problems that can be used to develop classroom assignments and projects. ReMoDD will publish an API (application program interface) that will allow a variety of tools to retrieve artifacts from the repository directly. For those who wish to use ReMoDD in a stand alone mode, we will develop web-based user interface software to make the artifacts easily and intuitively accessible, both in terms of performance and content. Specifically, we will develop software to present different views, browsing, and query interfaces to the user depending on their needs and objectives. For example, one type of user may be interested in finding all artifacts related to a specific domain (e.g., telecommunication systems). Another type of user may want to explore the available design patterns with sample implementations. And yet another user may want a general introduction to MDD and want to take a virtual tour through the different types of artifacts in the repository. We will need to investigate different types of storage media and structural organization of artifacts to best support these types of views and query needs.

The development of ReMoDD will be a collaborative effort involving researchers from Colorado State University (CSU) and Michigan State University (MSU). The team will also work with members of the MDD community to collect and evaluate candidate ReMoDD artifacts, as well as publicize the repository artifacts for dissemination purposes. An advisory board comprising MDD researchers and practitioners will provide the guidance needed to ensure that ReMoDD becomes a sustainable resource that significantly improves MDD research productivity and the quality and relevance of educational material.

## 2 Goals, Objectives, Targeted Activities

The objective of the project is to develop a community resource that provides a single point of access to shared artifacts reflecting high-quality MDD experience and knowledge from industry and academia. The aim is to facilitate sharing of relevant knowledge and experience that improve MDD research productivity and education.

We aim to collect and make available MDD artifacts from industry, academia, and other public domain sources (e.g., artifacts produced by open-source projects). Items in the repository will provide data for our research and will be a resource for the entire MDD research and education community. Initially, ReMoDD will support research in the following areas:

- Research on modeling languages and modeling approaches (e.g., research on aspect-oriented modeling, model semantics).
- Research on model transformations.
- Research on model analysis techniques.
- Research on evaluating the quality of modeling artifacts.
- Empirical studies of modeling phenomena.
- Research on reusable forms of modeling experience (e.g., work on developing and using domain-specific modeling languages, domain-specific modeling frameworks, and patterns).

In the long-term we envisage that the repository will consist of related living archives of software engineering artifacts including the MDD archive (ReMoDD) that will be developed in this research. For this reason, the repository infrastructure will not be restricted to storing and manipulating only MDD artifacts.

**Ensuring Relevance to the MDD Community.** We will use two mechanisms to help ensure that the community resource we develop meets the needs of the MDD community: A project Advisory Board and a series of Repository Development Workshops (RDWs).

An Advisory Board comprising national and international MDD researchers and practitioners will provide the oversight needed to ensure that the repository provides the MDD artifacts that are highly relevant to the MDD community. Board members will provide feedback on the design of the repository, evaluate candidate artifacts, contribute artifacts, and publicize the ReMoDD repository. The following is a list of persons who have agreed to serve on the Advisory Board:

Academia Members	
Joanne Atlee	University of Waterloo, Canada
Don Batory	Univ. of Texas – Austin, USA
Jean Bezivin	University of Nantes, France
Lionel Briand	Carleton University, Canada
Doris Carver	Louisiana State University, USA
David Garlan	Carnegie Mellon University, USA
Jeff Gray	University of Alabama at Birmingham, USA
Mark Harman	Kings College, UK
Jean-Marc Jezequel	IRISA/INRIA, France
Kevin Lano	Kings College, UK
Robyn Lutz	Iowa State University and NASA JPL, USA
Atif Memon	University of Maryland, USA
Spencer Rugaber	Georgia Tech, USA
Perdita Stevens	University of Edinburgh, UK

Industrial Members	
Michael Barnett	Microsoft Research, USA
Brian Berenbach	Siemens Corporate Research, Worldwide
Roger Burkhart	Deere & Company, USA
Alexander Egyed	Tecknowledge, USA
Luis Pereira	Eaton Innovation Center, Worldwide
Bran Selic	IBM, Canada
Frank Weil	Motorola, Worldwide

Research Development Workshops (RDWs) will give members of the general MDD community opportunities to interact with the project team and thus influence the development of ReMoDD and its mission. The workshops will be held biannually at two major conferences: ICSE (International Conference on Software Engineering) and MoDELS (Model Engineering Languages and Systems). In the early phases of the project, the workshops will focus on eliciting requirements from the community (new and seasoned developers and researchers) and on reviewing the ReMoDD design. Later workshops will focus on (1) developing and discussing artifacts that will be used to seed the repository, (2) demonstrating how researchers and educators can interact with ReMoDD and (3) soliciting feedback that will be used to assess the effectiveness of ReMoDD and to improve the repository.

**Impact of ReMoDD on Research and Education.** ReMoDD will (1) provide research projects with artifacts such as models, model transformations, and code on which research products can be applied, (2) facilitate comparative analyses of experience related to MDD, (3) provide raw data on MDD artifacts, technologies, and practices to research programs, (4) support efforts related to collecting empirical data about modeling techniques, technologies and notations, and about implemented systems, (5) provide educators with materials that can be used in software engineering courses that cover MDD, and (6) be used to communicate MDD successes and failures to the software development community.

**Impact on Standards Bodies.** In addition, we envision that ReMoDD will also play a role with the standards bodies. The experience and knowledge captured by artifacts in the repository can be used by developers of MDD standards to shape standards and to illustrate application of the standards. The repository can also help promote use of standards through sharing of artifacts that conform to standards, and can help with the evolution of standards by providing information (e.g., experience reports, quality evaluations) that can be used to determine the effectiveness of the standards.

### 3 Project Description and Infrastructure

In this section we present our vision of the ReMoDD infrastructure. We anticipate that the initial view presented in this section will be modified and elaborated during the project based on results of a requirements analysis that will be carried out in the early phases of the project and on feedback we gather from the MDD community.

#### 3.1 Core ReMoDD Content-Related Concepts

To support the short-term MDD-specific goals and the long-term software engineering repository goals we have based ReMoDD on the following content-related concepts: Artifact, artifact relationship, and artifact cluster. Basing ReMoDD on these generic concepts makes it possible to use the repository to store a variety of artifacts.

An *artifact* is an information item that can be retrieved from the repository. It can be simple or complex. A simple artifact is the smallest unit of information that can be accessed within the repository. It is a set of tightly-

coupled elements that is stored and retrieved as a non-decomposable unit in the repository. Examples of simple artifacts are UML class descriptions, UML relationships, Java programs, metamodels, test cases, and method descriptions. Each artifact has a type that contains metadata about the artifact and that specifies the kinds of manipulations that can be carried out on the artifact. The kinds of manipulations supported by an artifact can be described in terms of an interface that specifies allowable operations in terms of their signatures and constraints on their behavior. Users of the repository can use the metadata in artifact types to determine the quality of the artifacts and the tools needed to manipulate the artifacts once they are retrieved from the repository. For example, the type of a source code artifact can include metadata that gives the programming language used to express the artifact, the version and author of the artifact, and provides information on the quality of the code (e.g., complexity metrics). It can also contain specifications of interfaces for analyzing, compiling, testing, and executing the code. An artifact type can also specify data integrity and access control rules that are applicable to all artifacts of the type. These rules can restrict the form of the artifacts and how they are accessed within the repository.

### 3.2 The ReMoDD Development Plan

The ReMoDD project activities are structured into the following phases:

- **Project Startup:** Activities in this phase will be primarily concerned with soliciting and analyzing detailed requirements from the MDD community and with developing a repository architecture that provides a balanced solution.
- **Repository Design and Implementation:** The primary activities in this phase concern designing, implementing and testing the repository.
- **Repository Seeding:** The primary activities in this phase concern collecting, packaging, and storing MDD research and education artifacts that will be shared via ReMoDD.
- **Repository Deployment:** The activities in this phase are primarily concerned with making the repository available to the MDD community and training of potential users.
- **Repository Evaluation:** The activities in this phase are primarily concerned with evaluating the effectiveness of the ReMoDD infrastructure and the use of its seed contents.

The Planning Grant will largely focus on the Project Startup and Repository Seeding efforts, both of which will be described in the Future Work section.

## 4 Evaluation

For the planning grant, the following evaluation activities are planned.

- Evaluation of the software interface for the repository will be performed. Prototype interfaces for the key elements of ReMoDD will be developed and made available to the Advisory Board members as well as our collaborators. The objective of this evaluation step is to determine what types of interactions will be most useful for ReMoDD users. An iterative process will be used to gather feedback and refine the interfaces appropriately. RDWs at ICSE and MoDELS conferences will be used to gather more concentrated feedback regarding the interfaces.
- Evaluation of the artifacts in the repository. As part of the seeding efforts, we will identify representative artifacts to be placed in ReMoDD. As part of the initial requirements phase, we are soliciting input from stakeholders as to what types of artifacts are the most in demand for researchers and educators. Based on this feedback, we will collect appropriate artifacts from researchers, educators, and industrial collaborators. Once placed in a prototype repository with our prototype interfaces, we can then evaluate both the interfaces and the artifacts and determine what types of metadata are needed to make the artifacts the most useful.

## 5 Outreach

The inherent nature of the ReMoDD project is outreach to the research and educational communities. In addition to the Advisory Board members, we have a list of twenty collaborators from industry and universities who are interested in using and contributing to ReMoDD. To increase the scope of outreach, we are also working with several minority institutions to engage minority faculty and students in MDD research and education. The full

proposal describes activities to leverage the contacts by the 3 PIs in HBCUs and other minority institutions to support underrepresented minority undergraduate students in MDD-related research projects. Special effort will be made to involve educational institutions with predominantly minority students in the development and use of the repository. The institutions with software engineering research and teaching faculty will be contacted directly and invited to participate in the Repository Development Workshops (RDWs). Finally, we have identified several global software companies who are enthusiastically interested in participating and contributing to ReMoDD. In many cases, the PIs have long-standing relationships with these companies, where industrial-strength data and projects have been shared with their respective universities. As such, the plan is to leverage these partners to gain additional industrial partners.

## 6 Future Activities

The planning grant is just now beginning. As such, our efforts will focus on the requirements elicitation for the project and the preliminary seeding of the repository.

**Project Startup.** The following are the major activities in this phase:

- Establish the project Advisory Board.
- Elicit and analyze requirements for the repository.

At the start of the project we will hold a meeting with members of the Advisory Board to discuss initial project plans and to establish protocols that specify how the PIs will interact with the board. It is expected that the PIs will meet with the board members at least twice a year to discuss progress. These meetings will be held at the ICSE (May annually) and MoDELS (October annually) meetings. An online mailing list will also be set up to facilitate communication with board members.

Elicitation and analysis of repository requirements is the major activity in this phase. At the start of the project we will conduct a survey that will help us determine the types of artifacts that MDD researchers and educators would like to access in an open repository. We will use the mailing lists of the MoDELS conference, the OMG and other MDD related groups to distribute the survey, and provide online survey instruments.

To facilitate the elicitation activity the PIs will organize a Repository Development Workshop (RDW) at the MoDELS conference, which is due to take place shortly after the start of the project. The objective of the RDW would be to elicit requirements and discuss requirements with members of the MDD community. Members of the MDD community will present and discuss their requirements for the repository. The PIs will meet after these meetings to analyze the feedback and plan future requirements and early design activities.

**Repository Seeding.** The following are the major activities in this phase:

- Seed the repository from the MDD models used to develop ReMoDD.
- Seed the repository with artifacts from PIs projects.
- Seed the repository with artifacts from Advisory Board members.

This phase is intended to provide an initial population of the repository to assess how well the clusters we initially identified are suited to shareable artifacts. During this phase, we will also determine the types of manipulations that should be allowed on artifacts. This will help us determine the types of tools that should be either directly brought into the system or made pluggable into the client version of the system. For example, when browsing a collection of design patterns, should a user be able to click on a pattern, bring up the list of fields, and then pull up a UML diagram editor to begin editing one of the templates in the pattern?

The initial seeding of the repository will also help us to assess whether our criteria for simple and complex artifacts need to be modified, or whether we need another category of artifacts. Given that several of our advisory board members are from industry or work with industry, we will be able to explore scalability issues in terms of volume and complexity of artifacts. At the conclusion of the initial seeding phase, we will have examples of all the key types of artifacts that we plan to support (e.g., requirements, design, implementation, testing models, code for models, transformations, metamodels, and patterns). During the remainder of the project, we will continue to solicit additional artifacts from the community to submit to the repository. Our RDW will provide one means for soliciting artifacts. We will also send out announcements to the community soliciting additional artifacts. A software module will be developed to provide an easy means for users to submit candidate artifacts. All artifacts will be evaluated for their integrity by members of the ReMoDD team before being added to the repository.

# **CRI: Collaborative Research: Building the Next-Generation Global Routing Monitoring System**

NSF 0551725;

Year 0 (starts August 2006)

Colorado State University(lead), University of Memphis, University of Oregon, and UCLA

Dan Massey, Lan Wang, Dave Meyers, and Lixia Zhang

[massey@cs.colostate.edu](mailto:massey@cs.colostate.edu), [lanwang@memphis.edu](mailto:lanwang@memphis.edu), [dmm@1-4-5.net](mailto:dmm@1-4-5.net), and [lixia@cs.ucla.edu](mailto:lixia@cs.ucla.edu)

<http://www.routeviews.org>; <http://netsec.cs.colostate.edu/RV>

## 1. Abstract

This project will develop the next generation Internet route monitoring system. The resulting community resource will provide essential Internet routing data to both the research and education community and network operations community. This work builds upon the existing RouteViews data collection system that was launched in 1998. Over the last few years, the RouteViews archive has become a major data source for network research community and numerous recent network routing research projects have benefited from it. However, RouteViews was initially put together in an ad hoc way using available software packages, and was not designed for research use. As such it has various defects and simply cannot continue to meet the anticipated research demands or be easily adapted for educational use. We plan to replace the current router software package by a dedicated lightweight and extensible data collector, rebuild the data archive with a new standard format that can meet today's and future need, enable efficient download, and provide near real-time distribution of the global routing information.

## 2. Introduction/Background

Understanding global routing is critically important for current and future network research. The existing Internet relies on the BGP routing protocol, but limitations in the existing system are also well known. Security remains an open challenge and is the subject of active research, routing convergence problems have been identified and solutions have been proposed, the research community has also shown growing interest and understanding on impact of routing policies, and efforts on next generation designs have been motivated by problems experienced in the current system and are often evaluated using data drawn from the current Internet. Although it is not clear how a successful design of next generation global routing protocol would look like at this time, one thing is clear: the data that can help us understand the current global routing system is essential to a wide range of network routing research efforts.

However, access to the operational Internet can be a big barrier to research community. Commercial Internet service providers can be reluctant to widely share their operational data. While these barriers can be overcome by determined efforts at gathering data from specific ISPs, a more fundamental challenge is obtaining global scale data. To truly understand and properly analyze global routing data, one needs data from a wide range of sites with different geographical locations and different types (tiers) of ISPs. Fortunately global routing

monitoring projects such as Oregon RouteViews have been providing such essential data to both operations and research community. RouteViews collects BGP routing data from numerous routers around the world and makes this data freely accessible.

Results on route damping, route convergence, routing policies, power-law topologies, new routing protocols, routing security, and design of new routing protocols, and so forth have all benefited from RouteViews data. RouteViews data was cited in papers from SIGCOMM 2005, 2004, 2003, and 2002; in papers from INFOCOM 2005, 2004, 2003, and 2002; and in SIGMETRICS papers from 2004, 2003, and 2002. CiteSeer reports 39 publications that directly cite RouteViews in their reference list and Google Scholar returns 288 documents listing RouteViews as references. In the operational community, RouteViews also serves as a valuable data source and there have been new efforts that utilize the log data for fault detection such as identifying prefix hijacking events. RouteViews also presents a great opportunity for education at both the graduate and undergraduate level. Access to and experience with actual data, including all of its noise and unexplained behaviors, can prove to be invaluable.

But the wide spread use of the data has also revealed substantial problems and limitations. For example, there are concerns over *data completeness*; the BGP peering sessions frequently reset, resulting in lost data and measurement artifacts. *Data integrity* is also a problem; the logs fail to identify, have been shown to miss important dynamics, and there is no security for the collected data. *Data accessibility* is an increasing concern as the volume of data continues to grow dramatically; access is extremely rudimentary and with limited data organization, there is no easy way to access the specific data without downloading large volumes of irrelevant data. In addition, the system is facing problems with *scalability and extensibility*; there is limited support for new protocol extensions and IPv6, challenges in scaling the data as more collection points are added, and an inability to add needed services such as more rapid data access that is needed by researchers in security as well as operators concerned with security. These limitations create obstacles in using the data for research and education, thus blocking the way to meet new operational as well as research. A new development to help RouteViews overcome these limitations is needed to prevent this community resource from becoming outdated and obsolete.

### 3. Goals, Objectives, Targeted Activities

Given the system has been around for many years, the above defects are inevitable, so is a necessary re-alignment of the collector system in order to allow the service to continue and to provide new services as the demands arise. We believe that time is overdue to re-engineer RouteViews, a critically important global routing collection system, to produce high quality of data and meet current and future demands. We formulated these demands into the following list of goals of the new RouteViews.

#### *Data Completeness:*

- The sessions should not fail due to transient problems (e.g. transient congestion).
- The collection system should be able to quickly re-establish the sessions when they fail.

***Data Integrity:***

- Holes in the update files are clearly marked.
- Routing tables can be re-constructed correctly based on update streams.
- Timestamps should be accurate and have a high resolution.
- Data can be collected with the same level of security that is available between production Internet routers.
- Information regarding data authenticity (for example whether the session is protected by MD5) is provided in the logs.

***High Accessibility:***

- Researchers, network operators, and students should be able to access the data they need without having to intensively hunt for data or download large volumes of irrelevant data.
- Researchers, network operators, and students should have rapid access to elements of the routing data.

***Scalability and Extensibility:***

- The data collection system should be scalable to a large number of peers.
- The log format should support new route attributes and new routing features.
- The system should be flexible in providing new services.

To achieve the above goals, we have decomposed the engineering challenges into the following three major tasks. First, we will develop modern data collection software to gather the data from production routers. Instead of providing the routing functionality as the current routing software suites, the data collection software focuses on maintaining peering sessions and collecting BGP updates. Second, we will design and develop a new data storage system and add new meta-data to facilitate efficient data download. Third, we will expand a pilot system that enables rapid access to essential data. Each of the three components of the system has a specific functionality. The first component supplies the data collected from production routers to the other components, while the other two components mainly interact with users – the data storage system provides archived data and new access system serves newly collected data to the users. Collectively, they form a reliable and scalable global routing monitoring system.

#### 4. Infrastructure and Facilities

Our project builds upon the existing RouteViews infrastructure. Currently RouteViews collects routing data using open source routing suites such as MRTD, Zebra or Quagga. BGP data collection efforts at RIPE RIS, commercial ISPs and research institutions also use the same routing suites. The data collectors act as BGP routers and peer with actual BGP routers owned by various ISPs. Although most collectors are physically located at Internet Exchanges Points such as the Oregon Internet Exchange, Palo Alto Internet Exchange, Equinix, Lynx, and WIDE to set up direct BGP peering sessions with operational routers, they also use multi-hop BGP peering sessions to collect data from remote ISPs. After establishing a BGP session with each of its peers, a collector receives BGP update messages from the peers and periodically dumps the received BGP updates and their routing table content to a data storage device. Users then fetch data files from RouteViews data store.

The project team will revise and improve nearly all aspects of the system. The team includes the existing RouteViews project at University of Oregon and is complemented by programmers at Colorado State University, UCLA, University of Memphis who will be developing necessary software and data organization systems in concert with the RouteViews team. The combined facilities at the four sites provide the resources for software development, testing, data collection, and data dissemination.

## 5. Project Description

Based on the experience from 7 years of RouteViews operations, problems discovered from our own use of RouteViews archive, and user feedback from broader community, we have developed a basic design for the next generation of RouteViews BGP monitoring system. The 3 major pieces in our design include a new data collection software package, a new data archive structure, and real time data access. The new collection software package can record all the necessary information that can be collected for understanding the operations of routing system. It also provides enhanced data integrity and system extensibility to support new functions. The new archive structure will include metadata to facilitate efficient access by different users. The rapid access will provide summaries of the global routing operations to enable operators and researchers to identify potential faults and take countermeasures.

Our first step is to design and implement an open source BGP data collector. In some ways, this problem is easier than building a new routing protocol suite. There is no need to implement code for managing forwarding tables, supporting other protocols, announcing routes, and so forth. The resulting code base should be dramatically simpler than current routing protocol suites. This simplicity allows the data collector to focus on features that are important to data collection such as ensuring log message integrity and accurate time stamps, supporting IPv6 and BGP enhancements such as graceful restart, and providing easy extensibility for new BGP features. In addition, there are new features that a data collector should support such as automatically removing updates due to a session reset between the data collector and its peers, improved logging of data collector downtimes and problems, and a new log format that better meets the needs of both researchers and BGP operators. All these features are (at best) secondary issues and low priorities for routing suites, but they are primary jobs and high priorities for data collectors.

Once the data collector has established stable peering sessions, the problem becomes how to structure and organize the resulting data. A typical data collector will peer with several BGP routers. For example, the RouteViews2 data collector at the Oregon Internet Exchange currently peers with 39 BGP routers and each peer sends a high volume of updates. We randomly selected the month of September 2004 to provide some sense of the scale. During this month, the collector received 226.9 million updates. The data collector also maintains routing tables for each of the peers and a typical routing table for each peer has over 180,000 routes. When measured over months and years, the volume of data can become overwhelming. It is essential that log messages are properly structured and the vast volumes of data are organized in a usable fashion. Even measurements over a small time scale can become overwhelming if the data is not well structured and well organized. During the same

randomly chosen one month period, the mean number of updates received during a 30 second interval was 2,626 and, in the most active 30 second interval, the collector received 186,031 updates.

Currently, RouteViews uses the MRT format and logs all updates into a single file. Largely due to limited flexibility of the data collector, new update files are written roughly every 15 minutes. The log format is out of date and does not include some essential data such as session down periods, table transfers and so forth. Routing tables from all peers are also written into a single file every 2 hours. After being uncompressed and converted to ASCII form, one routing table file has a file size of roughly 1 gigabyte. This makes the data inconvenient to use, especially by non-experts. In our experience as researchers, the existing structure requires us to download huge volumes of data to answer even simple question and only a very small fraction of this data is relevant to the research question.

Overall, the current data structure and organization leaves a great deal to be desired, but improvements in the current system are limited due to a lack of flexibility in the base data collector. By replacing the current routing software packages with a dedicated data collector, we can also address the problems with log structure and organization. Once we have developed the new structure and organization, we plan to develop tools that convert the very large volume of existing data into this new format.

#### 6. Indicators of Success and Major Accomplishments

No major accomplishments yet. The project is scheduled to begin in August of 2006.

#### 7. Publications/Patents

No publications yet. The project is scheduled to begin in August of 2006.

#### 8. Outreach Programs and Supported Students

RouteViews is one of the best known and most widely used sources of global routing data and results of this project will directly effect the continued availability of this essential resource. The project will include graduate student researchers at Colorado State University, University of Memphis, and UCLA who will work with the RouteViews team at University of Oregon.

#### 9. Future and Strategic Directions

Our project will begin work in August of 2006. Following the plan outlined in the proposal, work the data collector and data organization components will begin as soon as the project starts. Test data should be available from the project within the first few months of the project.

# **CRI: WORKIT: A Universal Wireless Open Research KIT First Year (August, 2005 to August, 2006)**

Prof. Henning Schulzrinne (PI)  
Columbia University  
[hgs@cs.columbia.edu](mailto:hgs@cs.columbia.edu)  
Award No: 0454288

Prof. Tom La Porta (CO-PI)  
Penn State University  
[tlp@cse.penn.edu](mailto:tlp@cse.penn.edu)  
Award No: 0453830

Mr. Scott Miller (CO-PI)  
Dr. Milind M. Buddhikot (CO-PI)  
Lucent Bell Labs  
{scm,mbuddhikot}@lucent.com  
Award No: 0454174

Prof. Elizabeth Belding-Royer (CO-PI)  
University of California, Santa Barbara  
[ebelding@cs.ucsb.edu](mailto:ebelding@cs.ucsb.edu)  
Award No: 0454329

## **Abstract**

The WORKIT CRI project is community resource development effort involving Lucent Bell Laboratories and three university research groups, namely Columbia University, Penn State and University of California, Santa Barbara. The 3-year project aims to create a universal Wireless Open Research KIT (WORKIT) to provide a flexible platform that can be used as a starting point for wide range of systems research in wireless networking. Specifically, the kit will provide reusable software and well documented off-the-shelf reference hardware for realizing wide range of mobile networking scenarios such as integrated 802.11+3G networks, mesh networks, high capacity mobile access routers and SIP-based multimedia networks. This report describes the WORKIT project in brief and details the first year activities targeted at creating the baseline kit to be released in September 2006.

## **Introduction**

Several technology trends promise exciting times ahead for research in wireless networking. Specifically, the client devices of varying sizes such as tiny smart sensors, small PDAs, and laptops equipped with one or more radio devices will proliferate. Wireless networks based on diverse wireless technologies such as 3G (1xRTT, EV-DO, UMTS: 144 Kbps – 14 Mbps/carrier, range of few miles, WIMAX (802.16: up to 75 Mbps, range of 1-5 miles), 802.11 WLANs (11 to 54 Mbps, short range 100-300 feet) will compete and co-exist to provide network access to the endusers. A mobile enduser with multi-radio device will be able to roam transparently across these diverse networks, preserve its application sessions, and receive a single bill for all its network access maintaining a customer relationship with a single service provider.

The convergence of these diverse trends leading up to ultimate emergence of an integrated, seamless wireless world will bring to light plethora of significant new research problems. Specifically, systems research that aims to prototype complete end-to-end solutions will prove daunting due to cross-layer interactions, complexity and scale of the problems.

Traditionally, a large cross-section of academic research in wireless networks has focused on analytical and simulation tools. Experimental research that prototypes realistic software and hardware systems, and instantiates entire end-to-end application or services is often a notable exception than a norm. The lack and need for experimental wireless tools and platforms which can be made available to the research community to foster such research was a common theme in the NSF Wireless Network Workshop Report [1]. Two representative quotes from the report illustrate this:

*“There is a critical need for provisions and procedures for procuring community tools for wireless networking. This would allow significant infrastructure for experimental research to be developed and to be reused by a wider academic community”.*

*“Remove the non-academic burden from the university research groups and allow a significant number of software tools to be reused by the community”.*

The academic groups often focused on specific aspects of an entire end-to-end system, lack time and resources to build realistic systems with components at all levels. Availability of a Universal Wireless Research Platform that is open-source, free, reusable and provides foundations for new exciting research in the field of wireless networks will serve as the first step in addressing these problems. Our project Wireless Open Research KIT (WORKIT) represents such a *Community Resource Development* project.

## Goals, Objectives, Targeted Activities

The WORKIT project brings together Lucent Bell Laboratories and three academic groups namely Columbia Internet Real-Time (IRT) Lab, MOMENT Lab at University of California, Santa Barbara (UCSB), and the Networking Research Center at Penn State, to exploit their prior and ongoing collaborations with the goal of producing a Wireless Open Research Kit (WORKIT). The produced kit will be based on successful *Integration of Two Access Technologies (IOTA)* project at Bell Labs [2, 3] and it will consist of a set of software tools and modules that can be used in a wide range of wireless, mobile and multimedia networking scenarios. The kit is intended to work “out-of-the-box” and will provide detailed documentation on the software modules, APIs, usage scenarios and off-the-shelf reference hardware.

By making the platform available to the research community, the project hopes to dramatically shorten the time to develop comprehensive system implementations of wireless related research ideas and avoid duplication of effort across the community. The availability of the kit will stimulate new systems-oriented research efforts that otherwise would be avoided due to complexity in building a system prototype foundation. For a large class of problems, it will also encourage validation of simulation results with working systems. The WORKIT kit will find applicability in many NSF funded programs and test beds. Additionally, the components of the produced kit can be used for educational courses on mobile networking as a hands-on lab tool or easy-to-use classroom supplement. The standardized kit, available as either software-only download or packaged on low-cost off-the-shelf hardware, will promote interaction across universities as well as with industrial research labs.

## Project Description

The WORKIT kit will provide following four basic capabilities: (1) *WLAN and Wide-area Cellular Interworking* (Figure 1): A complete baseline system foundation for the diverse issues in wireless networks such as mobility management, interlayer awareness, software algorithms for best network selection, reconfiguration, security, accounting, authentication, policy download/enforcement, and hybrid wireless networking. The WORKIT components here are (a) a 802.11+3G integration gateway, (b) a multi-radio gateway (when equipped with radio cards for 802.11a/b/g, EV-DO, 1xRTT, UMTS, HSDPA) gateway that can be deployed in mobile, transient hotspots such as vehicles and infostations mounted on poles, (c) multi-radio mobility client software that supports automated selection of best-connected radio interface and seamless roaming and session handoff across 802.11 WLANs and 3G networks, and (d) a Backend mobility server called *Home Server* to handle authentication, accounting, mobility management, and profile management for devices and end-users. (2) *Mesh module* (Figure 2) : A multi-radio mesh module that supports secure auto-configuration, monitoring, and optimized packet forwarding and routing, (3) *Mobile Access Router (MAR) and Group Mobility* (Figure 1): A module that provides support for multi-radio Mobile Access Routers (MAR) with link bonding features and efficiently supports high speed mobility of large groups of users. (4) *SIP multimedia module*: A module that provides software tools to implement, monitor and benchmark end-to-end high performance IP multimedia in infrastructure and mesh wireless networks.


### 1.1. First Year Project Activities

The thrust of the first year activities of WORKIT project was to produce the baseline kit and take necessary steps to announce its availability to the academic community. In that regard, Bell Labs team had the majority of the work items for the first year. The salient project activities are described in the

following.

**Internet dissemination:** Domain names [www.wor-kit.org](http://www.wor-kit.org) and [www.wirelesskit.org](http://www.wirelesskit.org) have been reserved for the WORKIT Project website. The first version of the website is online and accessible on the net. The baseline kit will be announced and available for download from this website around August 30, 2006. Another miscellaneous website <http://www.bell-labs.com/user/mbuddhikot/workit/index.htm> has served as the portal for the WORKIT project during early days of the project.

**Kit Planning and Coordination:** The Bell Labs IOTA toolkit implementation used Redhat 7.2 distribution with Linux 2.14 kernel. Redhat has since discontinued the distribution support. Also, our experience with this distribution indicated that it was less-than-satisfactory for embedded systems and for module dependency management. As such in the interest of ease of use and long term stability of WORKIT, we evaluated available Linux distributions. We selected Debian Linux as the best choice. Also, we decided to support the most current Linux kernel 2.6 for the baseline kit.



**Baseline WORKIT build:** The creation of the baseline WORKIT gateway has involved extensive cleanup, testing and packaging of the IOTA gateway code base on the Linux 2.6 under Debian distribution. We ported the Debian Linux on the SOEKRIS wireless gateway platform which will be available as a preferred hardware platform for the kit users. The major tasks here involved following software modules: (a) Automatic backhaul (Ethernet vs. 3G) detection (ifc) activation, (b) Stateful firewall (ipf), (c) Authentication state database (db), (d) Mobile-IP mode (mipagent), and (e) Simple-IP mode (udhcp, webauth). The WORKIT gateway will be ready by the Sept 1, 2006 release date.

The multi-interface mobility client software available in the baseline kit will run on Windows XP/2000. Currently, the refinement and testing of this module is in-progress and expected to be completed by the Sept 1, 2006 release date.

Early work on design, implementation and testing of the baseline *Home Server* component namely Mobile IP Home Agent (MIP-HA) will be completed by the Sept 1, 2006 release date. The team will also do some preliminary testing of additional commercial wide area cellular network data cards specifically Cingular HSDPA and Sprint EV-DO cards.

**Additional Workitems:** The Columbia team made early progress in design of two benchmarks: SIMPLEStone for characterizing performance of Presence servers and SIPstone for characterizing the performance of the SIP servers. Two technical reports describing these benchmark tools are available from the IRT lab website.

Ms. Yan Sun, a graduate student from Prof. La Porta's group in Penn State will be working full time during the summer'2006 with the Bell Labs team to assemble a WORKIT baseline kit from scratch. Her work will serve as a trial to iron out procedure required to reliably get the WORKIT hardware and

software to work out-of-the-box with minimal effort. She will also begin work on support for multiple 3G wireless backhauls such as Cingular HSDPA and Sprint EV-DO. One of the related projects in Prof. La Porta's group addressing topic of Group Mobility has produced a conference paper. The results in this work will be useful in the next year's task for support of group mobility in wide area cellular environments.

Mr. Krishna Ramachandran, a graduate student from Prof. Belding-Royer's group in UCSB, has worked on the project during the course of the 2005-2006 academic year. Specifically, he has developed a centralized channel selection algorithm for multi-radio mesh networks, and is in the process of developing a centralized routing protocol. He has also developed an implementation of AODV-ST, a hybrid version of the AODV routing protocol optimized for performance in mesh networks. Each of these solutions has been implemented and will be incorporated into the kit during the upcoming year. In addition, the UCSB group has experimented with and gained expertise about the Soekris devices which are preferred WORKIT platform through their incorporation into the UCSB MeshNet (<http://moment.cs.ucsb.edu/meshnet>) testbed.

## Indicators of Success and Major Accomplishments

The baseline kit will be announced and will be available for download from WORKIT project website around September 1, 2006. The WORKIT project will release following open source software that will be available for academic and non-profit use: (1) WORKIT baseline kit software which will include gateway software, multi-interface mobility client for Windows OS and early version of Home Server, (2) SIPStone and SIMPLEStone benchmark software from Columbia, and (3) Early prototypes of MeshClusters software from Bell Labs including joint work with University of California, Santa Barbara.

Two pieces of equipment that will be part of WORKIT: (1) The WORKIT mobile gateway, (2) WORKIT mesh-relay and mesh-gateway elements. Both these are constructed using off-the-shelf Soekris hardware platform. The parts lists for these equipments will be available on the WORKIT website by September 1, 2006.

## Publications and Patents

Note that a few of the publications are on topics for which the kit will release software modules in the second year of the grant. These publications, therefore, represent preparatory work in that regard.

1. K. Ramachandran, M. Buddhikot, G. Chandranmenon, S. Miller, E. Belding-Royer, K. Almeroth, "On the Design of Infrastructure Mesh Networks", *First IEEE Workshop on Wireless Mesh Networks, (WiMesh05), Santa Clara, CA, August 26, 2005*.
2. K. Ramachandran, E. Belding, K. Almeroth, M. Buddhikot, "Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks", *IEEE Infocom, Barcelona, Spain, April, 2006*.
3. P. Traynor, J. Shin, B. Madan, S. Phoha and T. La Porta, "Efficient Group Mobility for Heterogeneous Sensor Networks" *IEEE Vehicular Technology Conference (IEEE VTC Fall)*, September 2006.
4. V. K. Singh and H. Schulzrinne, "SIMPLEStone - Benchmarking Presence Server Performance", Department of Computer Science, Columbia University Technical Report.
5. H. Schulzrinne, S. Narayanan, J. Lennox and M. Doyle, "SIPstone - Benchmarking SIP Server Performance," Department of Computer Science, Columbia University Technical Report.

## Outreach Programs and Supported Students

The WORKIT grant currently supports following students: (1) Mr. Suman Srinivasan, Ph.D. student, Columbia University, (2) Ms. Yun Sun, Ph.D. student, Penn State University, (3) Mr. Krishna Ramachandran, Ph.D. student, University of California, Santa Barbara. Ms. Yun Sun will spend part of Summer'2006 with the Bell Labs team. Mr. Krishna Ramachandran has already spent two summers

working with the Bell Labs team and published two papers. A subset of this work on mesh routing and auto-configuration will be refined and included in the mesh module in WORKIT.

The WORKIT team has one female CO-PI and two female team members (one from Bell Labs and one from Penn State). The partner academic institutions will trial the kit in their labs and courses and attempt to develop small development projects for undergraduate thesis and course credits.

## Future and Strategic Directions

In the first year, WORKIT project focused on creation of the baseline kit rather than the full-fledged kit which will include Scalable Home Server, Linux version of Multi-interface Mobility Client, advanced Mesh Networking, Group Mobility and SIP multimedia networking support. The second year of the project is therefore necessary and critical to develop and deliver these features to make the kit comprehensive as planned and have maximum utility and impact. Following table shows in brief various work items for the WORKIT partners during the second year of the project.

Institution Name	Proposed work items
Lucent Bell Laboratories	<ol style="list-style-type: none"> <li>1. Home Server release: Release Date: Mar 1, 2007</li> <li>2. Linux Multi-interface mobility Client: Jul 1, 2007</li> <li>3. Release of Mesh module: Release Date: Jul 1, 2007 (joint with UCSB)</li> <li>4. Community Support and Web materials</li> </ol>
Columbia University	<ol style="list-style-type: none"> <li>1. SIP specific items: IRT SIP infrastructure servers, SIP benchmark tools</li> <li>2. Solicitation of community proposals and Evaluation of proposals</li> <li>3. Assembling of hardware kits: Release Date: Mar 1, 2007</li> </ol>
University of California, Santa Barbara	<ol style="list-style-type: none"> <li>1. Development, testing and release of Mesh Network monitoring tools</li> <li>2. Development, testing and release of mesh network auto-configuration, routing and mobility support features for mesh (joint with Bell Labs)</li> </ol>
Penn State University	<ol style="list-style-type: none"> <li>1. Support for multiple 3G backhails in Simple IP mode with automated selection</li> <li>2. Support for multiple backhails with mobility using Mobile IP</li> <li>3. Optional support for dynamic user assignment to backhails</li> </ol>

The WORKIT toolkit will be useful to instructors in universities and colleges to supplement in-class training on public wireless networking, mobile networks, AAA protocols, SIP and VOIP networks and Mesh Networking topics in undergraduate and graduate level classes. The kit will also serve as starting point for doing experimental measurements and new systems based research in above mentioned topics in wireless networks and avoid duplication of systems building effort. In fact, any discipline that requires remote connectivity for infield data measurements or network access for teams can benefit from WORKIT gateway and Mesh solutions. We also expect that numerous research ideas implemented using the kit will be absorbed back into the kit in the future release. We also hope that the WORKIT can find its way into National Research Testbeds (NRT) such as ORBIT and WHYNET.

## References

1. Mario Gerla, Tom La Porta (editors), "NSF Wireless Workshop Report", Chicago, IL, September, 2003.
2. M. Buddhikot, G. Chandranmenon, S. J. Han, Y. W. Lee, S. Miller and L. Salgarelli, **Integration of 802.11 and Third Generation Wireless Data Networks**, *Proceedings of IEEE Infocom 2003, April 2003*.
3. M. Buddhikot, G. Chandranmenon, S. J. Han, Y. W. Lee, S. Miller and L. Salgarelli, **Design and Implementation of a WLAN/CDMA2000 Integration Architecture**, *IEEE Communications Magazine issue on 3G+802.11 Integration, Nov 2003*.

# **CRI: A Linux Cluster for Research in Computational Nanotechnology and Education in Parallel Computing**

Award No. CNS-0551472; Mar 2006-Feb 2009

Florida A & M University-Florida State University College of Engineering

PIs: Namas Chandra, Usha Chandra, Ashok Srinivasan

chandra@eng.fsu.edu

uchandra@cis.famu.edu

asriniva@cs.fsu.edu

## **1. ABSTRACT**

The Departments of Mechanical Engineering at Florida A&M University-Florida State University College of Engineering (ME/FAMU-FSU), Computer and Information Sciences at FAMU (CIS/FAMU) and Computer Science at FSU (CS/FSU) propose to set up a Linux Cluster at the FAMU-FSU College of Engineering for developing space and time parallelization algorithms for molecular dynamics (MD) simulations of carbon nanotubes, and for education and training in the area of parallel computing. We seek a 71-compute node Linux cluster that will be installed at the College of Engineering and maintained by the college Computer Maintenance Services (CMS). FAMU is a Historically Black College and University (HBCU), while FSU is a majority institution.

## **2. INTRODUCTION/BACKGROUND**

We seek to establish a Linux cluster at the FAMU-FSU College of Engineering for developing space and time parallelization algorithms for molecular dynamics (MD) simulations of carbon nanotubes (CNTs). The cluster will also be used to enhance the experience and training of undergraduate and graduate students, especially under-represented minorities, in the area of parallel computing. Nanoscience and technology are based on fundamental understanding of the behavior of atoms and molecules and then manipulating them to obtain unprecedented properties. Computational nanotechnology is the enquiry of matter at the length scales of atoms and molecules through theory, modeling and simulation. Rapid strides in large scale computing methodologies combined with advances in hardware capabilities have made scientific computing a genuine investigative tool on par with experimental methods. This transformation is clearly evident in the field of nanotechnology where computation has assumed a leading role, partly due to the fact that controlled experiments at this very small length scale ( $10^{-9}$  m) are extremely challenging. Though a variety of computational techniques from ab-initio quantum mechanics to molecular statics/dynamics (MS/MD) to non-linear finite element methods (FEM) are used in computational nanotechnology, we will focus on developing parallel algorithms for MD.

MD is a numerical method that determines the positions and velocities of a given set of atoms as a function of time. MD is computationally intensive since it involves extremely small time steps of  $10^{-15}$  seconds and models individual atoms separated with a length scale of  $10^{-10}$  m. Thus, MD is limited to a few thousand atoms and a simulation time of a few nanoseconds. Realistic sizes of devices range on the order of  $10^{23}$  atoms and time scales of  $10^8$  seconds. Hence there is a big gap between the scales of MD simulations and that of experiments. The proposed cluster will enhance the MD capabilities (by reducing time and space limitations), through developing parallel algorithms. These algorithms will not only be useful in the area of nanotechnology, but in many other areas that use MD, for example, computational chemistry, molecular biology, and pharmaceutical sciences.

The PI has been working in the general area of computational mechanics and materials for well over eighteen years, funded by the Department of Defense, NASA, NSF, and private corporations. For the last seven to eight years, his research effort has focused on MD and multi-scale methods linking MD to FEM. The PI realized that the class of problems that can be solved using sequential code is very limited. For example, MD simulation of one CNT problem requires a million iterations for a simulation time of a nanosecond. This problem involves a turn around time of four weeks and thus is severely restrictive to solve any realistic problem. Consequently, in the last four years, the PI has been working with two computer science faculty (both co-PIs) to address this serious problem. This computer science and engineering team has already developed a parallel algorithm for the spatial decomposition for fine grained problems. The team is currently working on developing parallel algorithms for large time scale simulations. The computer science students and faculty at FAMU are also working on developing interactive scientific

visualization tools for parallel MD simulations. The research team is currently using a shared IBM-SP3 supercomputer at Florida State University (FSU) that is being phased out. Also, the turnaround time is ironically higher than that of a sequential simulation, due to the queuing system. Further, this FSU facility is not available to general faculty and students at Florida A&M University (FAMU), unless they are collaborating with a FSU researcher. In this proposal, we seek a dedicated cluster to improve our research capabilities, and also to provide access to students and faculty at FAMU for educational and training purposes in the area of parallel computing.

This section introduces the scientific problem of understanding the behavior of (CNTs) based composites using MD simulations. Here we briefly describe what makes CNTs so special, how MD simulations can be used to study their properties, the limitations of the present sequential algorithms, and our approaches to spatial and time parallelization.

## 2.1 Carbon Nanotubes

CNTs are long, thin cylinders of Carbon (diameters of  $10^{-10}$  m and lengths of  $10^{-7}$  m) that possess very unique properties and were initially discovered by Iijima in 1991. CNTs can be thought of as sheets of graphite rolled into cylinders. These are large macromolecules that exhibit stiffness much higher than steel, thermal conductivity higher than diamond, weight lighter than feather and strength greater than titanium. CNT can be either a conductor or an insulator depending on how they are formed (chirality). Due to the extremely large surface area/volume, CNTs can also be used as sensors. Thus CNTs exhibit a broad range of electronic, thermal and structural properties based on their diameter, length, and chirality. These nanotubes can be single-wall or multiple walls (cylinders inside other cylinders). CNTs are contemplated to be used in molecular electronics, structural composites, thermal coatings and sensors. In general, CNTs will be integrated within composites to achieve the unique combination of extra ordinary properties.

**CNT based Composites:** The properties of CNT based composites will depend on the type of the CNTs (diameter, chirality, nature of defects, single or multi-walled), functional attachments (length, type, density, chemical group) and the matrix. While one CNT is shown in Figure 1, typical composites will be made of hundreds of them oriented randomly or in specifically arranged order. The system will be affected by temperature-stress loading conditions that vary with time. MD simulations can be used to not only understand the science behind the problem, but also arrive at an optimal configuration for a specified application.




Figure 1. Schematic of the boundary conditions applied in the pullout test simulation. Left figure represents the simulation of a CNT connected by chemical attachments, being pulled at a specific rate mechanically. While one end of the attachment moves along with the CNT, the other end is fixed to the matrix. The mechanical force vs. displacement response (Figure right ) of these attachments then represents the mechanical behavior of the interface that we seek to understand.

## 2.2 Molecular Dynamics

Molecular dynamics is a computational simulation method that determines the position  $\vec{r}^i$  and velocity  $\vec{v}^i$  of every atom  $i = 1, \dots, N$  that is contained in the computational cell subjected to external boundary conditions (force, pressure, temperature or velocity). It used the basic equation of motion that solves for these  $6N$  variables ( 3 velocity and 3 displacement components for each atom).

Sequential computing using MD consumes significant effort even for limited system sizes and meaningful time periods of simulation. For example, to understand a system behavior for even a 100 nm<sup>3</sup> of material for one nanosecond, it requires five million iterations (time step 0.2 femto seconds) for about a million atoms. The computational effort is further increased when chemical interactions are required to be modeled accurately using complex multi-body potentials. For example, a problem involving the pullout of a CNT with 3000 atoms for 800,000 time steps takes about two days of computational time on a 2GHz single processor PC running Linux. Each run is often just a component of a larger multi-scale simulation. It is therefore untenable to perform these simulations without the use of effective parallelization. Thus, for these classes of problems, parallelization is not a luxury but a necessity.

### 2.3 Spatial Parallelization

For spatial parallelization, we derive our approach from the well-known techniques of domain decomposition parallelization. Traditionally, atom decomposition algorithm assigns N/P atoms to each processor without taking into account geometric information of atoms. Although this scheme ensures effective load balancing, it suffers from extra communication overhead since each processor needs to acquire all atom positions in order to calculate forces on its atoms. Spatial decomposition (cell based) algorithm overcomes this problem by division of the simulated system into spatial blocks or domains, each of which is allocated to a specific processor. Such a geometric allocation of domains minimizes the communication cost because only communication with adjacent processors is needed. But for this strategy results in load imbalance if the simulated system does not have a homogeneous distribution of atoms.

Goal of our decomposition algorithm is to ensure effective load balancing even for non-homogeneous systems with minimal communication cost. For achieving this, we decompose the domain based on the projections along the longest axis. The processors can be logically considered as a linear array having two neighbors each (except at the two ends). If the atoms are sorted, then the sorted array can be divided into equal blocks of N/P atoms, and each block assigned to a different processor. The initial sorting consumes significant time since it involves the entire domain consisting of all atoms. But subsequent sorting is done on local atoms, and hence does not cause much overhead.

Typically, total force acting on an atom is contributed by its first (two-body forces), second (three-body forces) and third neighbors (four-body forces). Therefore, the calculation of force on atom requires information about atoms up to and including their third neighbors. In distributed (parallel) environment it makes sense for a processor to duplicate the information of atoms not associated with it at the expense of memory and time so that it can calculate forces on its "own" atoms independently of other processors. Division of domain along sorted projection axis ensures that atoms on neighboring processors are also neighbors in the physical nanotube. As a result, each processor needs data from adjacent processors, for atoms that are within a distance of three cutoffs on the projection axis on either end; so each processor sends this "buffer" data to the neighbors.

### 2.4 Time Parallelization

The basic incremental time scale permitted in MD of the order of 10<sup>-15</sup> second, and a meaningful interpretation of results can only be made with a bare minimum simulation time of around 10<sup>-9</sup> second, requiring about a million time steps. Figure 1 also shows that the physical domain is not homogeneous (with uneven distribution of attachments), and even for a case of 64 processors, the number of atoms per processor is of the order of only 50 for a 3000-atom CNT. Such fine granularity leads to communication cost being an important bottleneck to effective parallelization.

Our approach to time parallelization is to have each processor simulate different time intervals. However, MD being an initial value problem, the processors will not know the initial state for their time interval until the previous processor has its determined its final state. We use fine-scale computations to find out a relationship between prior results and current simulation. We then use coarse-scale response of prior simulations, and their relationship to the current one, to predict the initial state for each time interval. We then perform accurate simulations in parallel. Once the simulation is completed, we verify if the starting states of all the processors are correct; if so we have achieved solving the problem for P time intervals in a single step. By using a novel predictor-verifier method, we

demonstrated in a model problem, that time can also be parallelized. This will be very useful in achieving long-time simulations.

### **3. GOALS, OBJECTIVES, TARGETED ACTIVITIES**

The goals of this project are to enhance research in the area of parallel algorithm development both in space and time domains for computational nanotechnology applications. In addition, the cluster will enable the Department of CIS/FAMU to train undergraduate and graduate students in the area of parallel computing. We plan to use the cluster in the following three collaborative research projects, namely, 1) Nanotechnology Exploratory Research (NER): Scalable Techniques and Algorithms for Massively Parallel Nanomaterial Simulation for Long Time Behavior 2) Developing fine grained spatial parallel algorithms to simulate the mechanical Behavior of Carbon Nanotube (CNT) based Composites and, 3) Developing interactive MD visualization tools for Sensor applications: These projects are funded by NSF, ARO (U.S. Army) and NSA (National Securities Agency) respectively.

### **4. INFRASTRUCTURE AND FACILITIES**

The Linux cluster has 71 compute-nodes with two Intel Xeon 3.0GHz processors per node with Infiniband interconnect. The compute-nodes also have an 8MB cache. The cluster will run under Fedora Core 4 and the software includes cluster management software and the Portland Group Development kit that includes compilers, parallel debuggers (up to a maximum of five users at a time) and performance analysis tools (on 64 processors).

### **5. PROJECT DESCRIPTION**

The three PIs have been collaborating successfully for the last three years and have been awarded research grants from NSF, AFOSR and NSA on three different projects that require parallel computing. This section describes the above three projects and the two courses to be developed at the CIS/FAMU department, namely, parallel computing and advanced computer architecture.

#### **5.1 Research Project 1 – NER: Scalable Techniques for Massively Parallel Nanomaterial Simulation for Long Time Behavior (PIs: Ashok Srinivasan and Namas Chandra)**

The PI Ashok Srinivasan and Namas Chandra have been awarded a grant from NSF CMS-0403746 titled “NER: Scalable techniques for massively parallel nanomaterial simulations for long-time behavior” and this section describes the features of this project. The investigators address a class of problems involving an assembly of carbon nanotubes, wherein interfaces play a key role. This class of problems pertains to long time scale simulations in which classical transition state theories are not applicable. The solution strategy is based on harnessing the power of massive parallelism. Conventional parallelization through spatial decomposition will not be effective since that will lead to fine granularity. The proposed effort is aimed at time-parallelization using a predictor-verifier approach. One of the key research issues is to develop appropriate predictors. Successful results from this endeavor can be integrated with multi-scale simulations that can predict material properties to time scales several orders of magnitude greater than that today. The above research effort has applications in the areas of nanocoatings, nanosensors, nanoelectronics and nanocomposites. In the current stage of the development of nanotechnology, computation (theory, modeling and simulations) is playing a leading role, compared to experiments, because of size effects. Some of the fundamental understanding of both physics and computations has potential use for a wider class of applications, including nano-biotechnology.

#### **5.2 Research Project 2 - Mechanical Behavior of CNT based Composites (PI: Namas Chandra)**

The PI Namas Chandra has been awarded a grant from AFOSR FA-9550-04-1-0202 titled “Mechanical Behavior of CNT based Composites”. CNTs possess extra-ordinarily high values of stiffness (1 TPa) and strength (40 GPa), and are contemplated to be used as fibers in polymeric composites. Unfortunately, these excellent mechanical properties of CNTs have not translated into that of CNT based polymer composite properties. One of the key problems is that, since CNTs are atomically clean structural elements, the nanoscale interfaces between CNTs and the polymer matrices do not bond properly and do not transfer the load. In this Air Force funded project, we explore the use of different chemical attachments (functionalizations) to achieve enhanced composite properties by increasing the bonding between CNTs and polymer matrices. Here, we plan to conduct numerical simulations of CNTs with and without surrounding polymeric matrices, when the nanotubes are subjected to mechanical loads. Typical computational cell contains about 3,000 to 10,000 atoms, depending on whether the polymer matrix is included or not. We simulate pulling the CNT fiber out of the matrix with different levels of interfaces all the way from no bonding (vanDer Waals alone) to many types of functionalizations. Force-displacement response of the interfaces using molecular dynamics simulations will determine the effectiveness of the interface in achieving the load

transfer. Many parameters such as the type, length and density of chemical functional groups on different diameters and chiralities of CNTs will be examined to determine the optimum combinations. In dealing with real composites, we need to use many CNTs oriented arbitrarily within a polymer matrix. Such system will involve at least 100,000 to 200,000 atoms and a million time steps. We will integrate both time *and* space parallelization for this problem. The algorithm and implementation showed near linear scale-up and showed an efficiency of 75% for a 500 atoms/processor and is much finer than any previous effort.

### **5.3 Research Project 3 - Carbon Nanotube Based Chemical-Bio-Radioactive-Explosive Sensors: A Computational Simulation** (PIs Namas Chandra and Usha Chandra, 2004-2006):

CNTs have unique thermo-mechanical properties that far exceed that of any man made or synthetic materials to date. Single or multi walled nanotubes possess a combination of properties never seen before. They have electrical conductivity higher than copper, thermal conductivity higher than diamond, stiffness much greater than steel, strength larger than Titanium and weight lighter than feather. For every gram of these materials, they have a surface area greater than 1875 square meters. In this project, we propose to examine the feasibility of these wonder materials to act as Chemical-Biological-Radioactive-Explosive (CBRE) sensors using sequential and parallel computational methods. CNTs by themselves are not capable of acting as sensors, but additions of specific chemical functional groups that bind with CNTs and attract offending species make it possible. Since there are plethora of possibilities of species and processing conditions, and also because, experimentally manipulating materials at atomic scales is extremely difficult, numerical methods appear to be the only solution. This effort involves the implementation of spatial decompositions of MD codes with different potential energy functions (e.g. Lennard-Jones, Morse, and Brenner). Each potential energy function requires different types of parallel algorithms for effective implementation. Thus, in this proposed work, we will extend the study to the application of CNTs to CBRE sensors. This project involves both U.S. graduate and undergraduate students (especially African American) spanning both Computer Science and Engineering.

### **5.4 Courses in Parallel Computing and Advanced Parallel Architectures to be offered in the Department of Computer and Information Sciences at FAMU**

As a part of this proposal, the Co-PI, Usha Chandra will be developing two courses, namely, Parallel Computing and Advanced Parallel Architectures. The first course on Parallel Computing will be offered as an advanced undergraduate elective course and co-listed as a graduate elective course also. The course on Parallel Computing will be offered in the spring semester and the course on Advanced Parallel Architectures will be offered in the fall semester on a regular basis and the course announcements will be made available to the FAMU-FSU College of Engineering and the other Science and Mathematics departments at FAMU in order to train students in interdisciplinary research.

## **6. INDICATORS OF SUCCESS AND MAJOR ACCOMPLISHMENTS**

In order to ensure that the research and education components of this project are successful, the PIs will track the following and will submit an annual report and a final report with the above details.

- the users within and outside the research group,
- the number of computing hours on the cluster,
- the number of thesis generated directly by the PIs and other indirect users of the cluster,
- the number of publications of researchers in the group and outside the group,
- the number of funded research projects, and,
- the number of students enrolled in the two courses from both CIS/FAMU and other disciplines.

## **7. FUTURE AND STRATEGIC DIRECTIONS**

We believe that this infrastructure will enable us, especially, CIS/FAMU to obtain new resources of research support. As an example, we would like to submit a Research Experience for Undergraduate (REU) proposal to NSF to facilitate training students (especially minorities) during summer in the area of scientific computing. We have strong relationships with other HBCUs through the ADMI organization and we would like to use this as a vehicle to recruit students for this REU project. This would also greatly strengthen our CIS/Masters in Software Engineering Sciences (CIS/MSES) program. Currently, we have about twenty graduate students of which twelve are full-time. The strategic plan for MSES program in CIS/FAMU is to have a sustained enrollment of 40 students and this infrastructure will greatly help us to recruit students for the CIS/MSES program.

# **CRI: Enabling Globally-Distributed Electronic Collaboration (GloDEC) for Expertise Studies and Human and Social Dynamics**

## **Research Planning Grant**

Proposal Number CNS0452180, Year 1

Florida International University

Irma Becerra-Fernandez, Ph.D. and Michael Prietula, Ph.D.

[becferi@fiu.edu](mailto:becferi@fiu.edu), [mike\\_prietula@bus.emory.edu](mailto:mike_prietula@bus.emory.edu)

[www.kmlab.fiu.edu](http://www.kmlab.fiu.edu)

### **1. Abstract**

The College of Business Administration at Florida International University (FIU), a 33,000-student minority institution with the highest enrollment of Hispanic students in the country; requests the necessary funds to establish a technology infrastructure (previously known as Globally-Distributed Electronic Collaboration (GloDEC) and now renamed as Representation, Education, Exploration, and Discovery REED) necessary to launch Project Ensayo. The two Co-PI's (FIU: Becerra-Fernandez and Emory: Prietula) expect to build a hardware and software infrastructure to examine the decision-making and organizational complexities that arise from unique, complex, and significant events associated with the management of disasters, such as recovering from a hurricane in South Florida. The infrastructure is also being designed to simulate, as well as learn, how organizational events, rules, and contexts evolve over the course of a large-scale disaster scenario. The focus of this project is a unique Center located at the Miami-Dade County Office of Emergency Management (OEM). This office is the lead agency in an emergency event and the Emergency Operations Center (EOC) is the site for all of the emergency management operations. The design of the underlying simulation architecture, called REED, is presented in the next sections. The REED infrastructure will serve to simulate the events that take place at the EOC during emergency events in Miami-Dade.

### **2. Introduction/Background**

The original scope of the planning grant was to enable the two Co-PI's to plan research activities and propose a technology infrastructure called 'GLODEC' (globally-distributed electronic collaboration). The planning grant would allow the Co-PIs to assess the feasibility of expanding the present Knowledge Management Laboratory at Florida International University (FIU) College of Business Administration to support doctoral studies in expertise and human and social dynamics (HSD).

This proposal comes from the Decision Sciences and Information Systems (DSIS) department of the College of Business Administration at Florida International University (FIU), a 38,000-student minority institution that has the highest enrollment of Hispanic students in the country. The DSIS department grants undergraduate, Master's, and Ph.D. degrees in management information systems (MIS).

One of the limitations of the planning grant proposal expressed by the reviewers was that it could have provided more detail on the activities proposed for the period of the grant, particularly with regard to the technology. This is needed because the two research areas (expertise studies and HSD) require quite different technological tools. For this reason, we have put a significant effort in refining our research objectives for the required infrastructure, and we will define the revised project scope in the next section. This proposal has the potential to influence the education and, perhaps, training of minorities. We have significantly expanded the potential impact of this technology infrastructure by refocusing its use not solely for PhD research, but also for the actual training of emergency response personnel, as will be described in the next sections. In addition, this infrastructure will also support educational activities as part of the undergraduate business and MBA program, thus having a significant impact on the preparation of Hispanic students for careers that require knowledge management (KM). Furthermore, we're planning to also focus some of our research efforts in themes related to KM and HSD in the Hispanic context. Some possible research themes include 1) Research into the transfer of expertise from traditional sources

to people with Hispanic culture 2) Human and social dynamics of Hispanic organizations; and 3) Expertise transfer in emergency management in a primarily Hispanic context.

We expect that this project will enable us to set up and test the infrastructure for inter-university and inter-organizational collaboration, including both private and public organizations, and to enable the research and training that will support degree programs that the faculty and graduate students, many of them representing minority groups, will benefit from.

### **3. Goals, Objectives, Targeted Activities**

Project Ensayo is a large, multidisciplinary effort that is focused on building an infrastructure to examine the decision-making and organizational complexities that arise from unique, complex, and significant events associated with the management of disasters, such as recovering from a hurricane in South Florida. The heart of Project Ensayo is an extensible, distributed, Web-based simulation architecture, called REED (Representation, Education, Exploration, and Discovery). The REED architecture is mission-oriented and specifically designed to underlie the modeling of the infrastructure Miami-Dade EOC. The project goal is to connect Ensayo to every key decision making office in the Miami area that is responsible for decisions under the scenarios addressed, within the EOC context. It is important to note that the REED architecture and Ensayo are envisioned to be multi-use: for research (e.g., determining social networks of communication emerging under different events, testing different responses rule sets, timing), education (e.g., training emergency response teams), exploration (e.g., determining inter-office rule optimization), and discovery (e.g., real time monitoring and advising during actual events).

The Representation components allow flexible definition of actors in the simulation, such as individuals, teams, institutions, events, or artificial agents. For example, if a defined actor is the Mayor of Miami's office, then the decisions made by that actor could be realized by individuals actually sitting in the Mayor's office, or as a surrogate entity representing the decisions of the office itself, without detailing individuals within it. Events are broadly and flexibly defined occurrences, exogenous or endogenous, to the simulation. For example, a hurricane may be characterized as a top-level event, which can spawn ancillary events leading to causal chains and consequences. On the other hand, data being sent from one office to another, or a request for information between offices, are also events. Education components help craft the scenarios and monitor the decision and communication traffic in a form that serves specific pedagogical goals for the associate offices. For example, a particular office, such as the Dade County Public Health department, can interact with REED on their own or with other key actors, such as the meteorologists from the National Hurricane Center who issue the evacuation reports. Interacting actors can actual individuals or simulated agents. Exploration components allow data and knowledge mining of the consequences of decisions, occurring in post-simulation analysis or in real time. These components also afford dynamic, data-driven capabilities to acquire data or decisions from specific sources based on events, decisions, or the changing structure of the emergent, cross-agency participants. Forcing new communication links and driving communication exchange are essential elements of these components. Discovery components allow for real-time and post-simulation analysis, allowing meta-analytics to be brought to bear. This, perhaps, is the most valuable set of components for learning and accumulating knowledge that crosses individuals, organizations, agencies, and events.

Ensayo and REED are designed to be XML-based and to interact with exogenous data (e.g., from the National Hurricane Center). The overall architecture is fairly typical, supporting scenarios specification (Scenarios), which includes information on decisions and decisions rules (RuleSets), information on organizational structure, communication and resources (StructureSets), and other ancillary information on driving events in the scenario (ScriptSets). REED also has the ability to define events, objects and rules at a meta-level (MetaREED). For example, when run as a pure simulation, MetaControlScripts will define a series of simulations (i.e., an experimental design).

#### **4. Infrastructure and Facilities**

We're currently defining the components of REED, the necessary technology infrastructure to support the needs of Project Ensayo.

#### **5. Project Description**

A critical part of natural disasters involve dynamic decision-making before, during and after the events (e.g., Samii et al., 2002a,b). The Stafford Act establishes the federal programs and processes used by the federal government for granting both emergency and major disaster aid to individuals, profit and not profit organizations, local governments and states. Depending on the type of disaster, emergency or major disaster, the President will engage federal government in emergency response activities (e.g. debris removal, temporary housing, distribution of food and medicine) that are estimated to require less than US \$5M, or more than this amount respectively.

In December of 2005, the US Department of Homeland Security (DHS) issued the National Response Plan (NRP) as a comprehensive framework for the management of emergencies of national significance within the National Incident Management System (NIMS), and Incident Command System (ICS) as the incident management organization (command, operations, planning, logistics, and finance/administration). A careful investigation of the government response to Hurricane Katrina, one of the major natural disasters in the history of the United States, reveals widespread dissatisfaction with the results following the government's overwhelmed capabilities to respond (GAO, 2006).

In fact the lessons learned from Katrina were very similar to those identified with Hurricane Andrew, which occurred ten years prior (showing that there hasn't been much advancement in this area) and include: "(1) clearly defining and communicating leadership roles, responsibilities, and lines of authority for catastrophic response in advance of such events, (2) clarifying the procedures for activating the National Response Plan and applying them to emerging catastrophic disasters, (3) conducting strong advance planning and robust training and exercise programs, and (4) strengthening response and recovery capabilities for a catastrophic disaster" (GAO, 2006). Local emergency response protocols are defined by the NRP and the ICS. The number of organizations involved in emergency response ranges depending on the magnitude of the disaster. For example, it's estimated that at around sixty organizations at the federal level alone, were involved in the recovery efforts following hurricane Katrina. Furthermore voluntary organizations, like for example the International Federation of Red Cross as well as local voluntary organizations, may follow a different set of protocols (Samii, Van Wassenhove, Kumar, and Becerra-Fernandez, 2002a, 2002b).

The State of Florida, considered one of the most effective in disaster management (REF), has designated the State Emergency Response Team (SERT) with the mission to "Ensure that Florida is prepared to respond to emergencies, recover from them, and mitigate against their impacts". The SERT has identified eighteen hazards that pose an emergency threat to Florida, which are: wildfires, thunderstorms, tornadoes, lightning, flood, terrorism, drought, heat waves, hurricanes, cold, animals, nuclear, hazardous materials, cyber attacks, information warfare, aircraft, and bombs. The division is organized into four bureaus: compliance planning and support, policy and planning, preparedness and response, and recovery and mitigation. The Bureau of Preparedness and Response is responsible for developing and maintaining the State's ability to effectively respond to a wide variety of threats and has two sections preparedness and response. The Response Section coordinates emergency response at the state level, and which provide the necessary technical assistance to county governments. The Miami-Dade County Office of Emergency Management (OEM) is the lead agency in an emergency event. The Emergency Operations Center (EOC) is the site for all of the emergency management operations.

In order to respond to a hurricane emergency, the EOC depends on a large array of organizations that are organized into three groups: the Public Safety Functional Group, the Human Services Functional Group, and the Infrastructure Functional Group.

Public Safety Group. The Public Safety Functional Group includes the following organizations: National Park Service, Florida Fish and Wildlife Commission, US Coast Guard, the Department of Environmental Resources Management (DERM), Miami-Dade Fire Rescue Dept., Florida National Guard, Animal Services, Miami-Dade Corrections Dept., Florida Dept. of Law Enforcement, Florida Highway Patrol, and Miami-Dade Police Dept. The Public Safety Manager coordinates the activities of the public safety functional group, with the assistance of the Public Safety Assistant.

Human Services Group. The Human Services Functional Group consists of the following organizations: Dept. of Human Services, Team Metro, Salvation Army, Greater Miami Convention and Visitors Bureau, American Red Cross, Miami-Dade Public Schools, Miami-Dade Voluntary Organizations Active in Disaster (VOAD), Dept. of Mental Health, Miami-Dade Housing Agency, Florida Dept. of Children and Families, Miami-Dade Health Dept., Florida Agency for Health Care Administrators (AHCA), and the Miami-Dade Fire Rescue Emergency Management Services. The Human Services Manager coordinates the activities of the Human Services Functional Group, with the assistance of the Human Services Assistant and the Special Needs Coordinator.

Infrastructure Group. The Infrastructure Functional Group consists of the Miami-Dade Solid Waste Dept, Miami-Dade Water and Sewer, South Florida Water Management District, Miami-Dade Transit-Regular Services, Miami-Dade Transit Evacuation, Miami-Dade Public Schools, Miami-Dade Public Works, Florida Dept. of Transportation, Miami-Dade Parks Dept., Agriculture Extension, City Gas, Miami-Dade Enterprise Technology Services Dept. (ETSD), BellSouth, and Florida Power & Light. The Infrastructure Manager coordinates the activities of the Infrastructure Functional Group with the assistance of the Infrastructure Assistant. The Operations Section Manager coordinates the activities of the three functional groups, with the assistance of the EOC Support Manager, the Operations Section Assistant, and the Planning Situation Assessment Assistant.

A number of additional organizations are represented in the periphery of the EOC: Florida Division of Emergency Management (DEM), neighbor counties emergency management liaisons including Monroe County, Miami Beach, North Miami, Homestead, Coral Gables, Hialeah, City of Miami, Florida City, Broward County, Martin County, and Collier County. In addition, representatives from the Air Force Reserve Base, the airports, and the Port of Miami are also included.

The EOC depends on a number of state-of-the-art tools that it uses to manage emergencies: (1) Hurrevac, a software developed jointly by the National Hurricane Center, the U.S. Army Corps of Engineers and FEMA to track tropical cyclones and provide a continuous flow of information to emergency managers; (2) SLOSH II, software developed jointly by the National Hurricane Center Storm Surge Group, the U.S. Army Corps of Engineers, USGS and FEMA with input from the state of Florida and several local emergency managers including Miami-Dade County. It demonstrates probable storm surge based on size, direction and forward speed of a storm; (3) SALT (Storm Action Lead Time), software developed by Miami-Dade County OEM and ETSD to provide a check-list of pre- and post-storm activities; and (4) Snapshot, software developed by Miami- Dade OEM to provide virtually instant information on damages caused by a storm or flood event.

During such events, critical decisions must be made that involve cross-organizational and cross-agency coordination, and sharing of data, information and knowledge. As these events and their contexts are infrequent and varied, the nature of the decisions, where they are made, who makes them, the data and information resources required to make and monitor them, and the location of available knowledge to drive them, are either unknown, unavailable, or both. People leave the organizations and the informal rules that serve as the “glue” that affords the very ability to function is lost. Furthermore, coupled with these events are also the unlikely, but not inconsequential, coincident disasters, which have been essentially unexplored. There are associated and substantial risks associated with coincident events of terrorist opportunity, for example, where airborne bio-terrorism agents are released during a hurricane. This would result in data being available, but essentially masked, within the structure and dynamics of unfolding events, resulting in few if any decisions made regarding the masked attack. One mechanism to explicate, educate, and replicate these rules, both formal and informal, is through a software environment. That environment is called Ensayo.

## **6. Indicators of Success and Major Accomplishments**

Ensayo and REED are being developed in conjunction with the Miami-Dade Emergency Operations Center. Such emergency centers are “organizations” but suffer from the lack of normal conditions that permit organizational learning in the traditional sense. The participants are fluid and carry little organizational knowledge, the events are rare and diverse, the complexity of interactions within and between organizations is high and time-sensitive, and the pressures and risks are enormous. We envision that Ensayo and REED will contribute to research and education to better allow, and even to assist, such centers in their time of need.

## **7. Publications/Patents**

Becerra-Fernandez, I. and Prietula, M. Project Ensayo: Integrating simulation, training, discovery, and support”, Presentation and Proceedings of the North American Association for Computational Social and Organizational Science (NAACSOS 2006), June 22-23 2006, Notre Dame, Indiana.

Government Accountability Office (GAO) (2006) Hurricane Katrina: GAO’s Preliminary Observations Regarding Preparedness, Response, and Recovery. March 8.

Samii, R. Van Wassenhove, L., Kumar, K. and Becerra-Fernandez, I. (2002a) IFRC: Choreographer of Disaster Management – The Gujarat Earthquake. Published by INSEAD, Fontainebleau, France.

Samii, R. Van Wassenhove, L., Kumar, K. and Becerra-Fernandez, I. (2002b) IFRC: Choreographer of Disaster Management – Preparing for Tomorrow’s Disasters. Published by INSEAD, Fontainebleau, France.

## **8. Outreach Programs and Supported Students**

The broader impacts of the proposed activity are:

- 1) To help prepare predominately minority students to work in areas related to expertise studies, knowledge management, and human social dynamics.
- 2) To help train emergency management personnel throughout state and local governments.
- 2) To support the educational needs of graduate students with state of the art collaboration technologies
- 3) To design and disseminating new inter-institutional Ph.D. courses.
- 4) To attract new students to the graduate programs at FIU.

## **9. Future and Strategic Directions**

The future tasks for this planning grant are:

- 1) To complete the planning activities of this grant culminating in the CRI technology infrastructure proposal next year.
- 2) To develop collaboration relationships with all the key private and government organizations involved in emergency management at the EOC.
- 3) To identify other funding sources that may strategically support the efforts of this grant.

# **MII: Infrastructure for Research and Training in Database Management for Web-based Geospatial Data Visualization with Applications to Aviation**

PI: Naphtali Rishé

High Performance Database Research Center  
School of Computer Science, Florida International University  
Miami, FL 33199; rishen@fiu.edu; <http://hpdr.ccs.fiu.edu>

Co-PI: Boonserm Wongsaroj

Computer Science Division, Florida Memorial University  
15800 NW 42 Ave  
Miami, FL 33054; bwongsar@fmuniv.edu; <http://www.fmuniv.edu>

## **Introduction**

Florida International University (FIU) is one of the largest majority-minority doctoral-granting universities in the United States. Nearly 70% of our students are minorities. The University has the largest contingent of Hispanic students of any doctoral-granting university in the country, and graduates the most Hispanic computer science and engineering students in the Nation. The High Performance Database Research Center (HPDRC) was founded in 1994, and is associated with the School of Computer Science at Florida International University. HPDRC conducts research on database management systems and the Internet dissemination of data, leading to the development of new types of database systems and refinement of existing database systems.

For more than 122 years, Florida Memorial University (FMU) has provided students with an education that otherwise may not have been available to them. As South Florida's only Historically Black College and one of the State's oldest academic institutions, FMU has a long history and commitment to the free exchange of ideas, the pursuit of knowledge, and the ongoing transmission of African-American history and culture. Currently, the college enrolls approximately 2,150 students, employs more than 250 full-time faculty, staff, and other support personnel, and maintains an annual budget of more than \$43 million. The institution is related to the Baptist church community of Florida and attracts the majority of its students from Florida, the eastern half of the United States, and the Caribbean islands. The college is accredited by the Commission on Colleges of the Southern Association of Colleges and Schools, chartered by the State of Florida, and is a member of the college Fund/UNCF.

FIU and FMU are urban universities whose surrounding community base is substantially comprised of under-represented minorities: 90% of the students of Miami-Dade County Public Schools are members of under-represented minorities, including 60% Hispanic, 28% Black non-Hispanic, and 2% Other. One goal of this project is to establish a regional outreach program to attract talented local minority students to FMU and FIU. Without the support of this project, those students would otherwise not be able to take advantage of the career and educational opportunities, or would have to attend an out-of-state university (a non-favorable choice of many of the local minority students).

The general goals of the project, now in its second year, are to provide an infrastructure that will enable FIU's HPDRC and FMU to perform database and Internet research and to better recruit and retain minority students through their B.S., M.S., and Ph.D. degrees. One goal of this consortium is to encourage FMU's students and others to enter FIU's masters or doctoral program in computer science, creating a pipeline of students from FMU, a four year Carnegie-III HBCU, to FIU, a Carnegie-I HSI. Students participate in in-depth research necessary to enhance TerraFly, a web-based geospatial data viewer stemming from research performed at FIU's High Performance Database Research Center (HPDRC) under grants from NSF and other agencies.

Researchers work in Affinity Groups that consist of undergraduate and graduate students, post-docs, and faculty members. These groups provide the framework that enables deepening knowledge of a field by providing a physical setting in a cooperative research-engaging environment. The project aims to enhance TerraFly's underlying data storage mechanism, client-server interaction, user interface, its ability to overlay additional information layers, and its ergonomics of use and maintenance. The following specific research issues are being addressed: Efficient and ergonomic dissemination of imagery with spatio-temporal data overlays, Data storage and querying methods for spatio-temporal data, Integration of heterogeneous data sources including relational and semantic databases and web sources with the spatial data, and Testbed applications using geospatial data dissemination for aviation needs.

The project-provided infrastructure is being used to perform this research, as well as the training of students at both institutions. The deployment of infrastructure at FMU enables FIU's distributed database research to be incorporated into TerraFly research and FMU's students to gain valuable hands-on experience with the infrastructure on a daily basis.

HPDRC maintains a WWW page describing its projects and staff at <http://hpdrc.cs.fiu.edu>. The TerraFly project web page is available at <http://TerraFly.fiu.edu>.

## **Fourth Year Accomplishments**

### **Research**

The research supported by the MII infrastructure and the collaboration between FIU and FMU has focused on several areas:

- Searching Spatial Databases.
- Native XML Database Management.
- Storing and Managing Raster-Based Geospatial Data.
- Semantic Modeling and Virtual Quadrilaterals for Geospatial Data Management.
- Data Clustering.
- Video Event Detection.

We have had particularly good results in the first area this year. An increasing number of applications require the efficient execution of nearest neighbor queries constrained by the properties of the spatial objects. Due to the popularity of keyword search, particularly on the Internet, many of these applications allow the user to provide a list of keywords that the spatial objects (henceforth referred to simply as objects) should contain, in their description or other attribute. For example, online yellow pages allow users to specify an address and a set of keywords, and return businesses whose description contains these keywords, ordered by their distance to the specified address location. As another example, real estate web sites allow users

to search for properties with specific keywords in their description and rank them according to their distance from a specified location. We call such queries spatial keyword queries.

An example of a spatial query would be 'find the nearest hotels to point [30.5, 100.0] that contain keywords internet and pool'. Unfortunately there is no efficient support for spatial keyword queries. Instead, current systems use ad-hoc combinations of nearest neighbor and keyword search techniques to tackle the problem. For instance, an R-Tree is used to find the nearest neighbors and for each neighbor an inverted index is used to check if the query keywords are contained. We show that such two-phase approaches are inefficient.

We are developing a method to efficiently answer spatial keyword queries, which is based on the tight integration of data structures and algorithms used in spatial database search and Information Retrieval. In particular, our method consists of building an Information Retrieval RTree (IR<sup>2</sup>-Tree), which is a structure based on the R-Tree. At query time an incremental algorithm is employed that uses the IR<sup>2</sup>-Tree to efficiently produce the top results of the query.

The IR<sup>2</sup>-Tree is an R-Tree where a signature is added to each node  $v$  of the IR<sup>2</sup>-Tree to denote the textual content of all spatial objects in the subtree rooted at  $v$ . Our incremental spatial keyword search algorithm, exploits this information to locate the top query results by accessing a minimal portion of the IR<sup>2</sup>-Tree.

To date, this work has the following contributions:

- The problem of spatial keyword search is defined.
- The IR<sup>2</sup>-Tree is proposed as an efficient structure to store spatial and textual information of a set of objects. Efficient algorithms are also presented to maintain the IR<sup>2</sup>-Tree, that is, insert and delete objects.
- An efficient incremental algorithm is presented to answer spatial keyword queries using the IR<sup>2</sup>-Tree. Its performance is analytically and experimentally evaluated and compared to state-of-the-art approaches. Real datasets are used in our experiments that show the dramatic improvement in execution times.

### **TerraFly**

Our TerraFly web-based spatial data visualization research, made public at <http://TerraFly.FIU.edu>, provides an outlet for outreach for our research. The website makes available approximately 20 TB of data provided through a Cooperative Research and Development Agreement (CRADA) with the United States Geological Survey (USGS), as well as agreements with Space Imaging and Navigation Technologies. The data available via TerraFly includes 1m resolution aerial photography of the United States and will soon include extensive Landsat 7 coverage of the populated world and 1ft resolution aerial photography covering urban areas around the United States. The TerraFly website, operated as a public service, presently receives over 10,000 unique users and has been featured in *Science* magazine, *Nature* magazine's web site, *Parade* magazine, the *New York Times*, and as a Yahoo! Pick of the Week and *USA Today* Hot Site, as well as other press coverage world wide. Feedback received from educators, public officials, government employees, and the general public has been overwhelmingly positive.

### **Involving Student Researchers through Affinity Groups**

Researchers involved in the database thrust work in Affinity Groups. These groups bring faculty members, post doctoral associates, graduates, undergraduates, and high school students together to pursue deepening research in a specific area. We have found that by bringing researchers of varying ability together, individuals are afforded the opportunity to act as both learner and mentor, thus enhancing education opportunities and personal growth.

### **Involving FMU and FIU Students and Faculty in the Research Effort**

FMU and FIU faculty and FMU and FIU students have performed collaborative research under the support of this grant. FMU and FIU students have made visits to the other institution's campus to explore this research, working along side one another during these visits and exchanging e-mails and phone calls to discuss research at other times. The collaboration has resulted in six papers which have been published and an additional five that have been accepted for publication. Each is co-authored by a combination of FIU and FMU faculty members; all of the papers are co-authored by students – some by students at both institutions.

### **External Research Opportunities**

The MII collaboration has enabled FMU's undergraduate students to compete for scarce slots in summer research programs. In the Summer of 2005, two FMU students, Ahmad Gallon and Brenton Williams, participated in the PET High Performance Computing Summer Program at FIU. In the Summer of 2006, two FMU students are participating in REU Sites. Brenton Williams is at the FIU REU Site performing research in Autonomic Computing. Ahamad Gallon will be participating at the Iowa State University REU Site.

### **Outreach Program to Schools**

We are developing an outreach program that will ultimately consist of both visits to FIU and/or FMU and a traveling 'show' that includes a presentation geared to the appropriate audience at schools. The presentation is followed by a hands-on demonstration of interesting database projects to which the students can relate, such as advanced 'virtual reality' demonstrations made possible by TerraFly.

### **Graduates**

Several FIU graduate students who had been supported by our previous MII award are now supported by this MII award. One of these students, Scott Graham, received his Ph.D. in computer science in the Spring of 2004; Dr. Graham is now the research coordinator for FIU's School of Computer Science. Another student, Xiangyu Ye, had received her MS in Computer Science under our previous MII award; she received her PhD in computer science under this MII's support in the Summer of 2004 and is now doing consulting work while working to establish her own business. Two other students, Dayanara Hernandez and Ian De Felipe, received master's degrees in computer science in the Spring of 2004; they are both pursuing their doctoral degrees at FIU. Three additional MII-supported students received their PhDs in computer science: Yanli Sun received his degree in the Summer of 2004; Mario Sanchez received his degree in the Spring of 2005 and is now a faculty member at Miami Dade College; Andriy Selivonenko received his degree in the Fall of 2005 and is now a Post Doctoral associate here at FIU. Bachelor degree recipients included Fabian Alcantara (FIU), Chemise Armstrong (FMU), Ian Cooper (FMU), Rafael Cuevas (FIU), Jason Good (FMU), Royel Haynes (FMU), Osigbemhe

Iyalomhe (FMU), Theirick Knight (FMU), Eric Kobrin (FIU), Fitzgerald Lawrence (FMU), Albert Marshal (FMU), Krista Merrill (FIU), Luis Pachas (FIU), Elvis Padron (FMU), Wayne Pinckney (FMU), Adelein Rodriguez (FIU), Yolanda Ruggiero-Cardoch (FIU), Rafael Santiago (FIU), Cedric Slacks (FMU), Chattam Smith III (FMU), Winston Smith II (FMU), Clevin Weekes (FMU), and Joseph Wilnigue (FMU). Many have accepted jobs in industry; Mr. Weekes will pursue his master's degree at Embry-Riddle Aeronautical University after he finishes a master's degree program at Virginia Tech, and Mr. Marshal, Mr. Padron, Mr. Good, and Mr. Wilnigue are pursuing their master's degrees at FIU.

## **Publications**

Sixty seven papers acknowledging our MII award have been published. Ten additional papers have been accepted for publication. Twenty six of these publications were co-authored by FIU and FMU faculty and students. A selection of these publications follows:

- Mario Sanchez, Naphtali Rische, Keith Morren, Wilnigue Joseph, "Towards the Storage of Geospatial Data with the Semantic Binary Model and Virtual Quadrilaterals", Proceedings of the 2nd International Conference on Cybernetics and Information Technologies, Systems and Applications, CITSA 2005, p. II-89.
- N. Rische, M. Gutierrez, A. Selivonenko, S. Graham, "TerraFly: A Tool for Visualizing and Dispensing Geospatial Data", Imaging Notes, Summer 2005, p. 22, vol. 20.
- O. Wolfson, B Xu, H. Yin, N. Rische., "Resource Discovery Using Spatio-temporal Information in Mobile Ad-Hoc Networks", Web and Wireless Geographical Information Systems, 5<sup>th</sup> International Workshop, W2GIS 2005, Lausanne, Switzerland, December 15-16, 2005. Springer Verlag Lecture Notes in Computer Science., p. 129, vol. 3833, (2005).
- Shu-Ching Chen, Mei-Ling Shyu, Chengcui Zhang, Min Chen, "A Multimodal Data Mining Framework for Soccer Goal Detection Based on Decision Tree Logic", International Journal of Computer Applications in Technology, p. , vol. , (2005). Accepted
- Naphtali Rische, Malek Adjouadi, Ouri Wolfson, Maxim Chekmasov, Dmitry Vasilevsky, Scott Graham, Dayanara Hernandez, "XML-Based Semantic Database Definition Language", Proceedings of the 7th International Conference on Enterprise Information Systems (ICEIS 2005), p. 197.
- Naphtali Rische, Yanli Sun, Maxim Chekmasov, Andriy Selivonenko, Scott Graham, "System Architecture for 3D TerraFly Online GIS", IEEE Sixth International Symposium on Multimedia Software Engineering, p. 273, (2004).

# MII: Hardware-Software Integration for the Design of Real-Time Prototypes Merging Assistive Technologies to Neuroscience

Award Number: 0426125

PIs: Malek Adjouadi, Armando Barreto, Prasanna Jayakar, Ana Pasztor, and Gustavo Roig  
Department of Electrical and Computer Engineering, Florida International University  
[adjouadi@fiu.edu](mailto:adjouadi@fiu.edu), [barreto@fiu.edu](mailto:barreto@fiu.edu), [Prasanna.jayakar@mch.com](mailto:Prasanna.jayakar@mch.com), [pasztora@cs.fiu.edu](mailto:pasztora@cs.fiu.edu),  
[roigg@fiu.edu](mailto:roigg@fiu.edu) [www.cate.fiu.edu](http://www.cate.fiu.edu)

## 1. ABSTRACT

Our vision for this NSF-MII project is to continue consolidating cross-disciplinary research and educational programs as a catalyst for our students to develop their creative thinking by bringing in synergy the fields of applied image and signal processing assimilating software development to hardware designs into the practical fields of neuroscience and human-computer interfaces (HCI). This research mission is supported through a modern infrastructure with parallel and distributed processing capabilities and with equipment to support assistive technology and neuroscience research. In the merging of these technologies, we see a productive ground for the development of new methodologies and scientific discovery that (1) will enhance the HCI research prospects in order to respond effectively to the call of the NSF on the issue of universal accessibility focusing on visual impairment and motor disability; and (2) will meet the impending needs in neuroscience as we elicit both the functional mapping of the brain, and the causality of key brain disorders towards neuro-rehabilitation, focusing on epilepsy and autism.

## 2. INTRODUCTION/BACKGROUND

FIU continues to enjoy major growth, enrolling close to 38,000 students making it the largest university in South Florida, and with almost 1,000 full-time faculty and 113,000 alumni to date. Sponsored research funding (grants and contracts) from external sources for the year 2004-2005 totaled nearly \$80 million. The University is ranked as a Research University in the High Research Activity category of the Carnegie Foundation's prestigious classification system. FIU has two campuses, the 344-acre University Park campus in western Miami-Dade County, and the 200-acre Biscayne Bay Campus in Northeast Miami-Dade County, and a 40-acre Engineering Center located one mile away from the University Park campus. Just recently this year (March 2006) the Florida Board of Governors voted to create a School of Medicine at Florida International University. The Center for Advanced Technology and Education (CATE) was inaugurated in the fall of 1993 with the support of the NSF-MII program and has evolved over the years into a modern infrastructure with strong links to local hospitals and industry and to several partnering academic institutions. Since its inception, the CATE center has supported an unprecedented number of students both graduates and undergraduates, resulting in a strong student pipeline that serves a large number of underrepresented groups and women in electrical and computer engineering. Our infrastructure is modern and has experienced growth in scope and in instrumentation through strong joint collaboration programs with local hospitals and industry partners. The CATE center continues to serve as a resource infrastructure for research and education in CISE-related disciplines as we strive to recruit and educate more students from underrepresented minority groups, women and persons with disabilities all the way to the Ph.D. level. The CATE center provides a multidisciplinary research and educational environment with strengths in the following areas:

- Image/Signal processing and computer vision
- Human-computer interfaces in the development of assistive technology research for visual impairment/blindness and motor disability
- Neuroscience applications
- Computational designs for real-time processing

Furthermore, with strong collaboration with our medical colleagues, through a joint Neuro-Engineering Program, our faculty and students are enjoying access to modern equipment to perform advanced research in the challenging field of neuroscience.

## 3. GOALS, OBJECTIVES, TARGETED ACTIVITIES

Our goals since the inception of the CATE Center in terms of human development take strength from the following:

- An infrastructure that supports critical research in CISE related disciplines

- Outreach programs that attract underrepresented students into graduate studies and retain them through Ph.D. studies.
- Opportunities for advanced research in CISE-related disciplines with critical applications to Medicine and Assistive Technology Deployment, which has led to joint programs, faculty lines, and major joint proposals.
- Applications and implementations of research findings in neuroscience and assistive technology weighed in terms of the societal impact they yield.

The main objectives of the CATE center are as follows:

- Create a unified infrastructure, fostering an environment that supports cross disciplinary initiatives and education, joint collaborations and shared use of modern equipment and facilities.
- Establish an integration of multimodal technologies, exploiting and complementing the strengths of each modality, to create a cohesive platform providing a holistic view of CISE-related applications in computing, assistive technology and medicine.
- Provide researchers access to an entire gamut of technologies including 3Tmagnetic resonance with functional and diffusion tensor images, nuclear medicine studies, ultrasonography, electrophysiology including direct recordings from cerebral cortex, and optical and fluorescent spectroscopy.
- Ensure that these integrative efforts will benefit from computational paradigms that confront the heavy computational requirements as we address key issues such as spatial co-registration, 3-D reconstructions, inverse solutions, pertinent design problems rising from integrating diverse sensing modalities, and successfully overcoming challenges of high speed processing.
- Apply the knowledge gained in brain studies to begin understanding the functional brain mappings of subjects with epilepsy and autism, as we elicit new insights into the brain structure, its functional correlations and dynamics, both in its normal physiologic state and under specific pathologic conditions.

#### 4. INFRASTRUCTURE AND FACILITIES

Major Equipment currently available at the CATE Center and its affiliate laboratories includes:

- Center for advanced Technology and Education- (FIU) <http://www.cate.fiu.edu/>
  - Computer network with parallel and distributed computing capabilities
  - High resolution Electrical Source Imaging system with 256 electrodes (ESI-256)
  - A MagLink for integrating EEG recording under a Magnet for MRI acquisition
  - Eye-Gaze tracking systems (1 fixed, 1 portable) as HCIs for Motor Disability.
  - An automatic book reader computer interface as an assistive technology tool for blind persons
  - Distributed system for real-time processing
  - Mural display for visualization
  - Computing facility housing a Multisite fMRI database for pediatric epilepsy
  - Near-Infrared Camera System
- The Digital Signal Processing Laboratory –FIU <http://dsplab.eng.fiu.edu>
  - 3D sound based assistive devices for navigation
  - Human/computer interface for visual disability
  - EGT/EMG HCI for motor disability
  - Wavefront Analyzer
  - Affective computing system
- The Brain Institute – Miami Children’s Hospital (MCH) <http://www.mch.com>
  - High resolution EEG and EP, intracranial recording directly from the brain.
  - Optical 32 channel near infra-red spectroscopy system.
  - Functional MRI and Diffusion tensor imaging for tractography
  - 3T MRI [planned for installation].
  - Single photon emission computed tomography.
  - Optical probe for spectroscopy during surgery
  - Near-infrared spectroscopy

## 5. PROJECT DESCRIPTION

This project implements methodologies and creates initiatives designed to: (a) engage students into full participation in science, engineering, and technologically-driven disciplines, with a student pipeline designed to take them all the way to the Ph.D. level; (b) make the most of comprehensive outreach and recruitment programs to address and redress the attraction and retention of women, minority students, and persons with disabilities into graduate studies.

The overall structure and organization of this NSF-MII proposal is envisioned with the intent to:

- Establish a modern infrastructure with seamless integration to existing equipment and with valuable access to additional modern equipment and facilities of research partners extremely difficult to acquire otherwise and too expensive to budget through federal grants.
- Provide our students with an environment that is conducive to: (1) cross-training among disciplines, (2) experimental evaluations, and (3) feasibility studies using real-world settings from which prototype designs are moved to the realm of practicality.
- Develop new curriculum in relation to the proposed research themes, with engineers, computer scientists, neuroscientists, and radiologists working together for an integrated approach to research, teaching and training of our students.
- Strengthen our outreach programs for enhanced potential at recruiting outstanding graduate students from a diverse population, emphasizing the merits of the research.
- Create an environment where the merits of research are introduced early in the academic careers of students, with the added potential to generate creative methodologies worthy of the new-generation of Ph.D. graduates.
- Promote enrichment programs and resources that extend the notion of “access”; resulting in novel assistive technologies and tools to alleviate or remove the constraints imposed by the nature of the disability. The notion of usability is thus broadened by embracing the belief that designing for the disabled population directly benefits the population at large, since it illuminates usability issues that may not have been otherwise apparent.
- Establish a dissemination process that extends to other institutions of higher learning, including community colleges and K-12 institutions, enhancing as a consequence the interchange of ideas, and increasing awareness and recognition of findings and consequential contributions.

## 6. INDICATORS OF SUCCESS AND MAJOR ACCOMPLISHMENTS

- With the help of the leadership at FIU and Miami Children’s Hospital (MCH), Dr. P. Jayakar at MCH and Dr. Adjouadi spearheaded the efforts in creating the joint Neuro-Engineering Program which includes a faculty line (½ MCH ½ FIU) for the 2003-2004 academic year and potentially another line (full line MCH) for the 2006-2007 academic year, and recently a multimillion dollar 3-Tesla magnet is to be installed to advance our research goals in neuroscience.
- Secured the second largest College of Engineering endowment by the Ware Foundation (\$1.42M) (PI: M. Adjouadi, Co-PIs: P. Jayakar, and A. Barreto), celebrated both at the Torch Ceremony and at the President’s House here at FIU.
- 1 US patent granted and 2 patent applications pending.
- Produced as author or Co-author 63 Journal publications (including 3 book chapters and 7 book monographs), 85 conference publications with 19 of them in workshops, posters, and invited talks.
- FIU is now a member of the BPC-A: Computing Alliance for Hispanic-Serving Institutions, [PI: Dr. Ann Gates - UTEP], Co-PIs: M. Adjouadi - FIU, Mohsen Beheshti CSUDH, Desh Ranjan-NMSU, and Nestor Rodriguez – UPRM in order to attract and graduate more students and potentially more faculty in the computing disciplines.
- The CATE center was directly involved in securing more than 14 million dollars in federal funding, strengthening further our potential for research and for attracting the next generation of Ph.D. students in CISE related disciplines
- The CATE center has secured 8 prestigious graduate fellowships: 7 NSF graduate fellowships for 6 female students and 1 male), and 1 Intel graduate fellowship for a male student.
- Established through federal funding 5 ongoing educational/ research labs (3 within the College of Engineering, 1 at Miami Children’s Hospital, and one with the Multicultural Services Office in the Graham Center on main Campus).
- Graduated 3 female Hispanic Ph.D.s, with one more defending June 20<sup>th</sup>, 2006, and another defending in the fall of 2006. The student pipeline continues to strengthen in order to maintain this rate of success or even increase it in the number of future Ph.D. graduates, or MS graduate that pursue their Ph.D.s elsewhere in the Nation.

## 7. PUBLICATIONS/PATENTS

### 7.1 Publications

The following accomplishments are made for the period of May 2005 – May 2006. We had 1 book chapter, 20 journal publications and 30 conference articles published in proceedings, all acknowledging this NSF award CNS-0426125. Select Journal publications are listed below:

1. M. Cabrerizo, M. Adjouadi, M. Ayala, Chapter 5 titled: An Application of Eigensystem and Frequency Analysis in Brain Functional Mapping, Progress in Brain Mapping Research, NOVA Science, ISBN 1-59454-580-4, 2006.
2. M. Alonso, A. Barreto, J.G. Cremades, and M. Adjouadi, "Image Pre-compensation to facilitate computer access for users with refractive errors", **Behaviour & Information Technology**, May-June 2005, vol. 24, no. 3, pp. 161-173.
3. M. C. Zabawa, M. Adjouadi, and N. Rishe, "SystemC Co-Design for Image Compression: Fast Discrete Cosine Transform using Distributed Arithmetic, **WSEAS Transactions on Computers**, Vol. 4, No. 6, pp. 477-484, June 2005.
4. F. Gui, N. Zong, M. Adjouadi, "Dynamic Neural Network Based Algorithm for Context Awareness in Mobile Computing", **WSEAS Trans. on Communications**, Vol. 4 (8), pp. 629-636, August 2005.
5. M. Adjouadi, M. Cabrerizo, M. Ayala, and N. Mirkovic "Seizing lesions in 3-D", **IEEE Potentials**, Vol. 24, Issue 5, pp. 11-17, December 2005.
6. M. Rossman, M. Adjouadi, M. Ayala and I.Yaylali, "An Interactive Interface for Seizure Focus Localization Using SPECT Image Analysis", **Computers in Biology and Medicine**. Vol. 36(1), pp. 70-88 January 2006.
7. Li, C. and Barreto A., "Profile-Based 3-D Face Registration and Recognition". **Lecture Notes in Computer Science**, (Springer-Verlag), LNCS 3506, pp. 484-494, May 2005.
8. Zhai, J., and Barreto, A., "Stress Detection in Computer Users Through Noninvasive Monitoring of Physiological Signals", **Biomedical Science Instrumentation**, vol. 42, pp. 495-500, 2006.
9. Li, C., Barreto, A., Chin C. and Zhai J., "Biometric Identification Using 3D Face Scans," **Biomedical Science Instrumentation**, vol. 42, pp. 320-325, 2006.
10. Faller, K. J., and Barreto, A., "Simulation and Real-Time Implementation for Teaching 3D Sound". **Computers in Education Journal**, Vol. XVI, No. 2, April – June 2006, pp. 36-43.

### 7.2 Patent Applications

1. Artificial Neural Network Design and Evaluation Tool, US Patent Application Serial No. 11/203,510, October 12 2005. (with Dr. Melvin Ayala)
2. Configurable, Multimodal Human-Computer Interface System and Method, US Patent Application Serial No. 11/176,812, July 2005. (With Anaelis Sesin, Melvin Ayala, and Mercedes Cabrerizo).

## 8. OUTREACH PROGRAMS AND SUPPORTED STUDENTS

### 8.1 Outreach and Recruitment Activities 2005- 2006

Our outreach and recruitment efforts are led by Dr. Gustavo Roig, (Co-PI) and Director of the Center for Diversity in Engineering, and Ms. Stephanie Strange, Director of recruitment and retention. Their coordinated efforts with the participation of the PIs and students of the CATE center have yielded:

The Engineering and Computing Gala 2006: The FIU College of Engineering and Computing held the fourth annual Engineering and Computing Gala. The Gala is the colleges' premier outreach event. The event is the interface of the college, the university, high schools, middle schools, elementary schools and the community at large to meet, learn and celebrate engineering education and engineering as a profession. Each year the college works with the local school districts to coordinate field trips for the day, to visit the college. The students create and implement contests, events and "hands on" engineering projects. Also, the research and learning labs are made available for tours. The activities, resources and the pedagogic methods used to promote interest in math, science and engineering education, are available to educators throughout the event. This year the Engineering and Computing Gala attendance exceeded 1000 students. The students came from a variety of schools reflecting Miami's diverse population demographic. Miami is made up largely of an immigrant population with many "first time in college" students. Most of the students visiting are from traditionally underrepresented minorities in science and engineering education.

Miami Dade Schools Math and Science Division/ FIU College of Engineering and Computing Partnership: The FIU College of Engineering and Computing works extensively as a partner with the Miami Dade Schools Math and

Science Division. Through the regional section of the national program SECME (Science Engineering Math and Communication Enhancement Program) faculty, students and staff collaborate on projects, provide resources and participate in the many events, contests and workshops that are a part of the SECME program. The winners of the regional SECME Olympics are then able to compete on a national level. The Miami Dade Schools Advanced Academic Internship program has been a part of the college for two years. Each year between 6-10 Miami Dade high school seniors are placed with College of Engineering and Computing professors to assist in research and be mentored throughout the year. Summer programs such as a dual enrollment summer course jointly sponsored by the FIU College of Engineering and Computing and the Miami Dade Schools have been successful in the college.

School Visits (Elementary, Middle, High Schools): On an average; faculty, engineering and computer science students and administrators regularly coordinate visits to schools in Miami Dade and Broward counties each semester to promote engineering and computer science. Presentations vary with each department and presenters introduce a wide array of educational/research topics. The following is listing of the schools visited: Broward County High Schools: Stranahan High School, Coral Glades High School, Taravella High School, Plantation High School. Miami Dade County High Schools: Coral Reef; Coral Park; and Miami Springs, Michael Krop; Booker T. Washington; Barbara Goleman; and Miami Senior. Miami Dade Middle Schools: Highland Oaks; Miami Dade Elementary Schools: Campbell Drive

## 8.2 Supported Students

The NSF-CATE Center supported an unprecedented number of students to date:

- 89 graduate students (40 females and 49 males): 21 Ph.D.s with 12 females and 9 males; and 68 MS students with 28 Females and 40 Males
- 77 undergraduate students (37 females and 40 males)

Degrees Granted to Students supported by the CATE Center include:

- 11 obtained their Ph.D. degree (9 here at FIU: 6 females and 3 males), and 2 went on to get their Ph.D. degrees elsewhere after obtaining their MS degree here at FIU (1 female at MIT-Harvard, 1 male at UF);
- 56 obtained their MS degree (23 females and 33 males), and • 72 obtained their BS students (34 females and 38 males), with
- 8 NSF graduate fellowships (6 females and 2 males) and • 1 Intel fellowship (1 male student).

The CATE center has also provided financial support for 3 Postdoctoral students, 1 Research Scientist from Miami Children's Hospital, and 3 consultants.

## 9. FUTURE AND STRATEGIC DIRECTIONS

Our strategic research directions under this infrastructure award are structured to (1) effect long-range improvement in *Human Computer Interface* (HCI) research, with the objectives to impact significantly the theoretical and technological frontiers of the assistive technology field, which is destined to bridge the world of the able-bodied to persons with disabilities focusing on visual impairment and motor disability, as we seek universal access; (2) formulate new applications of neuroscience as we merge imaging and signal processing techniques to meet impending needs for new developments in bio-signal processing and neuro-rehabilitation as we elicit both the functional mapping of the brain, and the causality of key brain dysfunctions. In these research endeavors, engineering and scientific aspects relating to hardware/software integration, information processing, mathematical modeling, system development and design, and computational requirements are brought to bear in the implementation phases. The testing phases are supported through feasibility studies, such that design prototypes will meet the challenges imposed by real-world problems. Thus the thematic basis of this research is one that coalesces the research strength of key disciplines while seeking system integration, hardware-software assimilation, and where new theoretical developments are weighed in terms of the societal impact they yield, and with complementary technologies brought together such that their contributions are consolidated and their limitations overcome.

# CRI: Infrastructure for Multi-Agent Decision-Making Research

Proposal #CNS-0453923, Year One

Harvard University

Barbara J. Grosz, Avrom Pfeffer, and Stuart M. Shieber

[grosz@eecs.harvard.edu](mailto:grosz@eecs.harvard.edu), [avi@eecs.harvard.edu](mailto:avi@eecs.harvard.edu), [shieber@deas.harvard.edu](mailto:shieber@deas.harvard.edu)

<http://www.eecs.harvard.edu/ai/ct>

## Abstract

Widespread use of the Internet has fundamentally changed the ways in which we use computers, not only as individuals, but also within organizations. Settings in which many people and many computer systems work together, despite being distributed both geographically and in time, are increasingly the norm. These increasingly prevalent, heterogeneous group activities of computer systems and people frequently require decision-making on the part of autonomous-agent systems or the support of decision-making by people. A wide variety of issues arise in the decision making that accompanies such activities. Grosz and Kraus (Grosz et al. 2004) developed the game Colored Trails (CT) to support their investigation of multi-agent decision-making in task settings. CT provides the basis for development of a testbed that would support a wide-range of investigations of human decision-making and comparisons among computational strategies. The infrastructure grant supports making the CT testbed available for wide and open distribution to the scientific community.

## Introduction

Our objective is to develop and distribute a research testbed for use as a community resource by a broad range of researchers for the study, modeling, and development of decision-making strategies appropriate for multi-agent task settings of various sorts. This testbed will facilitate the analysis of decision-making strategies in the contexts of group tasks and activities carried out not only by groups comprising solely human agents (in settings similar to, but typically more complex than, those of games in economics like the ultimatum game) or solely computer agents (for instance, software information agents, online auctions), but also by heterogeneous groups of human and automated agents (for example, for systems administration, hospital patient scheduling). It will enable the development and testing of decision-making strategies for automated agents that must operate in groups that include people.

The testbed will be based on a conceptually simple but highly expressive game called *Colored Trails* (CT), in which agents, working individually or in teams (depending on parameter settings), make decisions about how to deploy their individually held resources. As described in detail later, Colored Trails has been designed as a parameterized family of games, in which *decision-making* and *limited resources* play intrinsic roles. CT is the first testbed to be designed to investigate decision-making in heterogeneous groups of people and computer systems. Already, a prototype implementation of CT has been used in several different experiments by researchers at two institutions. For instance, through experiments with a two-player version of CT, we have shown that it is possible to model human decision-making strategies with sufficient accuracy that automated agents using these models outperform agents that act in the game-theoretically optimal manner (Gal et al. 2004). We also have initial results with a four-player version showing differences in the ways people play and they ways they design computer agents to play (Grosz et al. 2004). Finally, we have shown in recent experiments that people play qualitatively different strategies when playing CT in contrast to a game-theoretically equivalent game that

does not display the CT task scenario. In particular, the CT user interface leads to much less strategic play, showing that the game is a good proxy for human interactions in task-oriented domains.

## **Background**

Widespread use of the Internet has fundamentally changed the ways in which we use computers, not only as individuals, but also within organizations. Settings in which many people and many computer systems work together, despite being distributed both geographically and in time, are increasingly the norm. Examples of group activities in which computer systems participate, whether as autonomous agents or as proxies for individual people, include online auctions, hospital care-delivery systems, emergency response systems, military systems, and systems administration groups.

These increasingly prevalent, heterogeneous group activities of computer systems and people—whether competitive, cooperative, or collaborative—frequently require decision-making on the part of autonomous-agent systems or the support of decision-making by people. A wide variety of issues arise in the decision making that accompanies such activities. How do groups decide who should perform which tasks in service of which goals? How are coalitions formed? How do agents coordinate their behavior? How do agents select an appropriate recipe (method) for achieving a goal from among the set that might be applicable? How are tasks allocated among agents capable of performing them? How do agents reconcile their current intentions to perform actions in service of group goals with new opportunities that arise? How do agents avoid conflicts in their plans for achieving goals? When and how do altruistic or cooperative behaviors arise? What incentives can be provided to encourage cooperation, helpfulness, or other group-benefiting behaviors when agents are designed independently or serve different organizations?

Grosz and Kraus (Grosz et al. 2004) developed the game Colored Trails to support their investigation of multi-agent decision-making in task settings. CT enables investigating the key interactions among *goals* (individual and group), *tasks* required to accomplish those goals, and *resources* needed to perform the tasks. It allows the modeling of these phenomena and exploration of their interactions and the ramifications of different decision-making strategies. CT provides the basis for development of a testbed that would support a wide-range of investigations of human decision-making and comparisons among computational strategies. In particular, it was designed to enable human decision-making to be studied both in groups comprising only people and in heterogeneous groups of people and computer systems, as well as for computational strategies to be studied both in settings where computational agents interact only with other such agents and in heterogeneous settings.

## **Goals, Objectives, Targeted Activities**

We aim to develop the CT testbed for *wide and open distribution* to the scholarly community. CT is unique in its heterogeneity, including not only its being suitable for play by people and computers and mixed groups, but also its ability to model competition, cooperation, and collaboration by essentially self-interested agents. Because of the generality of the game and its broad parameterizability, there has already been great interest in using CT for a variety of research projects. Such activities have been hampered by the unavailability of a robust, distributable version of the system. The present effort will develop the infrastructure needed to eliminate this obstacle. Indeed, it was in large part on the basis of repeated requests for a distributable, deployable CT system from a broad range of researchers from computer science, artificial intelligence, economics, and psychology that the present project was envisioned.

## Infrastructure and Facilities

The software infrastructure that we are developing, like the prototype system that we have already used to perform initial experiments, is a client-server system that allows multiple agents, both computer and human, to interact to play any of a wide variety of negotiation games selected from the parameterized family of games that we collectively refer to as “CT”. Experiments performed to-date have been on-site, with each human subject assigned to a computer console on a local-area network. As we detail below, we are developing the software infrastructure to allow the system to scale to wide-area networks so that experiment subjects can participate off-site, over a wide-area network. For example, subjects may participate from home via a standard web browser, or even from a variety of mobile devices.

## Project Description

*Structure of the Game* Colored Trails is played by two or more players on a rectangular board of colored squares. The rules are simple: Each player is given a starting position, a goal position on the board, and a set of chips in colors taken from the same palette as the squares. Players may advance toward their goals by moving to an adjacent board square. Such a move is allowed only if the player has a chip of the same color as the square, and the player must turn in the chip to carry out the move. Players may negotiate with their peers to exchange chips. Communication is controlled; players use a fixed but expressive messaging protocol.

The scoring function, which determines the payoff to the individual players, is a parameter of CT game instances and may depend on a variety of factors. At its simplest, it would consist of a weighted sum of such components as: whether the individual reached the goal, the final distance of the agent from the goal, the final number of chips the agent held. The scoring function may be varied to reflect different possible social policies and utility trade-offs, establishing a context in which to investigate the effects of different decision-making mechanisms. For example, it may include a component determined by other players' goal achievement, and by varying the relative weights of individual and group good in the scoring function it is possible to make collaborative behavior become more or less beneficial.

Despite the simplicity of the rules, play of CT is able to model a broad range of aspects of task situations in which a group of agents perform actions; it allows for scenarios in which agents act as individuals or in teams or both. Traversing a path through the board corresponds to performing a complex task the constituents of which are the individual tasks represented by each square. Different colors represent different types of tasks. The existence of multiple paths to a goal corresponds to the availability of different “recipes” (Grosz & Kraus 1996) or methods for achieving goals. The possession of a chip of a particular color corresponds to having the skills and resources needed for a task, and being able to deploy them at the appropriate time. Not all players get chips of all colors, just as agents have different capabilities, availability, and resources. The exchange of chips corresponds to agents producing the intended effects of actions for each other; in different domain settings an exchange could be taken to correspond to agents providing resources to each other, doing tasks for them, or enabling them in some other way to do tasks they otherwise could not do.

The game environment may be set to model different knowledge conditions as well. For example, varying the amount of the board an agent can “see” corresponds to varying information about task constituents or resource requirements, whereas varying the information players have about each other's chips corresponds to varying information agents have about the capabilities of others. Various requirements on player objectives, goal squares, and paths correspond to different types of group activities and collaborative tasks. To distinguish cooperative settings from those in which agents act independently, the scoring function may have a significant factor in which an agent's reward depends on others' performance.

*Testbed Development* We plan to enhance the game infrastructure in two stages. The first stage will focus on enhancing both the flexibility and the robustness of game play, including the following changes:

- We have restructured and re-implemented large portions of the code to bring them up to reasonable software engineering practice.
- To support repeated games, we have begun to add persistent storage of player and game information.
- In experiments with the current GUI, all players have had full board and chip information. We plan to complete the support of partial information in both the server and the clients.
- We have developed a rudimentary text-based meta-server for configuring, initiating, and monitoring games. We plan to considerably enhance these control capabilities to provide a graphical control panel through which one can centrally configure multiple simultaneous games for their game parameters and users. It will enable starting new games, monitoring their status, and terminating them.
- We will develop a non-player version of the GUI client for use in monitoring games in real time.
- The server's robustness and performance will be optimized.
- We plan to improve the GUI client's ease-of-use, especially for multiple-player games, which may involve the display of considerable amounts of information. We have begun to work on more compact ways to organize and filter this information.
- We plan to further generalize the programmable client skeleton, so it can be used not only as the basis for programmable clients, but also for the GUI client. This will allow an even more robust framework for programming clients with diverse interfaces.

We have opened the system's source code for non-commercial purposes through an initial alpha release. We will update this release as milestones are reached. This will allow researchers to download the code and use it as-is, or to tailor it for their own research goals.

In the second stage of development, we plan to open the system for use as a platform for agent-writing competitions, in the spirit of the Trading Agent Competition (TAC, <http://www.sics.se/tac/>). The basic setting would be a centralized meta-server controlling a collection of local servers running games with potentially remotely distributed clients. This will require the following changes:

- The server's communication and synchronization capabilities will be scaled to allow game play over wide area networks. In particular, the server will have to graciously handle network delays and failures.
- We will augment the meta-server with an online database of game configurations, game status, and results. This database will reflect an up-to-date snapshot of the games being carried out at any particular time.
- We will open the meta-server for Internet access, allowing easy browser-based configuration, monitoring, and control.
- We will add security mechanisms to the meta-server to allow different access rights for setting up games and game parameters, initializing, joining, and monitoring games, and viewing accumulated results.
- While each particular game may be limited to a handful of players, the meta-server should be able to robustly handle a load of multiple simultaneous games, and additional user requests. We will scale the meta-server to support the expected number of players.

### **Indicators of Success and Major Accomplishments**

An initial prototype of the Colored Trails system has already been used to obtain our first machine learning results (Gal et al. 2004). Experiments with alternative user interfaces have been conducted, as well as initial work on repeated games. Experiments are planned for CT games that introduce player uncertainty. Development of an entirely new CT code-base is well under way; the alpha release of our system is available for download at: <http://www.eecs.harvard.edu/ai/ct>.

## Publications

Research using Colored Trails has been published in Grosz et al. 2004, Gal et al. 2004, Marzo et al. 2004, Gal et al. 2005, Talman et al. 2005, and Gal 2006.

## Future and Strategic Directions

The next year of our efforts will concentrate on the development of the Colored Trails testbed towards the end of supporting our research goals. Development work falls into three broad categories: System scalability, system robustness, and system functionality. Specific efforts with these categories may closely interact. Scaling the system to support large numbers of players includes providing web-based interfaces and the construction of a meta-server to coordinate multiple experiments. Successful scaling demands that the system be built on top of a robust network framework; the system will be built to gracefully handle network delays and failures. The system will gain new functionality as required by our planned experiments. Such functionality includes the implementation of game-state “views” that allow the experimenter to specify what information is known to each player. Other functionality concerns experiment support code that is responsible for selecting which human subjects will play each other in each CT game. These and other features will allow us to easily perform experiments on a variety of game scenarios. These scenarios include games of partial information, repeated games, and games that involve multiple players; further, the expanded feature set of the CT platform will allow for easy replacement of user interface components.

## References

- Y. Gal. Reasoning about Rationality and Beliefs. PhD thesis, Harvard University, Cambridge, MA, 2006.
- Y. Gal, A. Pfeffer, F. Marzo, and Barbara J. Grosz. Learning social preferences in games. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-2004)*, pages 226–231, 2004.
- Y. Gal, B. J. Grosz, A. Pfeffer and S. Shieber. Colored Trails: A Formalism for Investigating Decision-making in Strategic Environments. *IJCAI Workshop on Reasoning, Representation, and Learning in Computer Games*, Edinburgh, Scotland, August 2005.
- B. J. Grosz and S. Kraus. Collaborative Plans for Complex Group Action. *Artificial Intelligence*, 86(2):269–357, 1996.
- B. J. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Havlin. The influence of social dependencies on decision-making: Initial investigations with a new game. In *Proceedings of the Third International Joint Conference on Multi-Agent Systems (AAMAS-2004)*, pages 782–789, 2004.
- F. Marzo, Y. Gal, B. J. Grosz and A. Pfeffer. Social Preferences in Relational Contexts. *IV Conference in Collective Intentionality*. Siena, Italy, October 2004.
- S. Talman, Y. Gal, M. Hadad and S. Kraus. Adapting to Agents' Personalities in Negotiation. *Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Utrecht, the Netherlands, July 2005.

Title: Minority Institutions Infrastructure -

## **MII: RESEARCH INFRASTRUCTURE IMPROVEMENT**

**Proposal # EIA-0324818**

**Project Year: 3**

**Department of Systems and Computer Science  
College of Engineering, Architecture and Computer Sciences  
Howard University  
2300 Sixth St. NW  
Washington, DC 20059**

**Principal Investigator: Ronald J. Leach**

**Co-Principal Investigators:**

**Harry N. Keeling**

**John Trimble**

**rjl@scs.howard.edu**

**hkeeling@howard.edu**

**trimble@scs.howard.edu**

**URL(s):**

**www.howard.edu**

**www.howard.edu/ceacs/Departments/CompSci/index.html**

**www.scs.howard.edu**

**www.scs.howard.edu/~rjl/Welcome.html**

**http://hufast.howard.edu/**

**http://hufast.howard.edu/FreedmenKnowledgeClient**

**http://groups.yahoo.com/group/syscomphd (work in progress)**

**http://pluto.scs.howard.edu/~blegand/**

**http://www.husec.net/**

### **1. Abstract**

The continued abysmal production of African-American Ph.D. graduates in Computer Science is a serious matter for national competitiveness. Howard University proposed to develop the infrastructure sufficient to offer a Ph.D. program in Computer Science and to take the lead in doubling the number of such Ph.D. graduates within a few years.

### **2. Introduction/Background**

The critical problem that Howard University wished to address in this proposal to the National Science Foundation was the critical shortage of African-American Ph.D.-level Computer Scientists. Howard University is uniquely qualified to address this area of critical national need because of the overall research environment at the university and the research growth in the Department of Systems and Computer Science.

We must address this problem at a time of severe decreases in the number of undergraduate students majoring in Computer Science nationally. Microsoft's Bill Gates is just one of many in the industry who have raised this issue. In his visit to campus as part of his six-campus visits, he emphasized

the importance of improving the appeal of the Computer Science major. (Howard University sent more students to Microsoft in 2003-4 than any other university in the country.)

### **3. Goals, Objectives, Targeted Activities**

The primary goal of this project is to improve the Computer Science research infrastructure of Howard University in a level that is sufficient to be able to offer a competitive Ph.D. in Computer Science. A secondary goal was to become a national model for the development of new Ph.D. programs at small universities.

### **4. Infrastructure and Facilities**

The Department of Systems and Computer Science has teaching and research facilities in two adjacent buildings on the main Howard University campus: Downing Hall and the Mackey Building. The facilities are largely separated by function, with the systems serving primarily undergraduates lodged in Downing Hall and the facilities primarily used by graduate students in the Mackey Building.

The department has three dedicated laboratories in Downing Hall used for undergraduate instruction: a 32-seat Microsoft Windows-based PC laboratory used for basic undergraduate instruction; a 16-seat Linux laboratory used for basic undergraduate instruction; and a 17-seat advanced laboratory running a variety of operating systems, including both servers and clients, including Windows .NET server, other Windows systems, and Linux. A departmentally controlled, but shared, multimedia laboratory is used for interdisciplinary instruction in multimedia and has just begun to be used for game development. Several College-wide computing facilities are also housed in Downing Hall.

The facilities for graduate students in the Mackey building include a graduate student research facility and office area, a distributed systems laboratory, and a telecommunications laboratory.

The large graduate student research facility and office is designed for multiple uses. It includes a large open area for supported graduate students, with smaller, more private offices for students working on specific funded research projects.

The distributed systems laboratory includes a collection of Sun Solaris workstations and a server; these systems are being converted to Linux this summer. There are two small clusters available; one for research, and one for systems development.

The department's telecommunications laboratory contains state-of-the-art systems for network modeling and simulation. Additional state-of-the-art equipment for research and education in wireless computing has been obtained within the last month from an industry source and is waiting for installation.

### **5. Project Description**

The department had two primary goals: develop the infrastructure for a new Ph.D. program in Computer Science and to become a model for other universities, especially HBCUs and MSIs, to develop their own new Ph.D. programs.

As part of the program development process, we have examined a large collection of published material from the Computing Research Association, NSF, and other sources that are not university-specific. In addition, we have examined the publications of several universities with relatively new or small Ph.D. programs. Some of the data is available, at least internally, from the Yahoo group given above.

Much of the data we collected about specific Computer Science Ph.D. programs can be useful in determining appropriate administrative processes. Of course, information about *process* is more useful than demographic data about very large programs in well-established institutions. Comparing a

department's research goals to, say, the Media Laboratory at MIT, which has approximately \$40 million dollars in its annual budget, is not likely to be helpful when setting program goals for smaller programs. For programs at smaller institutions, the implementation of a new program must take into account that program *need* is almost ahead of new program resources.

Based on this research and on our own experience, we have developed a preliminary paper with the following major section areas:

- The early stages of program creation:
- Academic Issues
- Political Issues
- Administrative Issues
- Procedural issues
- Funding issues
- Phasing in the programs

This work is being refined, and it is premature to publish such a paper, or series of papers, until the department's Ph.D. program is given final approval by the Howard University Board of Trustees.

## **6. Indicators of Success and Major Accomplishments**

We feel that in this difficult budgetary climate for higher education, getting approval for any new program that requires additional resources is extremely difficult. The department's success in this area is reported in the next few paragraphs.

There are over 216 Computer Science Ph.D. programs in the United States. While the need for more African-American Ph.D. Computer Scientists is critical, there is no need for a low-quality program that produces non-competitive Ph.D.s. Howard University prides itself on its academic ratings – such as Carnegie Level I Research Institution, Carnegie Doctoral Extensive, etc. and the less academically validated U.S. News and World Report “Best Buy” – and, in the current climate for higher education funding, is reluctant to support programs that are not likely to be of high quality or are not in the current strategic plan for the university.

As part of its strategic plan, known as the Strategic Framework for Action, and its successor plan, the Strategic Framework for Action II, Howard University has determined that the only new graduate programs it will support will be interdisciplinary ones. There is a single university-wide exception to this – the proposed Ph.D. program in Computer Science. This exception is due to the university's understanding of the importance of Computer Science to its research goals and aspirations.

A proposal to offer a Ph.D. in Computer Science had received positive evaluations at several levels prior to the awarding of the five-year infrastructure award. The proposal had been approved by the Deans of the College of Engineering, Architecture and Computer Sciences and the Graduate School. It was approved by the Executive Committee of the Graduate School and then by the entire Graduate Faculty of the University, after which it was transmitted to the Office of the Provost in the summer of 2003.

A new Provost was appointed in early September, 2003, first as an interim and then permanently. He was concerned about three matters relevant to the department's proposal to offer the Ph.D. in Computer Science: the need for additional reviews done under his guidance, the relative skewing of the department's faculty toward the junior ranks caused him concern about the availability of senior faculty within the department, and the balancing of the need for resources against the need to move forward.

In a meeting with all concerned parties, the Provost asked for an additional internal review of the proposal by the executive Committee of the Graduate School. This review was highly satisfactory. An external review was also strongly supportive and hence all concerns about quality issues have been resolved.

The Provost noted in a follow-up meeting in December 2005 that the issue of what had been a potential weakness – the small percentage of full professors – had ceased to be a matter of concern because three of the department’s Assistant Professors were promoted to Associate Professor with tenure. Their professional growth in research, together with that of their colleagues, demonstrated the ability to take on Ph.D. students and train them properly.

The final issue, resources, had been discussed at many meetings about the progress of the proposed Ph.D. program’s approval. The Provost had been assured that the proposed program would focus on highly limited areas that were also areas of departmental research strength – distributed systems and software engineering. Still, he took the position that the program needed resources to succeed and that such resources must be provided at Dean’s level. In short, he forced two Deans to reallocate some existing funds, which they agreed to do.

The agreement was that the department needed additional resources in the areas of faculty, technical staff, administrative staff, and student support. Dean Johnson has agreed to provide the following support from the budget from the College of Engineering, Architecture and Computer Sciences: an additional tenure-track faculty position beginning in the 2007-8 academic year, and to provide an administrative position when a retirement occurs in the near future.

Dean Taylor has agreed to provide a technical staff position beginning in the 2006-7 academic year by transferring the position of the person the department strongly wishes to have on its payroll from the Graduate School’s current budget. Full funding for two new Ph.D. students will be made available from Howard University funds as soon as possible after the proposal is approved in the near future.

With these commitments, we are confident that the university will approve a Ph.D. program in Computer Science very soon. (Dean Johnson stated to faculty candidates earlier this spring that he believed we would have approval well within a year when we were interviewing for new faculty positions.)

On another note, the department has successfully applied to have a chapter of Upsilon Pi Epsilon created on campus and this honorary society has been well received by graduate and undergraduate student inductees. This honor reflects the quality that the department’s programs demonstrate to national evaluators.

## **7. Selected Publications/Patents**

A representative list of publications by the Principal Investigator in the most recent project year is given below. In addition to the publications of the Principal Investigator and co-PIs, many department faculty members received some support from NSF as part of this infrastructure improvement project, in the form of support for their graduate or undergraduate students; travel; or software or equipment purchases to aid their research. Publications from the co-PIs, and any additional department member’s publications funded at least partially as a result of this grant will be incorporated into the annual progress report for year 3.

- a. Leach, R., and Harry N. Keeling, “Knowledge Engineering in a Hybrid Systems Engineering Curriculum,” *Frontiers in Education Conference, IEEE, Indianapolis, Indiana, October 19-22, 2005*, pp. F3F13-F3F14.
- b. Leach, Ronald J., and Caprice L. Ayers, “The Psychology of Invention in the Computer Field,” *Psychology of Programming Interest Group Workshop, Brighton, England, June 27-July 1, 2005*, pp. 131-145.
- c. Leach, Ronald J., “A Model Relating Software Quality and Reuse,” *Crosstalk, April, 2006*, pp. 27-28.
- d. Leach, Ronald J., “Can this COTS-Based System be Saved?” *Vol. 9, No. 4, pp. 38-44, PC/104 Embedded Solutions, Fall, 2005*. This article was designated as a “Software Feature.” Also appeared in *Military Embedded Systems, Fall 2005*.

## **8. Outreach Programs and Supported Students**

The department's planned outreach projects as part of this infrastructure project were highly tied to the approval of the proposed Ph.D. program. As such, the department has limited its outreach to recruiting trips, to the National Society of Black Engineers conference, and to the annual ADMI conference. An entry has been placed in Peterson's Guide and the success of this (expensive) entry is currently being evaluated.

The department has followed the approach of using student support dollars to support the department's research infrastructure and to aid in the education and training of African-American students. This has meant, for example, that student support was used to bridge the gap between funding of faculty members who were between grants. In general, the funds were frequently used to provide the most leverage for strong research ventures.

It has also meant providing a partial scholarship to a senior student whose family home in the ninth ward of New Orleans was devastated by Hurricane Katrina to allow the student to enter the workforce and provide some support to his family before attending graduate school.

## **9. Future and Strategic Directions**

The department expects to be very busy next academic year after the expected final approval of the department's proposed Ph.D. program by the Howard University Board of Trustees.

The activities needed for getting a Ph.D. program off to a good start are more than enough to keep any small department busy for the next year.

The multiple set of papers describing the comparative research necessary for the creation and development of a new PhD. program will be completed.

Outreach will be a much more important part of the project once the Ph.D. is approved. Multiple recruiting trips and increased interactions are expected.

The department expects to complete final negotiations momentarily with a new hire – an African-American female with a strong research record in the area of distributed systems, which is one of the key research areas. Outreach and recruiting should be more successful with a faculty member who can serve as a role model to the largest majority of our students.

# **CRI: Planning SWAMI – A Model of Support for a Women and Minority PhD Pipeline**

NSF 0551583, 3/1/06–2/28/07, year 1 of 1

Hunter College of the City University of New York

Virginia Teller (PI) and Stewart Weiss (co-PI)

Department of Computer Science

{vteller,sweiss}@hunter.cuny.edu

<http://www.compsci.hunter.cuny.edu>

Mary Flanagan (co-PI)

Department of Film and Media Studies

[mary@maryflanagan.com](mailto:mary@maryflanagan.com)

<http://tiltfactor.org>

## **ABSTRACT**

The investigators at Hunter College CUNY, a four-year, public, urban, minority institution with over 20,000 students, will use this award to plan a comprehensive infrastructure model for attracting and retaining female and minority students in the computer science discipline. During the grant period we will plan and prepare the larger infrastructure acquisition grant proposal, with the goal of creating a well-designed and well-supported PhD pipeline at Hunter College, one that focuses on research, creativity, and diversity within a robust computer science education.

## **I. Introduction/Background**

The investigators will draw on Hunter's unique resources to create a successful infrastructure upgrade and acquisition scheme based on a detailed assessment of needs in the areas of facilities and support systems. A set of hypotheses about female and minority recruitment and retention will be converted into a formal study to test and evaluate the effectiveness of the enhanced infrastructure in attracting and retaining the target student population.

The project is a collaboration between the Tiltfactor Laboratory in the Department of Film and Media Studies and the Computer Science Department. The Tiltfactor research group promotes research, scholarship, and creative practice centered on the social impact of digital technology, in particular, gender and computing. The investigators will also draw on the resources of other NSF-funded projects at Hunter such as the Louis Stokes Alliance for Minority Progress (LS-AMP) and the Gender Equity Project (an ADVANCE institutional transformation award).

Planning activities include designing the research study, developing an assessment plan and gathering baseline data, exploring outreach to high schools and community colleges, determining the requirements for and designing a state of the art undergraduate instructional computer laboratory, arranging support services such as tutoring and mentoring, exploring innovative curricula, solidifying collaborations to create undergraduate research opportunities, forming an external advisory committee, visiting computer science departments that have made progress in increasing female and minority enrollments, and organizing a retreat to generate consensus, cohesion and synergy for the infrastructure acquisition proposal.

## **II. Goals, Objectives, Targeted Activities**

The specific objectives to be achieved during the planning grant period are the following:

1. Develop an assessment method for each proposed infrastructural “intervention” and gather baseline data relevant to each one.
2. Determine the requirements for an undergraduate instructional laboratory that will help attract women and minorities into the computer science major and design a laboratory that meets these guidelines.
3. Arrange different types of support services such as tutoring and mentoring.
4. Investigate possible alternative curricula for experimental purposes.

5. Redesign two faculty laboratories for the express purpose of facilitating undergraduate participation in research.
6. Solidify collaborations both within the Computer Science Department and with other departments at Hunter to create undergraduate research opportunities.
7. Form an external advisory board composed of experts in various aspects of the project such as assessment, laboratory design, and women and minorities in the STEM disciplines.
8. Explore outreach to New York City high schools and CUNY community colleges.

In addition, two planning activities are specifically aimed at gaining the knowledge and experience needed to prepare a CRI large infrastructure acquisition proposal:

1. Visit computer science departments that have made progress in increasing enrollments of minorities and women.
2. Organize a Computer Science Department retreat with collaborators to generate consensus, cohesion and synergy before writing the full proposal.

### III. Project Description

#### *Research Hypotheses*

During the planning phase we will refine the hypotheses listed below and design an experimental study to test their effectiveness over the five year period of a full grant.

1. *Mentors matter.* Role models and mentors can have a significant impact on recruitment and retention. Can we attract new students to the discipline through programs such as visiting scientists and peer mentors? The lack of role models and mentors in the professoriate constitutes a significant barrier to producing minority STEM graduates. During the planning period we will design a local version of "MentorNet," the National Electronic Industrial Mentoring Network for Women in Engineering and Science (<http://www.mentornet.net>). In addition, we will organize mentoring partnerships with faculty and visiting scientists. For example, four faculty and/or visitors per year or summer of each of the years of the grant could be partnered with undergraduate research assistants to share their research techniques, working methods, and enthusiasm. Moreover, a new infrastructure means we can competitively recruit top notch faculty as positions become open in the future. The hiring of women and minority candidates will provide more mentors and role models.
2. *Attitudes matter.* The socio-cultural aspects of retention must be considered. Will the SWAMI program have positive psychological effects on users' motivation, attitudes, self-efficacy, confidence, and future intentions regarding careers in IT? We also cannot overlook the fact that computing field today has a serious image problem that must be overcome if we are to succeed in convincing the very brightest college students to study computer science.
3. *Context matters.* Will involvement in gender and minority-centered computer science research positively affect student interest? We will specifically build the SWAMI project on prior research conducted at Hunter and will synthesize findings from the Gender Equity Project, a recent assessment of retention and teaching in our department, and the Tiltfactor research, which has focused on programming instruction for two cohorts, middle school girls and introductory college courses. The software systems Flanagan and her collaborators have developed involve web-based collaborative experiences, build confidence, and provide content learning and reinforcement through feedback systems.
4. *Pedagogy matters.* Will the transformation of pedagogy, in particular, introductory course pedagogy, shift the student makeup? What methods such as pair programming, adopted in the Tiltfactor laboratory environment, will work best for our students? Pair programming is an example of collaborative learning, where two programmers work side-by-side at one computer, continuously collaborating on all aspects of the same problem: design, algorithm, code, and testing. Pair programmers are less likely to select a poor or inefficient approach to the task. NSF has funded several pair programming initiatives; of special interest to us is the research of McDowell and colleagues on the retention of women in computer science through the pair programming approach. Their results reveal compelling evidence that students working in pairs are more likely to complete introductory computer science courses, tackle challenging mathematics problems, **produce** higher quality programs, and enjoy the subject matter. McDowell et al. (2003) suggest that pair programming may

increase the representation of women in computer science because it fosters women's interest and promotes their success and confidence in the college classroom. Pair programming emphasizes collaboration and has been shown to be more effective than solo programming.

5. *Research matters.* Will enabling cutting edge faculty research attract and retain our target student population? The scope of the SWAMI project encompasses the compelling research areas of our faculty. For example, Hunter has faculty who are experts in the field of computer graphics, using computer vision for innovative techniques in 3D rendering. Computational linguistics is also an area of potential student interest. Computer science faculty members conduct diverse and fascinating research, publish in the top tier journals in their fields, and receive NSF grants (including an NSF CAREER award) for research, equipment and training. Facilitating these endeavors in part means providing a suitable technical and personnel infrastructure. It also means recruiting high quality, interested students to participate in these projects. The crucial common thread linking the disciplines incorporated into the SWAMI project is the incorporation of students into research projects.
6. *Interdisciplinary focus matters.* Can we recruit the target population through interesting and synergistic research projects among the diverse faculty at Hunter? Will cross-disciplinary collaboration work to interest students in computer science? Some research at Hunter College specifically addresses the target students' backgrounds through the design and development of educational software systems aimed at understanding and assisting underprivileged students. There are several projects that feed into the larger research questions of SWAMI, including instructional programs targeting younger women. New infrastructure will enable synergistic collaborations between the Computer Science Department and the Departments of Biological Sciences, Film and Media Studies and Geography as well as the Linguistics PhD Program at the CUNY Graduate Center.
7. *Facilities matter.* Research has shown that laboratories with state of the art equipment and technology based instructional systems are basic requirements for quality STEM education. Moreover, student perceptions of high quality are crucial to the attractiveness of any program. Through a detailed needs assessment in the planning phase, we will determine what facilities, both new and upgraded, and staffing levels will be adequate for the SWAMI project. In addition, because a stated goal of CRI is to extend the set of individuals and departments that are able to conduct innovative research, we will also conduct a needs analysis of the support necessary to foster innovative scientific discovery and collaboration among faculty and students in the related disciplines that will be included in the full research proposal.

#### *Planning Activities*

The activities listed in Section I are described in greater detail below.

**1. Assessment and Baseline Data.** We will seek the services of an outside consultant to assist in developing an adequate assessment plan for the SWAMI project. Both formative and summative evaluation measures will be established for each intervention to be included in the full CRI proposal. The research hypotheses listed above will be tested during the five-year CRI infrastructure acquisition project. Since the SWAMI infrastructure will also enable scientific advances that would not otherwise be possible, we must devise a way to measure the SWAMI impact in this area as well.

A major effort during the planning grant period will be devoted to collecting baseline data on every measurable aspect of the SWAMI project. This information will be gleaned from several sources, and some of it is already available. In Spring 2006 the Computer Science Department instituted its own system for tracking enrollments semester by semester as a means of projecting the demand for courses. This system will include a student survey each semester and will track majors and graduates as well as course enrollments.

**2. Laboratories.** Planning adequate computer laboratories for instruction and research are a top priority for the SWAMI project. Such facilities can figure prominently in increasing the number of minority science majors. In this area the SWAMI project has a valuable resource in Co-PI Stewart Weiss, who has an undergraduate degree in architecture. Combined with more than two decades in academic computer science, Weiss has the background and experience to tackle the difficult problem of laboratory design.

Weiss will spearhead a laboratory planning project on two fronts. On the undergraduate front the challenge will be to determine the requirements for an instructional laboratory that will attract women and minorities into the computer science major and design a laboratory that meets these guidelines. In addition, two faculty research laboratories – the Computer Science Faculty Research Laboratory and the Tiltfactor Laboratory – must be redesigned for the express purpose of facilitating undergraduate participation in research. We anticipate that the

Project Kaleidoscope manual *PKAL Volume III – Structures for Science: A Handbook on Planning Facilities for Undergraduate Natural Science Communities* will be a useful source of information in this endeavor.

**3. Support Services.** A logical place at Hunter to locate some of the support services such as tutoring and mentoring required for SWAMI is the Physical Sciences Learning Center (PSLC). This center, which is a part of the CUNY / New York City LS-AMP project, currently provides tutoring in physics, chemistry, and astronomy. There are computers with MS Office and Internet services and a library containing physical science textbooks and videos. Marilyn Rothschild, Director of the Center, is a former member of the Computer Science Department. We will work with her to plan an expansion of PSLC services to encompass the needs of computer science students.

**4. Curriculum Development.** The Computer Science Department has already adopted several measures aimed at improving the learning climate for 100-level students. For example, there are multiple entry routes into the major. The recommended route now starts with CSCI 127, an overview course, before CSCI 135, the introductory programming course. Alternatively, better prepared students can take CSCI 135 as their first course or exempt it completely on the basis of advanced placement computer science credit or a test-out exam given at Hunter. We also started a new one-credit laboratory course, CSCI 136, in Spring 2006, which gives CSCI 135 students the option of taking a supervised, hands-on laboratory course emphasizing practical computer skills and pair programming in conjunction with CSCI 135.

Another approach will be to explore alternative curricula that could be used for experimental purposes during the multi-year CRI project. Additionally, there are provisions in the MARC Curricular Improvement Phase II Implementation grant that Hunter is applying for in September 2005 to revamp the content and teaching methods used in eight computer science courses over a four-year period from 2006 to 2010.

**5. Research Opportunities.** Opportunities for undergraduate participation in research are a crucial component of the proposed SWAMI model. Our approach is to concentrate on interdisciplinary, collaborative research experiences that convey the excitement of discovery along with the usefulness of computer science in other disciplines. Current interdisciplinary projects in the Computer Science Department include: CUNY CoLAG, a collaboration of linguists and computer scientists using computer simulation to model first language acquisition; a joint project in bioinformatics with the Department of Biological Sciences exploring solutions to protein folding problems; and an urban planning project with the Geography Department that incorporates human perspectives – how people perceive and interact with their environments. This proposal also outlines our incipient work with the Tiltfactor laboratory in the Department of Film and Media Studies. These and other interdisciplinary collaborations will be solidified during the planning grant period.

**6. Advisory Board.** It is essential to have an external advisory board to oversee and monitor the SWAMI project. During the planning phase we will establish a committee composed of experts in various aspects of the project such as laboratory design, women and minorities in computer science and the STEM disciplines, pedagogy, and assessment. We will ask our advisory board to help us devise a management plan for the five-year project, critique our plans as we create them, and review the full proposal before we submit it. The board members must be willing to commit themselves to periodic meetings with the SWAMI team and to lend their collective expertise for the purposes of guiding the project and improving its outcomes.

**7. Outreach.** Outreach will target students at three levels – high schools, community colleges and Hunter undergraduates. At the high school level the focus will be on the College Now program. College Now is a collaboration between The City University of New York and the New York City Department of Education that enables qualified high school students to experience college by enrolling in undergraduate courses and attending college-sponsored events.

The Hunter component of College Now encompasses 32 public high schools in the five boroughs of New York City, including elite schools such as Hunter High School and Bronx High School of Science. Hunter has a liaison – a counselor or teacher – at each of these schools. Qualified students come to Hunter during the summer session and the fall and spring semesters to take special “cohort” classes composed only of high school students. A College Now student can earn up to 16 credits at Hunter. We will work closely with Clare Pistey, Director of Outreach and Special Programs, to expand the College Now offerings at Hunter to include computer science and to advertise Hunter’s computer science program at the College Now high schools.

At the community college level, CUNY recently submitted a proposal to NSF’s Broadening Participation in Computing initiative. The proposed project involves an alliance of six CUNY campuses – including Hunter – and consists of a comprehensive program to assist and support students transferring from a two-year degree program to a

four-year computer science program. Components include a summer bridge program, tutoring, mentoring and peer mentorships. Special emphasis will be placed on recruiting African American, Hispanic and female students majoring in computer science who are transferring from a CUNY community college into a CUNY senior college such as Hunter.

At Hunter, Teller, who is the Chair of Computer Science, has taken steps to interest entering freshmen in the field. Virtually all freshmen register for a block program in their first semester that offers a set of courses on various themes, and one of these blocks, "Changing Society Through Technology", includes CSCI 127, which is an overview of computer science. In Fall 2005 she taught the Orientation Seminar for this group of students. It is hoped that initial, albeit small, efforts such as these may help counteract the image problem that plagues the computing community today.

**8. Visits.** We plan to visit institutions that have made significant progress in increasing enrollments of women and minorities in computer science in order to get a first hand look at successful models. Our tentative strategy is for Flanagan to visit Carnegie Mellon University because of her expertise in gender issues. Teller will visit a school such as the University of Maryland Baltimore County, where great strides have been made in creating a PhD pipeline for minorities in the sciences. Weiss is expected to attend a Project Kaleidoscope workshop on laboratory design or a seminar on a similar topic related to innovative ways of utilizing limited spaces effectively. In addition, as PIs on a planning grant from a minority institution, the investigators will participate in the annual NSF CRI meeting for PIs from institutions with large awards at the beginning and again at the end of the planning grant period. The opportunity to meet and interact with these PIs will be extremely valuable in formulating and preparing our own large infrastructure acquisition proposal.

**9. Retreat.** Three years ago the Computer Science Department held its first retreat in connection with external program review. This highly successful, all-day meeting, which was held off-campus, proved instrumental in shaping the final form of the seven year strategic plan that we later submitted to Hunter's administration. We intend to organize another retreat for all of the key players in the SWAMI project to generate consensus, cohesion and synergy before writing the full CRI infrastructure acquisition proposal. As before, the retreat will be a one-day meeting held off-campus. The agenda will feature presentations and discussion related to the major components of the project. The funds budgeted for this purpose will allow us to invite some outside speakers to address SWAMI participants in their areas of expertise.

**10. Projected Timeline.** A chart showing a framework that indicates the type and anticipated duration of each task and activity during the planning phase will be included in our poster presentation at the CRI PIs meeting.

# CRI: A Research Infrastructure for Collaborative High-Performance Grid Applications

## EIA 02—02048, Fourth Year

Indiana University Bloomington

David Wise, Beth Plale, Andrew Lumsdaine, Geoffrey Fox, Randall Bramley,  
Kay Connelly, Dennis Gannon, and David Leake

dswise, plale, lums, gcf, connelly, gannon, leake at cs dot indiana dot edu  
<http://www.cs.indiana.edu/research/infra/>

### 1. Abstract

We are building an experimental infrastructure for distributed, high performance computing, in support of ten closely related research projects. These projects are extending the kind of location-transparency that the computational Grid provides for computation resources to the full spectrum of activities which end-users require. The services being explored include software development, parallel code middleware, distributed software components for scientific computing, security for parallel remote method invocation, managing large-scale data streams, and collaboration methodologies. This research builds on and extends our collaborations with several national Grid research teams.

### 2. Introduction/Background

The equipment forms a system capable of supporting the integrated range of activities needed in the research projects: parallel computation with clusters of SMP's, upgrades to both the internal and the external network connecting the department to remote instruments and data sources, a large high-speed storage device for staging and archiving, and infrastructure for collaboration. In addition to providing a set of important experimental support tools for our research projects, this facility has been designed to interoperate with the much larger Grid of resources that are being deployed nationally to support distributed, large-scale science. Altogether, our proposed facility represents a complete prototype development node for the NSF's new "Distributed Terascale Facility" (DTF) being installed at NCSA, NPACI, Argonne and Caltech. Several large production clusters similarly designed to work with the DTF are being installed at PACI sites, including Indiana University's central campus-wide computing center. The core system software we will install, the Oscar distribution of Linux and the Globus Grid and Access Grid software tools, are identical to the software that DTF and IU's production clusters will run.

### 3. Goals, Objectives, Targeted Activities

Several IU CS Department projects will be supported by the proposed facility. The ones detailed in this proposal are

**Arcee:** Basic work on parallel matrix algorithms that achieve high efficiency across many architectural platforms. specifically addressing locality within the memory hierarchy and parallelism.

**LAM and OpenMPI:** Middleware MPI implementations supporting hierarchical and fault-tolerant parallel computing.

**myLEAD:** Support for large-scale, grid-distributed data management.

**SURG:** Security for ubiquitous resources: user acceptance of ubiquitous and mobile technologies.

**HPJ:** High performance Java, creating a language platform for portable high performance coding.

**Grid Broker:** A reliable, robust publish/subscribe service for introducing fault-tolerance into the distributed Grid environment.

**Community Grids Collaboratory:** Advanced collaboration capabilities with applications to both distance education and distributed communities.

**CIMA:** The design of methodologies for remote instrument access and data management of the resulting extremely large data sets.

**Software Components:** A distributed software component model designed for applications that use parallel computing "nodes" in wide-area Grid environments.

**Science Portals:** A set of tools that allow programmers to build Grid distributed applications accessed and controlled from desktop environments and web browsers.

#### 4. Infrastructure and Facilities

##### 4.A Purchases.

##### 4.1. 8-node Aspen Systems Linux cluster (*thor* cluster).

Each node is a dual-processor system with 2.8GHz Intel Xeon processors with 2GB of RAM, with 4 PCIX slots, RedHat Linux EL4. It is connected to the SAN via a 2Gb/s Qlogic SANBlade 2300 Fibre Channel Host Bus Adapter. In order to support OpenMPI development it has connectivity via Myrinet, Infiniband, and Quadrics, as well as ethernet.

##### 4.2. Storage Area Network (SAN)

2.2TB of data storage. 16-port, 2Gb/s Qlogic SANBox2 Fibre Channel switch.

##### 4.3. Access Grid Node

With support from the department a meeting room was rebuilt to house the Access Grid Node for international video and multiconferencing. Cameras, sound, projectors are all provided by the grant, but dual use of the room is maintained with new lighting, wide whiteboards and a wide projection screen. Major components include 2 Hitachi Projectors, SX5600W; Canon VC-C4R, VC-C4 video camera and server, a PowerEdge 2650 Intel Xeon for audio/visual and a PowerEdge 2650 Intel Xeon for Video Out.

##### 4.4. The major purchase of the project occurred in the third year: a 144-node cluster purchased from Advanced Clustering Technologies in Kansas City, although Apple was awarded a bid to supply 17 (16 compute + 1 head) Apple Dual G5's.

- 128 Dual AMD Opteron (2.0GHZ) Compute Nodes (*odin* cluster) each with 4GB RAM, 40GB IDE drive, dual port Infiniband and dual 1Gb/s ethernet ports, Redhat EL4 Linux.
- 8 Dual AMD Opteron (2.0GHZ) Head Nodes (*odin/tyr* clusters) each with 8GB RAM, 40GB IDE drive, CDRW/DVD drive, fibre channel host adapter, and dual 1Gb/s ethernet ports, Redhat.
- 16 Dual Opteron (2.0GHZ) Web Service Nodes (*tyr* cluster) each with 16GB RAM, 40GB IDE drive, CDRW/DVD drive, fibre channel host bus adapter, and dual 1Gb/s Ethernet ports, Redhat.
- Apple G5 (*wotan* cluster) dual G5 (2.4GHZ) Head Node with 8GB RAM, 500GB Sata drives, CDRW/DVD drive, and dual 1Gb/s ethernet ports.
- 16 Apple G5 (*wotan* cluster) Dual G5 (2.4GHZ) compute nodes, each with CDRW/DVD, 80GB Sata drives, dual port Infiniband, Dual 1Gb/s Ethernet ports.
- 3.5TB of Fibre SAN RAID Disk Space.
- Remote Managed APC Power Strips for all devices in cluster.
- Serial port connections allowing console access to every device.
- BayStack Ethernet connectivity.
- UPS's to maintain power to critical nodes (Head Nodes/Web Service Nodes).
- Spare Parts kit to handle any hardware error.

##### 4.B. Staff support

One point cannot be ignored in reporting an infrastructure project on this scale. A new staff position was created specifically to support this equipment, with most of the salary contributed as part of the university's match of the award. With one person dedicated to its support, we have been able to sail through power failures, software upgrades, operating systems rolling in and out (Gentoo versus Redhat Linux), equipment failures and replacements, and other problems with which researchers and production staff ought not be bothered.

We thank Indiana University and NSF for including this support to keep the equipment invisibly alive.

#### 5. Project Description

5.1 The Common Component Architecture Forum (CCA) is a group of U.S. national laboratories and universities from around the world, developing methods of defining interfaces that will allow high-performance computing software to be more readily described, defined, and most importantly interchanged. The RI cluster has provided critical support in providing a testbed for CCA component research, and in disseminating CCA results in several tutorials at prominent national conferences.

5.2 The MxN problem is a special case of a general class of coupling parallel codes. This problem has become critical in areas like space weather modeling, fusion energy simulations, and other application areas. The general problem is to take existing HPC (and in particular, parallel) application codes that model some subphysics of a problem, and couple it with other subphysics codes to create a new, higher fidelity, integrated simulation. This

problem requires more than just a software solution, and introduces application science, applied mathematics, and computer science problems that are in turn coupled. The RI cluster is being used now to prototype and develop software frameworks for experimenting with this coupling problem, and has led to a funded contract with the DOE for coupling fusion simulations. This work has also fueled a set of proposals to the DOE's SciDAC program, including one to establish an Institute for Coupling Simulations at the University of Maryland.

5.3 The Proteus run-time system is now running on the new clusters. Proteus supports distributed software components by letting them dynamically negotiate communication protocols at run-time, then use the one selected as optimal. Proteus is currently being used to connect frameworks implementing the DoE's Common Component Architecture. Prompted by our NSF middleware project CIMA, we are exploring interoperability between Proteus and the Antelope Object Ring Buffer system for sending instrument-driven data across the network.

5.4 As a Web service farm the `tyr` cluster is being used for production-level support of the CIMA (Common Instrument Middleware Architecture) NSF-funded project. Currently the instrument data from five diffractometry labs and one national collaboration at the Advanced Photon Source facility at Argonne National Lab is being channeled to `tyr` via web service interfaces, and stored on the SAN hosted on that cluster. The data categories being streamed include frames from diffractometers; Humidity and temperatures for crystallography samples, the bays (rooms) containing the diffractometers, and rooms containing the instrument bays and sample preparation area; Streaming video of chemical samples being prepared, currently mounted sample being analyzed, general preparation areas and analysis workstations; Liquid nitrogen supplies; and Campus chilled-water supply to the labs.

This range of data and data types has validated the CIMA methodology, and in addition has provided critical support for a range of remote crystallography analysis. The overall large scale data-management system uses the `tyr` cluster as a front end and RI storage systems as intermediate stages to campus-level HPSS tape archival storage. This data system has been used for teaching and outreach activities also, ranging from groups of K-12 and college students being given an overview of X-ray crystallography to complete remote data acquisition and analysis being done by classes at Ball State and other universities. Diffractometry labs using CIMA and the RI equipment include ones at Purdue, the University of Minnesota, James Cook University, and the University of Sydney, and the ChemMat-CARS consortium beamline at Argonne National Lab. Current work is extending this web services methodology to data collection to mass spectrometry and long term photometric studies in astronomy.

5.5 The dQUOB/Calder streams project used the cluster resources for experimentally evaluating a new algorithm for fusing data streams together. The algorithm samples the rates of the data streams participating in the fuse operation, and builds a sliding window of a size based on a time interval. The actual size of the sliding window is correlated to stream rate. The algorithm was tested with data streams derived from mesoscale meteorology data streams obtained from our collaborators in the LEAD project. We built a data stream emulator modeled after realistic streams from the nine or so most popular instrument sources used in mesoscale meteorology forecasting. The emulator was used to generate a realistic workload under which the sliding window algorithm was executed.

5.6 As the scale and complexity of data-driven computational science grows, so grows the burden on the scientists and students in managing the data products used and generated during experiments. Products must be moved and directories created. Search support in traditional file systems is arcane. While storage management tools can store rich metadata, these tools do not satisfy the nuances of the individual computational science researcher working alone or cooperatively. We have developed a service-enabled personal workspace tool, myLEAD, that actively manages metadata and data products for users. Inspired by the Globus MCS metadata catalog and layered on top of the UK e-Science OGSA-DAI database interface tool, myLEAD provides capture, storage and search tools to the computational scientist. We use the NSF RI-funded cluster during an extensive experimental evaluation of the myLEAD personal metadata catalog. We set up a myLEAD server on one node in the `tyr` cluster and ran client nodes on multiple `tyr` nodes to simulate user activities of inserting and querying metadata. We also used the cluster to carry out a performance comparison between myLEAD and the Globus Toolkit Metadata Catalog Service (MCS).

5.7 The `thor` compute cluster and the `tyr` web-service farm have been used extensively in support of a several large NSF and DOE funded projects. These include the NSF "ITR Linked Environments for Atmospheric Discovery (LEAD)" project (ATM-03 31480), the Department of Energy project, "Middleware Technology to Support Science Portals: A Gateway to the Grid," the NSF NMI project "Middleware for Grid Portal Development (NMI-0330613)," the NSF Teragrid award "ETF Grid Infrastructure Group: Providing System Management and Integration for the TeraGrid (SCI- 0503697)" and the DOE Scidac project "Center for Component Technology For

Terascale Simulation Software.” The `thor` cluster was used to generate a simulated workload of meteorology data. The workload simulator processes generated-data streams that reflected the rate and size characteristics of meteorological data products. This simulated workload was then used to test the memory utilization of the Calder stream-processing system.

The `thor` cluster was also used for extensive performance and scalability tests on the Generic Application Factory and the Scientific Application Services created by it. The later are used in several weather-forecasting work flows in the Linked Environments for Atmospheric Discovery project.

Three nodes in the `tyr` cluster host the Karma provenance service, the WS-Messenger eventing service, and the `mySQL` data base used by these services. The `tyr` nodes, along with 108 nodes in the `odin` cluster, were used to run parallel simulations of ensemble WRF work flows for a comparative scalability study of the Karma provenance-collection framework

## 6. Major Accomplishments

6.1 The component architecture that debuted in LAM/MPI v7.0 has been carried forward into the OpenMPI project and been greatly expanded. Whereas LAM/MPI had four component types, OpenMPI has nearly thirty. It allows significant functionality for end users, system administrators, and third party developers. The framework allows both the run-time selection of which modules are used in the execution of a parallel MPI application as well as the arbitrary parameterization of those modules. This allows users to extensively customize their application and tweak performance characteristics in order to achieve maximum performance. System administrators can set defaults both for which modules are used as well as the “tweakable” parameters that are passed to them at run time. This allows even novice users to achieve maximum performance without needing to know the details for a given parallel environment. third party developers can utilize published interfaces to write their own modules that become part of OpenMPI at run time. This allows independent researchers to experiment in their given field of study without needing to know detailed knowledge of the rest of the OpenMPI implementation. For example, algorithm researchers can implement new versions of MPI collective routines simply by implementing a new, stand-alone collective component; it is not necessary to add code anywhere else in OpenMPI.

OpenMPI is an open-source, production-quality implementation of the Message Passing Interface (MPI), and was released November 2005. OpenMPI includes support for all of MPI-2. OpenMPI intends to continue and build upon LAM/MPI's widely-established reputation, has long included in all Linux distributions that provide clustering/high performance computing support (Red Hat, Mandrake, Yellow Dog / Black Lab, Debian, etc.) as well as multiple BSD distributions. We hope to replace LAM/MPI with OpenMPI in all of these venues.

The research infrastructure played an integral role in the evolution of the OpenMPI project. It continues to be used as a significant development and testing platform for the component architecture mentioned above, MPI parallel checkpoint-restart support, and native high-speed/low-latency Myrinet, Infiniband, and Quadrics networking support.




Figure 1: Cholesky factorization on 2.8Ghz Xeon.




Figure 2: TLB misses for tests from Figure 1.

6.2 The Arcee project is developing a paradigm of block-recursive programming on matrices to develop locality both for chip-multiprocessing and for more classic distributed multiprocessing that now uses OpenMPI. The future of on-chip multiprocessing has multiple cores demanding more locality among their cache loads; RAM won't matter among those processors. Arcee's quadtree algorithms can deliver that locality implicitly.

Figure 1, as an example, shows time (normalized to cycles/flop) for C-source Cholesky factorization competing with Intel's MKL (BLAS3) library, which uses machine-coding techniques developed at the University of Texas to preload cache. Arcee's C code accomplishes that implicitly with its recursion and limited intrinsic coding only for its 32x32 base case. Processor speed are plotted and enlightened by Figure 2 that shows how the OS-dependent TLB misses nearly disappear under the paradigm. Arcee's contribution is that no special coding provided or anticipated these behaviors. The serendipitous improvements just happen.

A new insight at the macroscopic level is still under development. Just as cache locality from recursive quadrants allows excellent performance at the low level, so also does it allow locality of access among processes at the macroscopic level. That is, big blocks---say 1024x1024---need reside only in one distributed processor while they are updated. The insight is that each need reside only on one distributed memory for almost all of the computation. This opens a new paradigm for parallelism.

## 7. Publications

Seventy-one articles and seven books or dissertations used the resources described here (listed in the formal report). In addition, software supported by this equipment includes LAM/MPI (a high-quality open-source implementation of the Message-Passing Interface), OpenMPI (released in November 2005), and the parallel Boost graph library. These are all distributed widely via the internet.

## 8. Outreach Programs

Jeffrey M. Squyres wrote a series of nineteen monthly columns in *ClusterWorld* magazine on the development of MPI as it evolved into OpenMPI with the support of this infrastructure. As the cluster of eight dual-Xeon ages, it has been made available on a limited basis to several outside organizations for software testing. The release of OpenMPI in November 2005 is a culmination of major sharing of the `thor` cluster which has several different networks installed (ethernet, In niband, Myrinet, and Quadrics). In 2005 Precision I/O was allowed to test an alternate TCP stack kernel to support the OpenMPI effort. We temporarily configured four nodes with two Myrinet cards each to allow them to test their new 'dual-rail' mx driver for these cards.

The NSF middleware project 'Common Instrument Middleware Architecture' is creating Web service interfaces to scientific instruments and relies critically upon the RI facility for long-term, reliable data services. Collaborators include Purdue University, two laboratories in IU's Biology Department, an observatory managed by IU's Astronomy Department, the IU Molecular Structures Center in the Chemistry Department, and an X-ray crystallography lab at the University of Southampton. As part of a DOE contract for the Common Component Architecture Forum, we are collaborating with the Princeton Particle Physics Lab and Oak Ridge National Laboratory in porting, optimizing, and componentizing codes for the simulation of magnetically con ned fusion energy.

The RI facility supports this effort by giving us ready access to large clusters where we can experiment with different communications protocols and systems implementations. The `thor`, `wotan`, and `odin` clusters are used to support the OpenMP development program with Lawrence Berkeley Labs and other vendors. Its initial release in November 2005 was a significant milestone in this project.

Kay Connelly's research used project-provided Palm PDAs in two major projects. DIMA is a nutrition-monitoring application for dialysis patients. In this project, we performed a user study to determine that elders and the ill can physically use PDAs as well as young, healthy people. Based on this work, we organized a successful workshop at ACM's SIGCHI meeting focused on how to perform user studies in non-traditional environments. The second project was on minimizing unwanted cell phone interruptions. This work culminated in a PhD dissertation and showed that electronic calendars are a good predictor of people's availability under certain conditions.

## 9. Future and Strategic Directions

The future for an infrastructure project of this genre is best described by the purchases under consideration for the future. These fall into three groups:

1. Upgrading the `odin` cluster (in situ) from simple AMD Opteron to dual-core Opteron chips. Since our cooling capacity went from 20% to 95% saturation with the installation of this cluster, the availability of affordable low-wattage processors is critical.
2. Purchase of a small quad-core, dual-processor cluster. Demands of locality change with the redistribution of the memory hierarchy that is implicit in multi-core chips that is different from shared RAM of dual-processor nodes.
3. Further expansion of SAN capacity and support for data acquisition and streaming.

Fortunately, the network and campus connectivity already obtained through this project is adequate to support these simple additions.

# CRI: Planning a Heterogeneous Data Repository for Computing Research

Proposal #: 0454355; 2005-2006  
Louisiana Tech University

**PIs:** Sumeet Dua (Louisiana Tech), PI<sup>1</sup>; Vijay Raghavan (UL-Lafayette), Co-PI<sup>2</sup>; Stephen Winters-Hilt (UNO), Co-PI<sup>3</sup>; Yenumala B. Reddy (Grambling), Co-PI<sup>4</sup>; Hilary W. Thompson (LSUHSC), SF<sup>5</sup>; Charlie McMahon (LSU), SF<sup>6</sup>; Seth Pincus (Children's Hospital), SF<sup>7</sup>.

**PI E-mails:** <sup>1</sup>[sdua@coes.latech.edu](mailto:sdua@coes.latech.edu), <sup>2</sup>[raghavan@cacs.louisiana.edu](mailto:raghavan@cacs.louisiana.edu), <sup>3</sup>[winters@cs.uno.edu](mailto:winters@cs.uno.edu), <sup>4</sup>[ybreddy@gram.edu](mailto:ybreddy@gram.edu), <sup>5</sup>[hthomp@lsu.edu](mailto:hthomp@lsu.edu), <sup>6</sup>[cmcmaho@lsu.edu](mailto:cmcmaho@lsu.edu), <sup>7</sup>[spincus@chnola-research.org](mailto:spincus@chnola-research.org)

## ABSTRACT

The goal of the CRI planning project is to assemble a multi-disciplinary and multi-institutional team of computer and information scientists, engineers, biologists, and clinical researchers which will develop a NSF-CISE-CRI large infrastructure acquisition proposal for building a heterogeneous data repository of Biomedical related data that will enhance cutting-edge computer science research in the State of Louisiana and beyond. One of the overarching goals of the fields of databases and data mining in computer science research is *the investigation of algorithms for the improved integration and mining of enormously large data sources*. The heterogeneous nature of biomedical, biological, and clinical data—ranging from free-text notes and instrument readings to confocal microscope images—provides an excellent research infrastructure and test-bed for development of novel computing technologies. The inherent dimensionality of the repository will stimulate the development of novel research outcomes covering essential areas of computer science that are at the heart of the current data-rich, but information-poor, technological revolution. The proposed consortium—an amalgamation of engineering, scientific, and clinical practitioners throughout Louisiana—builds on existing strengths within the higher education community and focuses on potential follow-up economic developments and investments in the private, state, and federal sectors. The development of a large infrastructure acquisition proposal involves researchers and students from a variety of disciplines and several Universities, including minority institutions. The project will also serve as a learning experience for students in interdisciplinary science courses. Involvement of underrepresented minority students will be specifically encouraged at various stages of the planning project. Once completed, the project will also have strong potential for extended outreach to several other institutions in the State, including a K-12 outreach component. Furthermore, existing partnerships with regional and national industrial partners will also be strengthened, in support of the economic development mission of the State of Louisiana (Louisiana Vision 2020). Additional ties to Louisiana companies and spin-off companies from the proposed collaboration will be pursued to improve the transfer of technology to commercial use.

## INTRODUCTION/BACKGROUND

Over the past decade, clinical and research-oriented medical disciplines have become increasingly data-intensive. The advances in automated data collection technologies in these domains have led to an unprecedented growth in the size, complexity, and quantity of collected data, a large proportion of which is currently inaccessible for analysis by computational scientists. Furthermore, in direct health-care applications the computerized medical record is on the verge of becoming practical, and soon after will be a business necessity for health-care providers, further swelling the data growth. This growth will lead to further renewal of demands for the development of novel technologies designed for the organization and mining of this data to enhance computing and biomedical research.

One of the overarching goals of database and data mining research in computer science is the investigation of algorithms for the improved integration and mining of enormously large data sources. Such sources include biomedical, biological, and clinical data consisting of a mixture of data types ranging from free-text notes and instrument readings to confocal microscope images. As storage capacity continues to grow, hyper-spectral images, X-rays, and similar complex data types will also become abundantly available. These data provide a superb research test-bed for development of novel computing technologies. Areas of possible concentration are data integration, data mining, database modeling, query translation, query optimization, information storage and retrieval, rule-based databases, feature recognition and indexing, information relevance feedback, and content-based queries. Moreover, it is commonly understood that such an integrated, panoramic view of the data provides an opportunity for interesting biological insights and discoveries otherwise extremely difficult or impossible (and thus, previously

unlikely). Aside from issues of mere availability, biomedical data is also complex, highly context-dependent, and inherently heterogeneous. This dimensionality makes the rendering of insightful knowledge through interpretation of raw data a challenging computational task. Advances in high-performance computing and high-speed networking infrastructure have led to efforts to utilize this infrastructure for high-end computing challenges originating in biomedical sciences. According to a report submitted by the NSF Directorate for Biological Sciences Advisory Committee (BIOAC) in 2003, an effective cyberinfrastructure should “provide a means to establish: (1) the tools for capturing, storing, and managing data; (2) the tools for organizing, finding and analyzing the data to obtain information; (3) the connection of experimental and theoretical analyses and their interplay with simulations and complex models based on that information; and (4) the integration of disparate aspects of that information to provide a synthesis, a knowledge repository for further considerations”<sup>1</sup>. Clearly, an integrated storage infrastructure is the first and most important step necessary to address this challenge. The proposed infrastructure also opens up challenging and important avenues for advanced computing research.

## GOALS, OBJECTIVES, TARGETED ACTIVITIES




Figure 3: LA-HealthGrid Team Members and Collaborators

The *goal* of the proposed planning project is to assemble a team of engineering, scientific, and medical practitioners and researchers (Figure 1) which will collaboratively develop of a NSF-CISE-CRI large infrastructure acquisition proposal for building a heterogeneous data repository of biomedical-related data that will enhance cutting-edge computer science research in the State of Louisiana and beyond. Such a repository would also assist biological, biomedical, and clinical researchers in their data intensive investigative endeavors.

The *objectives* of the planning project include:

1. Identification of a multi-disciplinary and multi-institutional team of researchers and practitioners to undergo planning for the development of a heterogeneous data repository of biomedical and clinical data available in the State of Louisiana and beyond;
2. Investigation of the fitting architecture for such a repository, using *cutting-edge computing research* as the decision principle;
3. Identification of key computing research fronts in the development and utilization of this data repository;
4. Identification of new educational and outreach opportunities engendered during and following after the development of this data repository for K-12 students as well as State universities, and
5. Development of a competitive ‘large’ infrastructure acquisition proposal to be submitted under NSF/CISE/CRI program for funding consideration during the next review cycle.

The above project objectives are being progressively completed through the following key *targeted activities*:

1. Identify a team of computing scientists across the State and beyond, who will benefit from the proposed repository through the enhancement of their research and training goals;

<sup>1</sup> Wooley, John C., Chair. *Building a Cyberinfrastructure for the Biological Sciences (CIBIO). 2005 and Beyond: A Roadmap for Consolidation and Exponentiation*. NSF Directorate for Biological Sciences Advisory Committee (BIOAC). Subcommittee on 21<sup>st</sup> Century Biology. July 14-15, 2003, 4.

2. Identify partners from Louisiana's computer science, biological, biomedical, and clinical departments, centers, and institutions who will share their data for this investigative endeavor;
3. Identify sources for deployment, operation, and usability of high performance computing and collaborative visualization hardware for data storage and manipulation that will build a modular and imminently expandable skeletal cyberinfrastructure to enable this collaborative endeavor;
4. The design and development of cutting-edge computing applications, algorithms, and functional simulation models will provide a middleware musculature enabling interoperability and complex data analysis and integration in such heterogeneously distributed domains, and
5. Identify the significant impact areas in educational outreach and economic development during and after the completion of this project.

In an effort to address many of the above objectives, we organized a workshop entitled "Workshop for Healthcare Data Integration and Mining for Next-generation Medical Research, Education and Service in Louisiana:" from March 16-17, 2006 at Louisiana Tech University. More details on the workshop are available at <http://mycenit.latech.edu/LAdataworkshop/>; its objectives and accomplishments are discussed later in this document.

## **INFRASTRUCTURE AND FACILITIES**

This planning project does not involve acquisition of infrastructure. However, as a part of one of our strategic objectives we have identified and collaborated with the following key infrastructural resources that would be the key to the design and development of the full proposal and its effective implementation.

Louisiana Optical Network Initiative (LONI; <http://www.loni.org>): The Louisiana Optical Network Initiative (LONI) will create a statewide environment capable of integrating and aggregating Louisiana's considerable strength, which is currently distributed across its universities and industries. Employing "lambda" technology, (each lambda is capable of carrying 10 Gbits/sec), the initial deployment of LONI will provide 40 Gbits/sec network connections between Louisiana's research institutions and 20 Gbits/sec connections to medical centers in New Orleans and Shreveport. LONI's founding mission is to advance science, understanding, and stewardship for cutting-edge research by fostering interdisciplinary activities and synergy among professional, scientific, and educational organizations throughout the State. In addition to establishing the network itself, the LONI initiative recognizes the crucial need to develop Louisiana's grid computing environment. LONI will seed the development of these facilities by establishing mid-sized computing clusters located at LA Tech, UL Lafayette, SU-Baton Rouge, UNO, and Tulane University. In sum, the LONI initiative will construct a high-speed computational backbone (both in network connectivity and access to HPC resources) through which to effectively and efficiently implement a pseudo-federated data repository that will enable real-time data and tool interoperability. LONI is funded by the State of Louisiana for \$40Million. PI, Dr. Sumeet Dua, has been invited to serve as the secretary of the LONI scientific advisory council (the research arm of LONI) and also leads the LONI efforts at Louisiana Tech University.

Access Grid (<http://dmrl.latech.edu/Access-Grid/>): The creation of this broad and paradigm defining effort is also facilitated through the use of the Access Grid (AG). The AG, made possible by an NIH-NCCR-INBRE grant, has been constructed as an essential part of the global open source project aimed at developing a large-scale collaborative environment similar to video-conferencing rooms today, but scaled up dramatically in flexibility, functionality, and performance. The AG is capable of interacting simultaneously with any number of locations across the country and worldwide, providing the ability to send and display live video and audio streams from multiple sources in order to share work documents and complex scientific images (Figure 2). By creating a virtual environment in which a group of participants have the ability to hear and see each other, as well as exchange visualization in real time, without delays, this technology is highly useful for multi-institutional data sharing and collaboration. It will also aid in outreach by eliminating geographical and institutional barriers.



*Figure 2: The AG in operation, simultaneously displaying participants & data.*

IBM p5-575 Shared-Memory Supercomputing Grid: Louisiana Tech and four other partner institutions in the State (UL-Lafayette, LSU, Tulane and Southern University) have acquired an IBM supercomputer that will serve as constituents of the high performance computing (HPC) environment which will exploit LONI. With 14 node racks, each containing 128, 1.9 GHz p5 processors, connected by integrated gigabit interconnected network ethernet, these computation workhorses provide stellar performance of about 0.85 TeraFlops with undisputed reliability and

stability, and a robust programming environment and hardware features. Coupled with the LONI pipeline and the computational test-bed provided by the LA-HealthGrid collaboration and enabled by communication over the Access Grid network, advances in high-performance computing and high-speed networking infrastructure will lead to efforts to utilize this HPC infrastructure for high-end computing challenges originating in biomedical sciences.

Media Library/Repository: The media library developed by researchers at LSUHSC and LaTech is a basic data resource for developing clinical decision support algorithms using data mining. It is a multi-tiered, extensible database application system that supports the storage, retrieval, indexing, and processing of any type of generic large data object (media) along with its associated metadata. The database is implemented with a novel object-relational design that supports heterogeneous data types within the database. Features of the system include secure remote access using encrypted client-to-server communication, a user and media access control system, content-based query processing, and an import/export facility for collections of media and metadata using portable, application-independent XML catalogs, and multimodal search and retrieval of archived media. Initial tests included the storage, 3D reconstruction, and immersive visualization of image stacks from the lamina cribrosa of the cow eye. Legacy data are being migrated to the central database as time and resources allow. Centralized data storage has many benefits: performance, fault tolerance, increased availability, and efficient searches among heterogeneous data types. This media library is just one example of the type of flexible and valuable technologies and resources that an effective cyberinfrastructure will enable. Initial conversion to such systems can be difficult and time consuming, but the eventual research benefits will far outweigh the initial capital investment.

## **PROJECT DESCRIPTION**

---

The mission of the proposed data repository is to construct a nationally preeminent and internationally competitive infrastructure for enhanced computing research in the State of Louisiana and beyond. This repository will lead to the development of novel research outcomes in the areas of content-based image retrieval, data assimilation, design and development of efficient, high-performance algorithms for data integration, query translation, query and database optimization, data mining, information storage and retrieval, and data warehousing. The proposed consortium will amalgamate engineers, scientists, and practitioners into a multi-disciplinary and multi-institutional research team within Louisiana that builds on existing strengths within the community of higher education. It will provide high potential for follow-up economic development including investments from private sector, state, and federal sources. The proposed repository, once completed, would perform the following key functions: It will

1. Assimilate heterogeneous data from multiple data sources (data providers) into a hybrid-integrated schema of storage;
2. Provide many opportunities for development of novel computing technologies designed to augment and advance the repository;
3. Provide a test-bed upon which computing scientists can develop and test technologies for data integration, query optimization/translation, feature extraction, content-based queries, and distributed data mining;
4. Provide a portal for researchers to Extract, Transform, and Load (ETL) data through web-based tools, and
5. Provide an interface for access to this data using complex context-based queries.

## **INDICATORS OF SUCCESS AND MAJOR ACCOMPLISHMENTS**

---

Delays Overcome: Only two months after receiving funding approval for this planning proposal, hurricanes Katrina and Rita devastated the Gulf Coast and southern Louisiana. These storms not only caused major destruction at many of the institutions collaborating in this project, but it also sundered participating members from their research resources and reliable means of communication. By October '05, some contributors could not be located, while others were desperately contacting other participants, hoping to recover key data resources lost in the crisis. Despite these events, which could have resulted in serious setbacks and delays in the planning process, we continued to forge forward. Regular strategic communication and planning was entirely revived by January, shortly after which a major consortium event (LA-HealthGrid workshop) was organized, scheduled, and executed in March to outline the primary questions, needs, and capabilities of high performance computing and biomedical sciences in Louisiana. This meeting assembled a diverse company of computer and information scientists, biologists, clinical researchers, and industry representatives who were presented with the terms and challenges of developing the large infrastructure and heterogeneous data repository of biomedical-related data mandated by the NSF-CISE-CRI proposal. This meeting included original members as well as new contributors (see Outreach Programs and Supported Students).

Neither the storm nor the ensuing confusion and trouble have inhibited the progress of constructing the LONI optical network pipeline. To date, its construction, deployment, and operation is running entirely according to the original schedule. In short, despite incredible obstacles and difficulties beyond human control, the momentum of the planning stages continues to mount. The determination of the contributors and collaborators during this difficult time indicates that their intense tenacity and drive to complete their goals and objective is unparalleled.

**Planning Workshop and Establishment of LA-HealthGrid Consortium:** As a key component of this project, we organized a workshop entitled “Workshop for Healthcare Data Integration and Mining for Next-generation Medical Research, Education and Service in Louisiana” from March 16-17, 2006. A total of 59 leading domain scientists and engineers representing over 15 institutions and companies participated in the workshop. The purpose of this workshop was to identify and discuss the opportunities, challenges, and technologies involved in constructing a statewide medical data repository and grid network (LA-HealthGrid). The mission of the proposed data repository is to construct a nationally preeminent and internationally competitive infrastructure for enhanced computing research and consequently, healthcare delivery in the State of Louisiana. The following objectives were addressed at that assembly: (1) Review the state of and future directions of collaboration and associated technologies, with a particular focus on data capture, deposition, and use. (2) Identify computing challenges induced by advancements technological data generation. (3) Architect a collaborative, heterogeneous repository of healthcare data, based upon effectively exploiting and extending current investments and resources. (4) Focus on healthcare/biomedical applications by involving members biomedical and clinical research communities, but also cyber infrastructure (grid) and other information technology communities. (5) Assemble multi-disciplinary and multi-institutional team that will include medical professionals from diverse fields of healthcare and research. (6) Discuss the opportunities, challenges, and technologies involved in constructing a statewide medical data repository and grid network (LA-HealthGrid). (7) Plan a state-of-the-art infrastructure for enhanced computing research quality healthcare.

Institutional representatives were asked the following specific questions: (1) What type and degree of data is available for sharing and analysis? (2) What are the opportunities for collaborative approaches to data? (3) What could the healthcare and biomedical community expect from the computing research endeavors? (4) What should the goals of a collaboratory program be? (5) How should a collaboratory program be evaluated (metrics of success)? (6) What are the overriding impediments to the success of this collaboratory? (7) How can you partner in developing a collaboratory with broad scope, vision, and possibilities? The workshop proceedings were documented in detail (and issues demarcated between opportunities, challenges, and technology) and a report generated. The detailed report is available at the workshop website: <http://mycenit.latech.edu/LAdataworkshop/>.

## **OUTREACH PROGRAMS AND SUPPORTED STUDENTS**

---

The possibilities surrounding the completion of LONI and the creation of a biomedical and clinical health repository has generated copious interest across the state from diverse sectors. The LA-HealthGrid meeting alone attracted new contributors from statewide universities, investigators in critical research fields, and from corporations in the private sector. Presenters included members of the Louisiana Cancer Research Consortium, Louisiana Board of Regents, Louisiana Healthcare Review, Inc., and several universities. The opportunities for high performance computing and information-based, on-demand medical network infrastructure have also attracted the eager interest of representatives of IBM Corp., as well ARAICOM Life Sciences, LLC. HBCUs (Grambling State University and Southern University) have also been involved in this project. Although the planning project did not provide for student support, 11 students volunteered to manage the workshop and were afforded unique opportunity to interact with participants during and after the workshop.

## **FUTURE AND STRATEGIC DIRECTIONS**

---

The key, future directions for this project include, but are not limited to, the following:

1. Continued development and submission of a CRI proposal for design, development, and acquisition of the infrastructure.
2. Adopt and support a data sharing/integration standard to enable real-time data and tool interoperability.
3. Identify core partner roles, an external advisory committee (management plan), and training opportunities (educational plan).
4. Develop formal data-sharing agreements with key institutional partners to proactively prepare towards the implementation of the proposed infrastructure.
5. Develop strategic relationships with small businesses and industry partners to identify and generate economic development opportunities.

# **CRI: Infrastructure for Networked Sensor Information Technology**

**Proposal # 0551734, awarded 04/01/2006**

**New Mexico State University**

Amiya Bhattacharya, PI (amiya@cs.nmsu.edu)

Derek Bailey, Co-PI (dwbailey@nmsu.edu)

Jonathan Cook, Co-PI (jcook@cs.nmsu.edu)

Clinton Jeffery, Co-PI (jeffery@cs.nmsu.edu)

Debra Peters, Co-PI (debpeter@nmsu.edu)

URL: <http://pecan.cs.nmsu.edu>

## **Abstract**

Developing infrastructure and core capabilities in the information technology of wireless sensor network, this project attempts to enable research on a number of core systems and applicative areas such as component-based middleware design, protection of contextual privacy, augmented collaborative virtual environment, tracking and monitoring of livestock in rangelands, and scalable micro-monitoring of environment and climate. New laboratory and courses at undergraduate and graduate levels, new technology for distance education, instrumentation for ecological research, and improved land and livestock management are among the broader impacts.

## **1. Introduction**

Availability of low-cost embedded processor and radio hardware with on-board sensors in recent years, along with widespread research in development of operating system and programming platform for these embedded devices has opened up the possibility of cost-effective ubiquitous use of wireless sensornets made out of commercial off-the-shelf components. These sensornets are expected to be the enabling technology for pervasive/ubiquitous computing by integrating tiny sensor/actuator devices into the fabric of our everyday life.

We envision a rapid growth of information technology into a new direction with a goal to reduce the risk of use of sensornets in applications ranging from ordinary to critical, as well as to simplify custom programming and management of these sensornets. We call this new direction “networked sensor information technology.” The computing research infrastructure project at New Mexico State University (NMSU) is aimed at establishing core capabilities and expertise in this growing area of research, by designing and deploying a carefully chosen set of distributed sensing applications offering diversity in system requirements while posing the need for a set of unifying design principles.

## **2. Goals, objectives and targeted activities**

The sensornet infrastructure at NMSU will be housed at three different sites. The central location of the facilities is a new state-of-the-art PECaN (Pervasive and Embedded Computing and Networking) Laboratory, within Science Hall in the main campus at Las Cruces, where the Department of Computer Science is located. The sensornet design and programming environment at PECaN-Lab will be primarily used for the following set of research projects:

1. Component-based middleware engineering to support dynamic component execution and reprogramming;
2. Protecting contextual privacy and thwarting traffic-analysis attack in distributed sensing tasks;
3. Integrating real-time sensor data to a collaborative virtual environment for distance learning;
4. Tracking and monitoring of livestock on rangeland using multiple localization technologies;
5. Scalable and cost-effective micro-monitoring of climate on an ecological research site.

In addition to PECaN Lab, the sensor nodes and embedded computers pertaining to the collaborative virtual environment project will be located in the classrooms, undergraduate and graduate computer laboratories, faculty research laboratories, and some faculty and staff offices.

The main site for research related to livestock tracking and behavior monitoring will be at the Chihuahuan Desert Rangeland Resource Center, approximately 25 miles north of the NMSU main campus. The pasture we plan on using has an area about 1000 to 1200 hectares (100 hectare = 1 km<sup>2</sup>). A second NMSU owned site at the Corona Range and Livestock Research Center near Corona, NM (about 185 miles from Las Cruces) will be used as a test site during the final year of the project, as a larger grazing area of up to 3000 hectare (30 km<sup>2</sup>) will satisfy the scalability requirement of this project.

Using our collaboration with the Jornada Experiment Range Agricultural Research Service (JER ARS), we plan to use a pasture (12A) of 1200 hectare area for the large-scale environment monitoring project. The Jornada Experimental Range is a USDA facility approximately 25 miles north-east from NMSU main campus. The Long Term Ecological Research program funded by NSF is already underway at this location.

### **3. Infrastructure and facilities**

#### **3.1 PECaN Lab:**

The PECaN laboratory is to be used primarily for design, simulation, and prototyping required by all current and future research projects in wireless sensor networks. In addition, students enrolled in courses related to sensor networks will also use the laboratory facility. Three workstations (3.2 GHz, 1 GB memory) will support the design and development activities of the three Graduate Research Assistants. In addition, a higher capability computing server (3.2 GHz dual Intel Xeon processor and 2 GB memory) will be used for large simulations. A few tablet PCs will be used across all sites to run client programs for sensor data gathering and node re-programming.

The software environment we plan to set up in the laboratory is mostly based on open source software including TinyOS/nesc suite and their alternatives such as SOS and MANTIS OS. Simulation software such as ns2 and OPNET modeler will be used as necessary. The MathWorks MATLAB/Simulink software suite including Stateflow and Real-time Workshop would provide a platform for rapid simulation, prototyping and visualization for physical processes. While ns2 and OPNET modeler excel in network protocol simulation, Simulink is expected to help understand the interaction of the sensor network with the physical processes from which sensor data originates.

Our goal is to have a heterogeneous set of hardware compatible to the TinyOS environment. For indoor networks, we plan on using both Crossbow MicaZ MPR2400 motes and Crossbow TelosB or Moteiv Tmote Sky. Based on availability, Intel Imote<sup>2</sup> will be used as nodes with intermediate capabilities. Motes of all three varieties can communicate with each other as they use identical radio hardware (Chipcon CC2420). Several Soekris 4526 and 4801 embedded Linux gateways are to be used in conjunction to provide gateway functionalities. Also, to develop intelligent mobile gateways and systems for secure resource tracking, two Acroname Garcia Robots will be used, with mounted Stargates equipped with WiFi interfaces.

#### **3.2 Science Hall (for Virtual Collaborative Environment):**

The sensor-based Virtual Collaborative Environment project requires sensor node placement in indoor office, laboratory and classroom environments. We estimate approximately 80 locations that need monitoring. To support higher data rates (up to 250 Kbps), we opt for sensor nodes supporting IEEE 802.15.4 compliant radio. Currently, the choice is between Crossbow MicaZ MPR 2400 motes with MTS310 sensors and Moteiv tmote invent. Light, temperature, microphone and accelerometer data are to be used in this project.

A workstation (with 3.2 GHz dual Intel Xeon processors, 2 GB memory, 3D graphics acceleration and at least 10,000 rpm disks will serve as the design environment for the project. The virtual environment software has been under development at NMSU for the past few years. The workstation also needs to provide 24-hour web-based access to queries on raw sensor data as well as virtual representation of the indoor environment. Three Crossbow SPB400 Stargates equipped with MicaZ motes and a CF 802.11 adapter, or Soekris net4801 equipped with tmote sky are to be used as sensornet gateways.

#### **3.3 Chihuahuan Desert / Corona Rangeland:**

Although there is considerable difference in the area covered between the two rangelands, livestock tracking and monitoring calls for covering the herd rather than the area. Prior experience with cattle behavior indicates that interesting intra-herd movement and social behavior is shown only after the herd size exceeds a critical number within a range of 40 to 50. Also, the herd size typically lies between 100 to 500 in commercial ranches. Thus, we plan to build approximately one hundred cattle collars with sensor nodes.

To obtain intra-herd fine-grain localization and tracking data, we plan on using Crossbow 433 MHz MPR410 Mica2 or MCS410 Cricket motes. Crickets are essentially Mica2 motes with localization capability built

using a combination of RF and ultrasound technology. The choice will be contingent on the success of using received signal strength indicator for fine grain localization on the desert terrain. Success in weatherproof packaging of Cricket motes will also influence the decision. The compatible sensor modules we plan on using are Crossbow MTS400 and MTS420. Both are equipped with sensors for temperature, humidity, pressure and ambient light, as well as a 2-axis accelerometer which will be useful to determine whether an animal is grazing or resting. The main difference is that the MTS420 is equipped with a Leadtek 9546 GPS module with 10 meter accuracy. The gateway functionalities will be provided by ten Crossbow SPB400 Stargate gateways that are to be left in strategic positions such as near water resources to ensure eventual data acquisition. Cellular data service will be used for carrying sensor data between the embedded gateways. A workstation (3.2 GHz, 1 GB memory) will be equipped with the TinyOS mote programming environment and a web interface to the herd sensornet database.

### **3.4 Jornada Experimental Range:**

The 12 km<sup>2</sup> area of pasture 12A will require of the order of one hundred Crossbow 433 MHz Mica2 (MPR410) motes for coverage. Compatible Crossbow MTS400 environmental sensors or MDA320 data acquisition boards will be required for each mote, depending on the need for using external sensors such as the ones for soil humidity. The most likely candidate will be the Crossbow MEP410/510 weatherproof packaged and sensor-equipped motes. The sensor gateway functionality will be served by five Crossbow SPB400 Stargate gateways, which is an embedded Linux single board computer based on the 400 MHz Intel X-Scale processor. Experiments will make use of one or more of these gateways in static and mobile data acquisition modes. The pasture currently has IEEE 802.11b wireless LAN coverage (with directed antennas), and CF 802.11 adapters will be used for data transfer between the Stargates and the Jornada Range headquarters. A workstation (3.2 GHz, 1 GB memory) will be equipped with the TinyOS mote programming environment and a web interface to the Jornada sensor database.

## **4. Project description**

### **4.1 Middleware to support dynamic component execution**

While component-based software engineering has been thoroughly exploited in the design of sensornet programming environment such as TinyOS, its execution model does not support the components at runtime. In fact, the model does not separate the application from the operating system, and creates a statically linked code out of the necessary components at compile-time using aggressive code optimization. With the availability of next generation sensor node hardware, sensornet operating systems are emerging with more dynamic runtime support for the component-based programming model. Runtime support for dynamic component execution and reprogramming in SOS based on software memory protection, or support for POSIX compliant lightweight thread in MANTIS are worthwhile examples.

Using existing OS support, this project aims to develop an efficient infrastructure to break the static-linking paradigm of sensor applications and enable dynamically-linked component deployment. This is a move to dynamic component binding, easily manipulated as connecting named required services to named provided services on other components, or even to connector components that act as intermediaries. Part of this step will be enabling the downloading of individual components to a sensor node, and of rebinding the existing application to the new component version. This step will improve the efficiency of software updates, which till now must re-deploy the entire monolithic application. Of course, this must come at the cost of sacrificing some code inlining done by nesC compilation. Research on supporting component evolution will need to involve methods for component specification, verification of component compatibilities, interface specification and checking, and build environment adaptations (e.g., compiler changes) for creating independent components. While nesC already contains syntactic interface specification, adding semantic information would allow static property checking to be done, especially since the applications are compiled in a whole-program step.

### **4.2 Contextual privacy against traffic analysis**

In a sensornet environment that cannot be physically secured from intruders, channel snooping and traffic analysis by an eavesdropper exposes the sensitive information about the location and time of the triggering events. The open-architecture of the underlying RF technology (such as IEEE 802.15.4) makes it particularly easy for adversaries to build custom channel monitor based on either their own COTS-based nodes or more sophisticated spectrum analyzer or software radio technologies. While considerable efforts have gone into designing secure protocols for code and data dissemination over wireless sensornets, these encryption-based protocols render

inadequate for protecting total privacy of a sensing task. Thus contextual privacy has become a growing concern in sensor network deployment for sensitive tasks in public areas.

Obfuscation of the contextual information is a natural solution to the contextual privacy problem. Since both location and time of a sensed event constitute its context, obfuscation can be performed both spatially and temporally. Most researchers address the contextual privacy problem by focusing mainly on the location component of the contextual information. As a result, the spatial obfuscation techniques employed (such as controlled and probabilistic flooding, fake messaging and phantom routing) pose a burden on the power-constrained sensor nodes due to the inevitable increase in message transmissions. In contrast, temporal obfuscation by delayed reporting of sensed data is possible, if the sensing task has a less stringent requirement than real-time reporting. To the best of our knowledge, this possibility has not been exploited earlier in the context of sensor data dissemination protocols. Indeed, it is the uncertainty due to variable delay introduced in reporting the sensed events that achieves the desired obfuscation. Coalescing several messages at the originating sensor node prior to encrypted transmissions, is a delaying approach that has the unique advantage of battery power savings.

Based on the observation that there is a potential trade-off between contextual privacy and power consumption, we propose to create a framework to trade off battery power for protecting contextual privacy. A natural consequence is that the framework must be aware of application semantics. We plan to abstract both privacy and responsiveness needs of a sensing task and integrate this abstraction as sensor database system extension. This leads to an optimization under the following three constraints: (i) adversarial power (number of nodes, radio range, power source, and node mobility), (ii) application requirements (privacy, responsiveness, and data aging), and (iii) power consumption.

### **4.3 Sensor-augmented collaborative virtual environment**

Using a three-dimensional graphical interface over a network, a Collaborative Virtual Environment can represent a location from the real world, or it can place users in an imaginary or abstract place. Collaborative Virtual Environments are used in many application domains such as entertainment and medicine. NMSU is currently building a virtual version of the real physical space in Science Hall that makes up the Computer Science Department, for the purposes of distance education. Sensor network infrastructure will provide unique opportunities to explore the use of virtual environment augmentation in both real and virtual activities. Sensors providing door and light switch information will assist people in recognizing each others' different types of "away from keyboard" time. Temperature information throughout the department will assist users in both immediate and long-term decisions about where and when to work. Administrators and physical plant services will be able to use the information to plan allocate machine rooms versus human space allocation. Approximate measurements of activity levels in laboratories will assist both faculty and administrators in scheduling various aspects of the labs and projects. Sensor network information provides input for simple remote control tasks, such as turning off shared lighting based on usage patterns. In all of these real-world augmentation tasks, the collaborative virtual environment provides a natural mechanism for visual representation of sensor network data. But, sensor information has the potential to be even more valuable to distant users of the collaborative virtual environment, who do not have the option of walking down the hall to gather a piece of information.

Sensor-augmented virtual environments can enable virtual users to see "ghosts" of non-virtual user activity. For example, if a tutor is helping other virtual environment users, it may be quite apparent. But, if a tutor is on the phone or has a few real world students waiting outside their door, the sensor network information can provide the virtual student with the information they need in order to decide whether to wait or come back later. In general, such an environment can create a more immersive virtual experience, in which remote users are passively cued regarding local time or other conditions in the real world environment. Day and night may be easily programmed without a sensor network, but light levels and other cues can help remote users recognize different circumstances on-campus students and faculty may be in (for example, a thunderstorm, or a fire drill).

### **4.4 Livestock tracking and behavior monitoring**

Managed grazing rangelands are typically partitioned into several pastures that can range from tens of hectares up to several hundreds and beyond. Pastures are utilized and managed so as not to overgraze them and to give them periods of rest and recovery. Managing livestock to efficiently and equally utilize the rangeland is an ongoing area of research with many challenges. For example, one problem in larger and steeper pastures, especially in areas with few water sources, is in getting animals to move to more remote areas and utilize the forage growing in those areas. Solving problems such as the one above, and others, is to improve the general knowledge of livestock behavior and movement patterns. Some of this might be general knowledge, such as which supplements (e.g., mineral or molasses) work best to attract cattle to under-utilized areas, and some of the knowledge might be specific,

such as finding individuals that are more likely to range further from water or up to higher elevations than others, and investigating whether or not such behavior is hereditary. Relationships between livestock within one herd are also an important area of ongoing research, and can lead to better methods for selecting which individuals to separate into different herds and how interactions might influence range utilization.

The initial and specific deployment of an outdoor sensor network will be in monitoring cattle movements both at macro-level (locations on a given range) and micro-level (relative positions to other cattle in a herd). Previous work using GPS collars to track macro-level movements of livestock gets too expensive to collar more than a few animals, and therefore cannot provide intra-herd micro-level movement data. Low-cost wireless sensor nodes, possibly equipped with inexpensive low-resolution GPS receivers, will overcome this instrumentation barrier. Experience with domestic livestock monitoring will also lead to new research efforts in wildlife monitoring. New algorithm for tracking moving animals using multiple localization techniques is a key to this deployment. Interference with the livestock should be kept to a minimum, and thus extending battery life as long as possible will be an important research goal. Adaptive sampling techniques based on multiple sensor readings is an. For example, if an animal is at rest, sensor data can be sampled at a slower rate, allowing the sensor hardware to sleep longer between periods of activity.

#### **4.5 Scalable micro-monitoring of climate and environment**

NMSU is located within the Chihuahuan Desert, a high-desert ecosystem consisting of desert grasslands and shrublands, and having a history of extensive use by humans, primarily for grazing. The Jornada Basin Long Term Ecological Research Program (JRN LTER) has been investigating desertification processes since 1982. While significant progress has been made in understanding the causes and consequences of desertification, following key questions remain unresolved: (i) can we predict spatial and temporal variation in ecosystem properties related to desertification and grass recovery? (ii) how do we integrate diverse observations about vegetation, climate, soils, hydrology, and animal populations to accomplish this prediction? As the desert ecosystems can often have highly variant weather and climate differences even among nearby locales (e.g., a storm cloud can produce significant rainfall in one place while leaving another 100 meters away completely dry), detailed monitoring and data gathering is useful for scientific advancement. Yet, such type of data gathering has historically been very expensive both in instrumentation cost and human effort. A network of low cost wireless sensor nodes can solve this problem.

Integration of diverse environmental and climatic data is the focus of current JRN LTER studies. A major component of this integration is the collection, analysis, and integration of data across multiple scales, from individual plant and inter-plant boundaries all the way to the complete Jornada Basin watershed. Our sensor deployment will accomplish multi-scale data collection by recording fine-grained weather and environmental data, such as soil moisture and temperature. The baseline data that can be used to develop new predictive models, to test models derived from other data, and to provide ground truth data that can be used to calibrate remote sensing equipment (such as satellite images). From this, the current predictive models for desertification will be extended to predict variation in ecosystem properties rather than just average properties.

#### **5. Publications**

During the two months since the CRI project began, one research paper has been submitted to the “3<sup>rd</sup> Annual IEEE Communications Society Conference on Sensor, Mesh and Ad hoc Communications and Networks (SECON 2006),” which is currently under review. The paper was titled “In situ coalescing of sensor data for obfuscating contextual information” and was authored by the PI and a graduate student. This paper addresses some theoretical issues related to thwarting traffic analysis and protecting contextual privacy of the sensing task.

#### **6. Outreach programs and supported students**

In Spring 2006, a senior project class (CS 448) was offered by the PI, Dr. Bhattacharya, which introduced Mica2, TelosB and Cricket mote hardware and TinyOS programming to the undergraduate students. Students worked in groups of two or three on sensor calibration, data aggregation, localization and traffic analysis over wireless channels. One graduate student is currently working on the CRI project during Summer 2006.

#### **7. Future and strategic directions**

At present, the three-year CRI project on networked sensor information technology at NMSU is at a very early stage. Establishing core capability and expertise in sensor design and sensor tasking is thus the current emphasis. Following is a subset of new research directions that we believe the infrastructure project will enable:

- Secured sensor service over the Internet
- Large-scale community sensor network through sensor-node virtualization
- Sensor network-based instrumentation for National Ecological Observatory Network (NEON)

# MII: Computing Support for the Next Generation Declarative Programming Systems

CNS-0454066

Year 2005-06

New Mexico State University

E. Pontelli, D. Ranjan, T.C. Son, K. Villaverde  
( epontell | dranjan | tson | kvillave ) @ cs.nmsu.edu  
<http://www.cs.nmsu.edu/lldap>

## 1. Introduction/Background

The objective of this proposal is the acquisition of new computing equipment to support three inter-related research projects currently ongoing in the *Knowledge representation, Logic, and Advanced Programming (KLAP)* laboratory. The activities of KLAP include active research, supervision and mentoring of graduate and undergraduate students, seminars and colloquium series open to faculty members and students (some of the most significant names in the field of logic and constraint programming have visited KLAP), and other educational activities—e.g., KLAP is currently managing the CS curriculum's courses on logic and constraint programming, AI, and parallel computing.

## 2. Goals, Objectives, Targeted Activities

The research conducted in the KLAP lab is focused on design, implementation, and application of declarative languages—especially logic-based languages – and knowledge representation techniques. The interest in the development of systems capable of autonomous reasoning led to recent advances in programming paradigms centered on the notion of *knowledge*. These paradigms strive to reduce a substantial part of the programming process to the description of objects comprising the domain of interest and relations between these objects. The resulting description can be queried to establish truth or falsity of statements about the domain, or to find objects satisfying various properties.

Knowledge can also be updated, modified, and used to make decisions on actions which need to be taken in a given situation to achieve certain goals. Though different in style, all these languages are based on *classical logic*, developed to model mathematical reasoning. This severely limits their ability to represent commonsense knowledge and reasoning crucial for development of sophisticated knowledge based systems. In the last ten years we witnessed a rapid development of alternative logical systems, called *non-monotonic logics*, which allow new axioms to retract existing theorems (more adequate for commonsense reasoning). The combination of traditional logic programming and non-monotonic reasoning provides a solid computational media for knowledge-based applications, and has recently established itself in domains such as planning, intelligent agents design, semantic web, and semantic-based filtering. By now we have solid logical foundations for such systems, but there is a general agreement that more is needed to make them practically useful. In particular, we need to (1) Design systems which strike a balance between expressive power and efficiency of inference engines; (2) Develop efficient implementations of such engines; and (3) Develop application-specific methodologies for representing knowledge and




Fig. 1: Project Components

applications in such domains. The goal of this proposal is to address these problems, through the research conducted in three inter-related projects. All the three projects rely on the use of a logic-based paradigm called *Answer Set Programming*, which provides the capabilities of traditional logic programming along with the ability to express non-monotonicity. We propose the use of *parallelism* and *distribution* to tackle the issue of efficiency. We study how adequate knowledge representation methodologies can be constructed in this framework, and used in the context of planning and Reasoning about Action and Change (RAC) in dynamic and complex domains. We propose to apply the resulting framework to develop the core part of a domain-specific computational workbench dedicated to evolutionary biologists—where reasoning about action and change is used to automatically determine the steps required to conduct a phylogenetic analysis process. As shown in Fig. 1, the three projects are genuinely inter-related and are mutually supportive.

### 3. Infrastructure and Facilities

The main objective of this grant is to renovate the computing facilities of the Knowledge representation, Logic, and Advanced Programming (KLAP) laboratory to address the needs of the highlighted research projects. In particular, the core component of the requested facilities is the acquisition of a new Beowulf cluster to support the ongoing research in parallel execution of logic-based programming frameworks.

To accomplish the development of the described infrastructure, we requested in the budget (and we hired) one undergraduate research assistant, to help with the selection, acquisition, installation, and configuration of the equipment. The person hired for this task is Mark Chernault; Mark has extensive experience in system and network administration, and he has also a sincere interest in parallel computing – indeed, he is expected to start a Master’s degree in Fall 2006 with focus on this area.

The new cluster has been ordered and it is scheduled to arrive by mid June; it is based on Infiniband (20Gbit/sec), with 5 nodes, where each node has 4 dual-core Pentium D (2.8GHz) processors. The cluster will come pre-packaged with all the necessary development and management software.

### 4. Project Description

The focus of research entertained is along three interconnected and mutually supportive directions. These directions are described as three projects: (i) Parallel and Distributed Execution of Logic Programming; (ii) Reasoning about Action and Change (RAC) in dynamic and complex domains; and (iii) Development of a Domain specific Workbench for Evolutionary Biology. These projects have the common thread of logic-based technology running through them. All three projects will use a Beowulf cluster of shared-memory multiprocessors as computing platform, as it provides a cost effective test-bed for experimenting with issues related to parallelism, distribution, and intercommunication.

Parallel and Distributed Logic Programming: The first project deals with the design and development of efficient parallel and distributed execution models for logic programming. We are interested in designing execution models that can dynamically adapt to the exploitation of parallelism on *Shared Memory (SMPs)*, *Distributed Memory (DMPs)*, and combinations of both types of platforms (i.e., clusters of SMPs). This project includes:

- Design and development of execution models for logic programming languages on parallel platforms—covering aspects such as memory management, communication, and task scheduling and distribution;
- Experimental evaluation of such technology through the use of real-life benchmarks and low-level performance analysis (e.g., through the use of simulators to study locality and cache behavior).

The target language for our investigation is *Answer Set Programming (ASP)*, a logic-based paradigm which combines traditional logic programming with features of constraint programming and non-monotonic reasoning. ASP has grown popular not only in the logic programming community but also in non-monotonic reasoning and semantic web reasoning.

Logic-based Planning with Incomplete Knowledge and Sensing Actions: The objective of this project is to develop logic-based methodologies to allow reasoning about action and change in complex domains. In particular, our focus is on (i) dealing with incomplete information and sensing actions, (ii) dealing with actions with duration and delayed effects, and (iii) dealing with non-binary fluents, resource consumption

and production. These are all aspects that have been recognized as mismatches between the expressive power of existing action languages and the properties of realistic planning domains. We propose to investigate these issues in the context of a logic-based action representation, with an easy to understand semantics, and by developing computational models that support planning in complex domain using such action representation. The parallel and distributed engines from the first project will provide the computational backbone for the experimentation in this project—the compute intensive nature of this type of computations will provide ample opportunities for the exploitation of parallelism. In turn, the peculiarities of the computations required for planning provide opportunities for alternative approaches to transparent extraction of parallelism (e.g., pipeline parallelism) that will be also explored.

*Domain-Specific Workbench for Evolutionary Biology:* This project deals with the development of a computational workbench to facilitate the rapid development and deployment of analysis processes by evolutionary biologists. The framework relies on two key ideas: (i) use of a visual *Domain Specific Language (DSL)*, which allows biologists to describe analysis processes using their own natural concepts and processes; (ii) use of an agent-based architecture, to automatically generate from each DSL program a *plan* to implement the analysis process—where the actions correspond to invocations of existing bioinformatics services; (iii) management of bioinformatics resources and tools as semantic web services, with automated composition and monitoring of services.

## **5. Indicators of Success and Major Accomplishments**

**Research Accomplishments:** the following research activities and accomplishments have been achieved during the Year 2005-06:

### *Parallel Logic Programming:*

- We have completed the design and implementation of a complete parallel inference engine for answer set programming. The engine is designed to run on a Beowulf cluster, and it has been validated on a large pool of complex benchmarks, with excellent performance results.
- We have performed a comprehensive study of the problem of task sharing during search parallelism; the study has identified dependencies between different strategies and properties of the execution, leading to the design of an adaptive task sharing strategy which applies the most appropriate task sharing methodology depending on the state of the execution. The hybrid strategy has been integrated in the current engine.
- We have initiated a similar investigation to related scheduling strategies to properties of the different benchmarks.

### *Reasoning about Actions and Change:*

- We have completed the design of the design of a new action language, called AM, for representing and reasoning about actions and change in domains with multi-valued fluents.
- We have developed a state-of-the-art conformant planner; the planner has been implemented in C and it has been shown to be competitive to existing technology. In particular, the planner implements a novel approximation of the entailment relation in presence of conditional plans.
- We have started investigating the interaction between conformant planning (as implemented in our system) and parallelism. We have identified two forms of parallelism that can be exploited during planning (a form of search parallelism and a form of deterministic parallelism). We have constructed a number of prototypes with excellent preliminary performance results.

### *Workbench for Phylogenetic Inference:*

- We have completed the implementation of a UDDI registry to handle bioinformatics web services.
- We have wrapped a number of bioinformatics tools and data repositories into web services and developed the corresponding WSDL descriptions.
- We have drafted the design of a visual data-flow language for the description of analysis processes in phylogenetic inference and the implementation of an editor and a compiler is in progress.
- We have designed a basic planning scheme to map high-level description of analysis processes to sequences of calls to bioinformatics web services.

## Software Products:

- We developed the first ever complete implementation of unrestricted aggregates in logic programming, as an extension of the Smodels answer set programming system.
- We developed a state-of-the-art conformant planner
- We completed the development of the first release of ASP-Prolog, a semantically well-founded integration of answer set programming and Prolog programming, and we demonstrated its use in the context of semantic web service composition and question-answer with deep reasoning.
- We developed a debugging system for answer set programming, based on the notion of *justifications*.
- We completed a prototype parallel answer set programming system capable of dynamically adapt task sharing strategies depending on the structure of the computation.

## 6. Publications/Patents

- T.C. Son, T. Phan, M. Gelfond, R. Morales. "Conformant Planning for Domains with Constraints." *National Conference on Artificial Intelligence AAAI*, pp. 1211--1216, 2005.
- T.C. Son, T. Phan, M. Gelfond, R. Morales. "An Approximation of Action Theories of AL and its Application to Conformant Planning." *International Conference on Logic Programming and Non-Monotonic Reasoning*, pp. 172 -- 184, 2005.
- T.C. Son, T. Phan, X. Zhang. "Reasoning about Sensing Actions in Domains with Multi-Valued Fluents." *Studia Logica* (2005) 79: 135--160.
- T.C. Son, C. Baral, T. Nam, S. McIlraith. "Domain-Dependent Knowledge in Answer Set Planning." Accepted for *TOCL*.
- T. Phan, T.C. Son, C. Baral. "Reasoning, Planning with Sensing Actions," Accepted for *TPLP*.
- M. Balduccini, E. Pontelli, O. Elkhatib, H. Le. "Issues in Parallel Execution of Non-monotonic Reasoning Systems." *Parallel Computing Journal*, 31:608--647, 2005.
- Dovier, E. Pontelli, G. Rossi. "Set Unification Revisited." *Theory and Practice of Logic Programming*, (to appear).
- T.C. Son and E. Pontelli. "Planning with Preferences using Logic Programming'." In *Theory and Practice of Logic Programming*, (to appear).
- Dal Palu, E. Pontelli, D. Ranjan. "Sequential and Parallel Algorithms for the NCA Problem on Pure Pointer Machines." In *Theoretical Computer Science*, 352:108--135, 2006.
- T.C. Son and E. Pontelli. "A constructive semantic Characterisation of aggregates in answer set programming", In *Theory and Practice of Logic Programming*, (to appear).
- E. Pontelli, K. Villaverde, H. Guo, G. Gupta. "PALS: Efficient Or-Parallel Execution of Prolog on Beowulf Clusters'." In *Theory and Practice of Logic Programming*, (to appear).
- E. Pontelli, K. Villaverde, H. Guo, G. Gupta. "A technique for efficient exploitation of search parallelism on share-nothing platforms." *Journal of Parallel and Distributed Computing* (to appear).
- Dal Palu, E. Pontelli, J. He, Y. Lu. "Identification of alpha-helices from Low Resolution Protein Density Maps." In *Computational Systems Bioinformatics*, Life Sciences Society, 2006.
- T.C. Son, E. Pontelli, T. Phan. "Answer Sets for Logic Programs with Arbitrary Abstract Constraint Atoms." In *National Conference on Artificial Intelligence*, AAAI/MIT Press, 2006.
- E. Pontelli and T. Son. "Justifications for Logic Programs under Answer Set Semantics." In *International Conference on Logic Programming*, Springer Verlag, 2006.
- O. Elkhatib, E. Pontelli, T.C. Son. "A Tool for Knowledge Base Integration and Querying." In *AAAI Spring Symposium*, AAAI Press, 2006.
- Dal Palu, E. Pontelli, J. He, Y. Lu. "A Constraint Logic Programming Approach to 3D Structure Determination for Large Protein Complexes." In *Symp. Applied Computing*, ACM Press, 2006.
- Y. Wu, E. Pontelli, D. Ranjan. "Computational Issues in Exploiting Dependent And-Parallelism in Logic Programming: Leftness Detection in Dynamic Search Trees." In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*, Springer Verlag, 2005.

- Dal Palu, A. Dovier, E. Pontelli. “A New Constraint Solver for 3-D Lattices and its Application to the Protein Folding Problem.” In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*, Springer Verlag, 2005.
- E. Saad and E. Pontelli. “Hybrid Probabilistic Logic Programs with Non-monotonic Negation.” In *International Conference on Logic Programming*, Springer Verlag, 2005.
- Dovier, A. Formisano, E. Pontelli. “A Comparison of CLP(FD) and ASP Solutions to NP-Complete Problems.” In *Int. Conference on Logic Programming*, Springer Verlag, 2005.
- O. Elkhatib, E. Pontelli, T.C. Son. “Justification and Debugging of Answer Set Programs in ASP-Prolog.” In *International Conference on Automated Debugging*, ACM Press, 2005.
- O. Elkhatib, E. Pontelli, T.C. Son. “Integrating an Answer Set Solver into Prolog.” In *Logic Programming and Non-monotonic Reasoning*, Springer Verlag, 2005.
- Elkabani, E. Pontelli, T.C. Son. “A System for Computing Answer Sets of Logic Programs with Aggregates.” In *Logic Progr. & Non-monotonic Reasoning*, Springer Verlag, 2005.
- H. Le and E. Pontelli. “An Investigation of Sharing Strategies for Answer Set Solvers and SAT Solvers.” In *EuroPar*, Springer Verlag, 2005.

## 7. Outreach Programs and Supported Students

The KLAP investigators have played a key role in the development of an educational program to recruit and sustain students, with a diverse background, throughout the gateway courses in the undergraduate program in Computer Science. The program has been highly effective in improving retention and progress. The pathways program is also supporting summer activities involving students from tribal colleges across the region and young women from local high schools. Resources of the KLAP lab are used in these events (graduate assistants and equipment).

The present grant has directly supported only one undergraduate student (Mark Chernault). Mark has been very effective in helping with the selection of the new cluster and he will lead the effort of installing and configuring it. Mark will transition to graduate studies in Fall 2006.

A number of graduate students have been indirectly supported by the grant, as they conduct research that is part of the described projects and they benefit from the equipment provided by the grant. During 2005-06, 7 doctoral students and 4 Master students participated in the research, and, of these, 1 doctoral student and 2 M.Sc. students completed their degrees.

## 8. Future and Strategic Directions

During the coming year, we plan to complete the acquisition of the equipment as indicated in the grant proposal. The cluster is expected to be functional by early Summer 2006, and the software currently developed will be ported on the new system. The additional equipment to be acquired includes a new printer and two multi-core servers.

The research project on parallel logic programming is expected to complete the study of scheduling techniques, and the result of this study will be integrated in the parallel engine. We also expect to modify the engine to accommodate for the presence of a hybrid cluster, applying different strategies (for scheduling and task sharing) depending on the location of the processors involved.

The research on planning will focus on developing new heuristics to speed-up the search of a plan; we will also continue the work on parallel planning, integrating the two different parallelization schemes within the same parallel engine and applying it to hard planning instances.

The research on the phylogenetic inference workbench will continue by implementing the editor and compiler for the visual language and implementing the underlying planner in charge of mapping high-level analysis processes to sequences of calls to web services.

# MII: Frameworks for the Development of Efficient and Scalable Knowledge-based Systems

CNS-0220590

Year 2005-2006

New Mexico State University

PI: Enrico Pontelli

Co-PIs: Roger Hartley, Hing Leung, Desh Ranjan, Son Cao Tran

( epontell | rth | hleung | dranjan | tson )@cs.nmsu.edu

## 1. Introduction/Background

In this report, we discuss the main components of the project funded by NSF MII grant CNS-0220590. In particular, we highlight the activities and some of the main accomplishments achieved during the Year 2005-06.

## 2. Goals, Objectives, Targeted Activities

The MII project is articulated in two inter-twined directions – a *research* direction and an *educational* direction. The overall objectives of the project are to sustain and grow the research capabilities of the Department of Computer Science, with particular emphasis on the development of collaborative research initiatives with underlying themes related to knowledge-based system, and to integrate the research efforts in a comprehensive educational plan, aimed at promoting participation of a diverse student population to Computer Science education and training.

The specific research and educational objectives of the project can be summarized as follows:

- **Research Objectives:** the theme of the research effort is the development of efficient and scalable knowledge-based systems; the specific objectives are: (1) Design of logic-based languages and software development methodologies for the construction and maintenance of knowledge-based systems; (2) Design of algorithms, data structures, and implementation methodologies for the creation of efficient and scalable knowledge-based systems; (3) Application of knowledge-based systems to different domains, including computational biology, semantic web, assistive technologies and sensor networks.
- **Educational Objectives:** the educational objective is the creation of a support infrastructure which promotes: (1) Recruitment of students with a diverse background; (2) tutoring and mentoring of incoming students to prepare them for success in the undergraduate CS program; (3) advanced training through courses and research experiences; (4) retention of students throughout the undergraduate program; and (5) improved graduation and transition to graduate studies

Emphasis is given to promote participation of Hispanic and Native American students to training and research.

## 3. Infrastructure and Facilities

**The Teaching Laboratory:** we have completed the construction of our teaching laboratory. The laboratory, housed in a new dedicated room, is equipped with 35 new Dell laptops, installed on movable stations, 3 computer projectors, a laser printer, a dedicate teacher stations (with a PC and 3 LCD screens), cameras, microphones and audio system for delivery of on-line courses, and a smart white-board. The room is covered by two a/b/g wireless access points. The room has already been employed for teaching various regular classes and participation in remote courses (e.g., a course on life science ethics from Iowa State University).

**Network Infrastructure:** we have continued the development and improvement of the Department network infrastructure. Gigabit Ethernet has been extended to all the labs, the teaching laboratory, the machine rooms, and many of the faculty offices. The network has been extended to reach the second floor of Science Hall, which now houses the Center for Research Excellence in Bioinformatics. Wireless network (a/b/g) has been extended to cover both the department as well as the classrooms wing.

**Beowulf Cluster:** although our budget included the acquisition of a new Beowulf cluster for the 2005/06 year, we decided to postpone this step to the next year. Thanks to another grant, we have ordered a smaller Infiniband-based cluster (expected to arrive at the end of May or early June) and we wish to evaluate such configuration before proceeding with the acquisition of a bigger cluster. In the meantime, to support our work on parallel processing, we have acquired two new Sun Servers, each with 6-core Sparc processors.

**Miscellaneous:** to support the new focus on sensor networks, we have contributed to the purchase of MICA motes and software (Mathematica).

## 4. Project Description

The research conducted under the auspices of this grant is focused on the development of a framework for building efficient and scalable knowledge-based systems. The research has been articulated according to the following main research threads: (1) *Languages and Programming Methodologies*, (2) *Efficiency and Scalability*, (3) *Applications*.

### **Languages and Methodologies:**

*Reasoning about Actions and Change*: one of the key areas of knowledge representation explored in this project is reasoning about action and change, with its application to problem domains such as planning, diagnosis, and argumentation. In particular, our interest is in reasoning with incomplete knowledge, sensing actions, multi-valued fluents, and complex/dynamic domains.

*Logic Programming*: our paradigm of choice throughout most of this project is logic programming, in its different incarnations (e.g., Prolog, constraint logic programming, and answer set programming). As part of this project, we investigate extensions of these paradigms to better fit the needs of knowledge-based application domains, as well as integration between different logic-based paradigms and other programming paradigms.

*Component-based Programming*: the research project relies on the notion of components for the composition of software elements and the reliable upgrading and evolution of software systems.

### **Efficiency and Scalability**

*Complexity and Data Structures*: this component of our research project is aimed at the development of formal abstractions of the various components of the execution models of the previously considered languages, the investigation of their computational properties, and the development of optimal data structures and algorithms.

*Parallel Execution Models*: data structures and algorithms are designed and implemented to support efficient parallel execution of the investigated languages and systems.

*High-Level Parallelism*: instances of parallelism are detected and exploited from specific applications in the context of reasoning about action and change (e.g., parallel planning).

### **Applications**

*Assistive Technologies*: we investigate the use of logic programming and knowledge representation techniques to develop tools to assist visually disabled individuals in effectively accessing complex on-line documents (e.g., Web pages and documents with mathematical content).

*Bioinformatics*: we investigate the use of constraint and logic programming to address different problems in the field of bioinformatics, such as bioinformatics service composition and protein structure determination.

*Semantic Web*: we apply knowledge representation methodologies and logic programming to address aspects of the semantic web, such as service composition and monitoring, and management of rule-based knowledge.

## 5. Indicators of Success and Major Accomplishments

*Reasoning about Actions and Change*:

- Development of a semantic approximation for conformant planning in presence of incomplete knowledge.
- Design and implementation of a new conformant planner, which significantly outperforms existing planners on various planning problems.
- Refinement of a language for preferences in action theories and implementation in constraint programming.

*Logic Programming*:

- We have completely revisited the semantics of logic programming with aggregates, advancing the state-of-the-art over previous proposals.
- We have developed a complete implementation of answer set programming with aggregates.
- We have extended the work on aggregates to the more general case of logic programs with abstract constraint atoms, leading to the first general semantics for logic programs with unrestricted abstract constraint atoms.

*Component-based Programming*:

- We have continued the development of the Dynamic Dynamic Loader, a modification of the GNU dynamic loader, to make it into an event-driven framework, which allows tool developers to “listen” to dynamic link resolution and redirect function calls to different symbols. The framework is used to implement HerculesC++, a reliable updating framework for C++.

*Complexity and Data Structures*:

- We have continued the investigation of a formal abstraction of and-parallel execution of constraint and logic programs, identifying lower bounds and efficient data structures for its resolution.
- Our work on pointer machines has continued, exploring the relationships between pointer machines and Kolmogorov complexity.

#### *Parallel Execution Models:*

- We focused on the investigation of methods to integrate different schemes for task sharing in search parallelism; this is required, as our studies have highlighted that none of the individual schemes has an optimal behavior on all types of programs. We have designed an algorithm and heuristics to allow the system to dynamically switch between task sharing schemes, depending on the structure of the computation tree.

#### *High-Level Parallelism:*

- We have initiated the investigation of schemes to parallelize forward planning algorithms, and applied them to the efficient conformant planner constructed by our group. We have studied the exploitation of two forms of parallelism – *search* parallelism (across the choice of different actions at each time step) and *next-state* parallelism (where we parallelize the computation of the successors of a given state of a trajectory).

#### *Assistive Technologies:*

- We have continued our study of how to use machine learning to improve understanding of the structure of a web page, and to use such knowledge to guide the non-visual presentation of documents. We have obtained excellent results in the context of handling of HTML forms and to rank document components according to relevance to the user's interest. We have started investigating the use of RSS feeds as a collection of information useful to provide aural summaries and help in aurally accessing specific types of web sites.

#### *Bioinformatics:*

- We have completed the development of a constraint solver operating with rigid structures in a 3d space, organized as a crystal lattice. The solver is efficient, capable of distributing executions on Beowulf clusters, and it has been applied to efficiently solve the protein structure determination problem (in this approximated version of the 3d space).
- We have developed a tool, using constraint logic programming, which applies constraints obtained from analysis of protein density maps to the primary sequence of the protein, in order to guide prediction of secondary structure components.

#### *Semantic Web:*

- We have investigated ways to adopt our integrated logic programming framework (which combines Prolog and Answer Set Programming) to compose semantic web knowledge-bases expressed using different languages from the RuleML family.

## **6. Publications/Patents**

### **Selected Publications:**

- M. Balduccini, E. Pontelli, O. Elkhatib, H. Le. "Issues in Parallel Execution of Non-monotonic Reasoning Systems", *Parallel Computing Journal*, 31:608-647, 2005.
- J.E. Cook and Z. Du. "Discovering Thread Interactions in a Concurrent System." In *Journal of Systems and Software*, pp. 285--297, vol. 77, no. 3, 2005.
- A. Dovier, E. Pontelli, G. Rossi. "Set Unification Revisited", *Theory & Practice of Logic Progr.*, (to appear).
- T.C. Son and E. Pontelli. "Planning with Preferences using Logic Programming", *Theory & Practice of Logic Progr.*, (to appear).
- S. Tambe, N. Vedagiri, N. Abbas, and J. Cook, "DDL: Extending Dynamic Linking for Program Customization, Analysis, and Evolution." In *Int. Conference on Software Maintenance*, 2005.
- Dal Palu, E. Pontelli, D. Ranjan. "Sequential and Parallel Algorithms for the NCA Problem on Pure Pointer Machines", In *Theoretical Computer Science*, 352:108-135, 2006.
- E. Pontelli, K. Villaverde, H. Guo, G. Gupta. "Stack Splitting: a technique for efficient exploitation of search parallelism on share-nothing platforms", In *Journal of Parallel and Distributed Computing* (to appear).
- Dal Palu, E. Pontelli, J. He, Y. Lu. "Identification of alpha-helices from Low Resolution Protein Density Maps", In *Computational Systems Bioinformatics (CSB)*, Life Sciences Society, 2006.
- T.C. Son, E. Pontelli, T. Phan. "Answer Sets for Logic Programs with Arbitrary Abstract Constraint Atoms", In *National Conference on Artificial Intelligence (AAAI)*, AAAI/MIT Press, 2006.
- E. Pontelli and T. Son. "Justifications for Logic Programs under Answer Set Semantics", In *International Conference on Logic Programming*, Springer Verlag, 2006.
- Dal Palu, E. Pontelli, J. He, Y. Lu. "A Constraint Logic Programming Approach to 3D Structure Determination for Large Protein Complexes", In *ACM Symp, on Applied Computing (Bioinformatics Track)*, ACM Press, 2006.
- Y. Wu, E. Pontelli, D. Ranjan. "Computational Issues in Exploiting Dependent And-Parallelism in Logic Programming: Leftness Detection in Dynamic Search Trees", In *12th International Conference on Logic for Programming Artificial Intelligence and Reasoning*, Springer Verlag, 2005.

- E. Saad and E. Pontelli. "Hybrid Probabilistic Logic Programs with Non-monotonic Negation", In *International Conference on Logic Programming*, Springer Verlag, 2005.
- O. Elkhatib, E. Pontelli, T.C. Son. "Justification and Debugging of Answer Set Programs in ASP-Prolog", In *ACM International Conference on Automated Debugging (AADEBUG)*, ACM Press, 2005.
- H. Le and E. Pontelli. "An Investigation of Sharing Strategies for Answer Set Solvers and SAT Solvers", In *EuroPar*, Springer Verlag, 2005.
- E. Pontelli. "Towards Automated Web Pages Structure Discovery", In *HCI International*, 2005.
- T. Son, C. Baral, T. Nam, S. McIlraith. "Domain-Dependent Knowledge in Answer Set Planning." In *ACM Transactions on Computational Logic*, (to appear).

## 7. Outreach Programs and Supported Students



The project has supported a variety of outreach efforts.

**Pathways in Computer Science:** this program, created thanks to the funding of the MII grant, provides educational, financial, and mentoring to Computer Science freshmen and sophomore students, from the entrance in the undergraduate program to the completion of the gateway courses. The objective of the pathways program is to provide a set of tools to provide personalized assistance to students in the gateway courses; these tools meet the following properties

- They are *uniform* across different courses.
- They offer different levels of assistance, depending on the background and level of preparation of the student.

The program started by targeting two gateway courses (intro to programming and discrete math) and it has been extended

during the last year to three more courses (intro to data structures, algorithms, and object oriented programming). Student input was sought during the development process. This input highlighted that students have many different concerns (i.e., academic, peers, family, work, cultural, and financial), all of which have an impact on the educational process, but some of which are not normally considered by support programs. This resulted in a program that uses a team-oriented framework, offers academic support and challenges in individual and group formats, creates employment opportunities, and takes advantage of faculty and peer mentoring.

While this is still the core of the program, it has been expanded to offer tools which are expected to lead to the long-term success of our students. The following main activities are currently included in the Pathways program:

- *Educational assistance:* this includes course-specific focus groups, peer-tutoring, time-independent self-support (web-based), and access to libraries of problems and reading material.
- *Mentoring:* peer mentors are provided to students that are identified by instructors and advisors as at-risk; mentors are trained to advise students on course-independent educational matters and to refer students to the appropriate university offices for personal problems.
- *Workshops:* offered by faculty members and graduate students, they deal with general topics (e.g., how to prepare for exams) and specific technologies (e.g., introduction to specific programs, like Eclipse and Netbeans).
- *Social events:* designed to build bonds among students and to network with faculty and graduate assistants.
- *Employment opportunities:* undergraduate students have been hired as tutors, peer-mentors, and web designers; some of the students employed are students that participated in the pathways program themselves.

The pathways program assists about 35 students per semester. Statistics have indicated a significant reduction in drop-out rates and an improvement in course grades.

**Summer Camps:** summer camps are offered to provide students with exposure to campus life at NMSU and to provide hands-on experience in various aspects of Computer Science. In Summer'05, we offered two camps:

- a 5-week residential camp, open to 2-year college students from various tribal colleges across the State, and providing education and practice in Computer Graphics, Bioinformatics, Web programming, and fundamentals of computer programming (in Java and Linux). The students receive college credits for the computer programming class. The camp was attended by 12 students.
- A 2-week non-residential camp, open to students from local high-schools. The camp is focused exclusively on bioinformatics, and it is composed of an educational part – exposing students to basic notions of molecular biology and tools for bioinformatics – and a research part – where students work in teams in solving a research task, and develop and present a research poster. 17 students attended the 2005 camp.

The 2006 summer camps will be organized along the same lines; at this moment, 10 students have been accepted for the 5-week camp, and 15 for the high-school camp.

**Scholarships:** we have developed a scholarship program opened to excellent students from traditionally under-represented backgrounds. The first scholarship has been awarded to a Native American student.

**Research Experiences:** two undergraduate students (S. Jain and A. Teufel) participated in research activities as part of two research projects (dealing with the bioinformatics application area of our research project).

**Young Women in Computing Program:** the investigators have been recently awarded a supplement to the MII grant, part of the Broadening Participation in Computing initiative. The supplement will be used to create a demonstration program targeting participation of young women in Computer Science. The program, targeting high school women, includes a 5-week summer camp (providing an introduction to programming using Alice, Robot programming using Lego Mindstorm II, and fundamentals of bioinformatics), summer workshops (e.g., library research, self-defense, study and note-taking skills), and monthly luncheons during the academic year – with social activities to network with faculty and students, and invited speakers (e.g., successful women researchers). A 12-student cohort has been recruited for this program, which will start in July 2006.

#### **Other Activities and Accomplishments:**

- Degrees awarded to participating graduate students: 2 Ph.D., 29 M.Sc.
- Three undergraduate students have been involved in research activities.
- New research-oriented courses: Planning and Reasoning about Actions and Change (Fall'05), Non-monotonic Reasoning (Spring'06), Advanced Software Engineering (Spring'06); we have also completed the establishment of a new core course on object-oriented programming.
- We have established collaborations with local industries to secure topics for the required Undergraduate Senior Capstone course (a project from Harris Corporation was used in Fall'05).
- Work is still in progress to develop articulation agreements with various 2-year colleges with significant Native American presence (e.g., Northern NM College and Crownpoint Institute of Technology). As part of this effort, the MII program will host during Summer 2006 a workshop of Computer Science instructors from tribal colleges across the nation, to discuss curriculum issues and educational methodologies.
- During Summer 2005, investigators have mentored 1 Native American student, part of summer bridges program.
- During Summer 2006, we will host one instructor from a major Native American college (David Broderick from Dine College in Tsaile, AZ) to investigate curriculum issues and integration of research and education.
- We have extended our recruitment efforts, including more regular visits to tribal colleges across the State (conducted by faculty members and by our program coordinator, Ms. Sandoval) and local high schools. We have established (leveraging on another grant) a scholarship program at the local NMSU Dona Ana Community College (8 students have been awarded in 2005-06), with the aim of establishing contacts with students that have the potential to transfer to a 4-year program in Computer Science.
- In the last year, we have successfully recruited 3 minority graduate students.
- The MII investigators have actively participated in promoting a consortium of Computer Science programs in Hispanic-serving institutions; the consortium has been recently awarded a BPC alliance grant.

## **8. Future and Strategic Directions**

**Research Directions:** The individual research components have clear directions along which they will progress. We propose to extend the integration between the different components of the project; in particular, we intend to make use of the methodologies developed for parallel execution of logic programming in the context of parallelization of reasoning about actions and change; we plan to make use of the novel inference engines (with aggregates, parallelism, modules) to support the needs of the other projects. We also plan to combine the work on component-based programming with the work on integration of logic programming frameworks, to create a more general programming framework; we propose to demonstrate its effectiveness in the context of composition and monitoring of semantic web services.

**Educational Directions:** New courses with high research content will be developed; a course on Semantic Web will be introduced in Fall 2006; an advanced course on parallel and distributed processing is in the planning for Spring 2007. Work is in progress to introduce team-based learning methodologies throughout the undergraduate program (following the affinity group model and with the help of the UTEP team).

**Outreach Directions:** The scope of the pathways program will be expanded over the next year to cover the majority of the core undergraduate courses, to include additional web-based material, to include links to existing on-line relevant courses, and to address the final step of facilitating transition to the graduate program. The Young Women in Computer Science program will be implemented for the first time during the 2006/07 year; significant effort will be invested to create links and interaction between the different outreach programs (e.g., sharing activities during the summer programs). Further work will be invested towards program articulation with 2-year tribal colleges.

# **CRI: ADVANCED SENSOR NETWORK APPLICATIONS FOR ENVIRONMENTAL MONITORING SYSTEMS**

Proposal No.: CNS-0454259  
First Year

John Kelly, Associate Professor, Electrical & Computer Engineering, [jck@ncat.edu](mailto:jck@ncat.edu)  
Jung H. Kim, Professor, Electrical & Computer Engineering, [kim@ncat.edu](mailto:kim@ncat.edu)  
Winser E. Alexander, Visiting Professor, Electrical Engineering, [winser@ncsu.edu](mailto:winser@ncsu.edu)  
Sung H. Yoon, Associate Professor, Computer Science, [syoon@ncat.edu](mailto:syoon@ncat.edu)  
Albert Esterline, Associate Professor, Computer Science, [esterlin@ncat.edu](mailto:esterlin@ncat.edu)  
Kenneth Williams, Associate Professor, Computer Science, [Williams@ncat.edu](mailto:Williams@ncat.edu)  
Dr. JungHee Kim, Assistant Professor, Computer Science, [jungkim@ncat.edu](mailto:jungkim@ncat.edu)  
Dr. G.B. Reddy, Professor, Natural Resources & Environmental Design, [reddyg@ncat.edu](mailto:reddyg@ncat.edu)  
Dr. Manuel Reyes, Associate Professor, Natural Resources & Environmental Design, [reyes@ncat.edu](mailto:reyes@ncat.edu)

North Carolina A&T State University  
Greensboro, NC 27411  
[www.ncat.edu](http://www.ncat.edu)

## **1. ABSTRACT**

The Departments of Electrical and Computer Engineering, Computer Science and Natural Resources & Environmental Design at North Carolina A&T State University are developing a research infrastructure for advanced sensor networks applications for environmental monitoring systems to provide experimental tools for graduate student research as well as for undergraduate student projects. We define an environmental monitoring system in this project as hardware and software systems to monitor the water quality in a watershed by collecting data from various kinds of environmental sensors through a wireless sensor networks. The proposed infrastructure will support research on wireless sensor networks, data collection for environmental monitoring, data mining and modeling related to water quality management.

## **2. INTRODUCTION AND BACKGROUND**

Non-point source pollution is a major cause of water quality degradation. Non-point source pollution comes from land runoff, atmospheric deposition, drainage, seepage, and hydrologic modification. The Environmental Protection Agency (EPA) has made significant progress in implementing Clean Water Act (CWA) programs and in reducing water pollution. Despite such progress, however, serious water quality problems persist throughout the country and agricultural operations are considered a significant source of water pollution in the United States. Even in North Carolina, 477 water bodies were identified as impaired and therefore, pollution sources and Total Maximum Daily Load (TMDL) protocols must be developed to restore the water quality in these water bodies. The pollution parameters must be measured on an hourly basis if possible and the data needs to be managed and modeled for such restoration processes. In order to accomplish this task, a wireless sensor network needs to be developed to monitor the parameters.

Currently data on measured parameters is being used for a study on watershed assessment in North Carolina. However, there is no timely way to automatically obtain measured parameter data. Therefore, there is a vital need to monitor water quality parameters in streams and wetlands. We plan to place sensors in streams and wetlands to monitor water quality parameters as a part of this infrastructure. Therefore, the main purpose of this project is to develop an infrastructure to support research on the

development of a wireless sensor networks for environmental monitoring with emphasis on water quality management.

### 3. INFRASTRUCTURE AND FACILITIES

The overall infrastructure will consist of indoor laboratories and outdoor testbeds. The indoor Laboratories include Intelligent Mobile Information Systems (IMIS) Lab in Interdisciplinary Research Center (IRC), and Image Processing and Computer Vision (IPCV) Lab in Cherry Hall.

We plan to develop an additional indoor laboratory in new space and it will be called the Sensor Network Systems Laboratory (SNSL). The SNSL will be used to construct a wireless sensor network environment. The lab will be equipped with many sensor nodes. Also, we plan to construct an outdoor testbed at a weather station located near the Department of Natural Resources and Environmental Design for field tests. Rugged wireless sensor nodes with environmental sensors will be installed in this outdoor testbed. The major hardware and software purchased during the reporting period are as follows:

(1) ECONet Weather Station:

- Including Data Logger, Modem, Air Temperature & Relative Humidity, Barometric Pressure, All-season Heated Precipitation Sensor, Soil Temperature, Soil Moisture, Solar Radiation Sensor, Photo-synthetically Active Radiation, Wind Monitor, Evapotranspiration, Power Supply, Tower and Accessories, etc.

(2) Crossbow Wireless Sensor Network Modules:

- Including MEP Mote Environmental System including microclimate sensor nodes, mote base station and Mica2 processor and Radio Board, Advanced User's Programming Dongle Kit for MEP Products, MSP Mote Developer's System including Mote Security packages, Mote Base Station, Mica2 Processor and Radio Board, MSP Programming Board, Mica2/Micaz Data Acquisition Board, Mica2/Micaz Sensor Board (Light, Humidity, Barometric Pressure, Seismic), Mica2, Micaz, Mica2dot PC Interface Board, Advanced Developer's Kit including Basic Contents, USB Webcam and WIFI Card, Stargate Daughter Card, IEEE 802.15.4 TelosB Mote with Temperature, Humidity and Light Sensors

(3) Campbell Environmental Monitoring Data Logger with Sensors:

- Including Micrologger CR3000, 16 Channel Relay Multiplexer, Loggernet, with Software, Temperature/PH Probe, RM Wind Monitor, etc.

(4) Intellweather Vantage Pro2 Weather Station:

- Including Davis Anemometer, Davis Rain Collector

(5) Desktop PC and Notebooks

### 4. PROJECT DESCRIPTION

The following research activities have been conducted during the reporting period.

**(1) Applications Layer Research – Environmental Monitoring (Dr. G.B. Reddy and Dr. Manuel Reyes)**

The School of Agriculture and Environmental Sciences (SAES) at North Carolina A&T State University has an initiative to develop a Soil and Water Quality document. Small-scale and limited resource farmers in North Carolina are the primary audience for this initiative. Specific programmatic focus areas include: Agricultural Best Management Practices, Non-Agricultural Best Management Practices, and Academic Preparation and Professional Development. The development and implementation of Total Maximum Daily Loads (TMDL) protocols for selected impaired water bodies in North Carolina is an integral component of these program areas. TMDL development requires the identification of agricultural and urban pollutants that impair a water body. These can be achieved by measuring pollutant loads in water bodies or by using watershed based computer models that simulate pollutant loads in a water body.

**(2) Treatment of swine wastewater with appropriate cost-effective technologies (Dr. G.B. Reddy and Dr. Manuel Reyes)**

North Carolina ranks second in U.S. in swine production. Apart from the swine production and economic gains for the state, North Carolina is burdened with the swine waste problems. Traditionally swine operations pump the washed wastewater from the swine houses into an anaerobic lagoon and spread the wastewater from lagoons on the agronomic fields. However, North Carolina legislature introduced a moratorium to ban the lagoon usage for new operations. We have constructed wetlands which are cost-effective, and use simple and passive technology to treat swine wastewater. We constructed six wetland cells (11 m wide x 40 m long) in 1996 and started the experiments beginning in 1997 to remove nitrogen, phosphorus, solids, chemical oxygen demand (COD), and pathogens. We plan to use the proposed environmental monitoring system to automatically collect data for the research on the use of these wetlands for treatment of swine waste water. Sensors of interest include the dissolved oxygen probe, the redox potential probe, the temperature probe and the pH probe.

**(3) Sensor Networks for environmental monitoring (Dr. Sung Yoon and Dr. Jung Kim)**

We have been developing a sensor network system that can be used for environmental monitoring. We are currently developing sensor networks that support multi-hopping capabilities. Monitoring software is written by MS visual C++.Net and Database is being built on a SQL server. Each mote in the proposed sensor network will collect data through an interface to on-board sensors or external sensors. A small scale of the sensor network system will be implemented and tested. Once the testing is successful, then a sensor network infrastructure for environmental monitoring will be developed in collaboration with the School of Agriculture and Environmental Sciences.

**(4) Localization using Sensor networks (Dr. Sung Yoon and Dr. Jung Kim)**

Knowledge about the locations of individual nodes can be a useful primitive in many routing protocols, it can answer the questions on the network coverage and it can assist the group querying of sensors. When reporting about the origins of events in network area or when network is used to the tracking of moving target, the node location information is essential. Location information is also fundamental part in many practical sensor network solutions, like smart highway, smart shopping mall, smart battlefield, industrial applications in logistics and robotics, or in medical applications. We are developing localization algorithms for sensor networks. We are using Crickets from Crossbow that can be used as the prototype of the localization systems. Cricket provides location information to attached host devices using RF and ultrasound accurately. Using these devices with sensor nodes, we can evaluate our routing localization algorithms and routing algorithms.

**(5) Fuzzy Dynamic Logic for Means-Ends Ascription and Adaptation and Learning for Tactical Behaviors (Dr. A. Esterline)**

The wireless network infrastructure will be used for research on Means-ends reasoning and learning classifier systems involving multiagent learning by Dr. Esterline.

**5. INDICATORS OF SUCCESS AND MAJOR ACCOMPLISHMENTS**

- (1) Nine faculty members in three departments (Electrical & Computer Engineering, Computer Science and Natural Resources & Environmental Design) have been participating in this project. During the funding period, 10 research activities have been undertaken, and 8 doctoral students, 13 master students and 1 undergraduate student have participated. Among the students participating in this program, 1 Ph.D. student graduated during the reporting period, 1 MS graduate student has completed their masters program and 1 undergraduate student has graduated with research experience under the supervision of faculty members. Three journal papers are in review for publication and three journal papers were published. 17 papers have been published in the regional, national and international conference proceedings during the reporting period.

- (2) A state of the art agrometeorological station (AS) was installed at the North Carolina A&T State University Farm. The AS is part of the NC Environment and Climate Observing Network (NC ECONet), which total 26 monitoring sites across North Carolina. The AS automatically measures the following agrometeorological variables: air temperature & relative humidity at 10 meters and 2 meters, barometric Pressure at 2 meters, precipitation at 1 meter, soil temperature at 6 levels below surface (down to 1 meter), soil moisture at 6 levels below surface (down to 1 meter), solar radiation at 2 meters, photosynthetically active radiation at 2 meters, wind direction and speed at 10 meters and at 2 meters, and evapotranspiration. Data can be accessed at



<http://www.nc-climate.ncsu.edu/cronos/index.php?station=NCAT>

## 6. PUBLICATIONS

### □ *Journal Papers:*

- [1] Ji-Hyun Lee , Sung H. Yoon , Jung H. Kim, “Region Scalable Image Coding Using a Post-Segmentation Approach”, International Journal of Computers and Their Applications, In review.
- [2] Sung Yoon, Hee Y, Choi, Jung H. Kim, "Efficient Neural network face detector for face recognition applications," Pattern Recognition, in Review.
- [3] Homaifar, A., Hawari, H., Esterline, A. C., Iran-Nejad, A., and Tunstel, E., “Soft Computing for Agent-Based Decision Making Using the Biofunctional Theory of Knowledge,” submitted to IEEE Journal of Fuzzy sets.
- [4] Hughes, J., Esterline, A. and Kimiaghalam, B., “Means-End Relations and a Measure of Efficacy,” Journal of Logic, Language and Information, Vol. 15, No. 1, 2006.
- [5] Kimiaghalam, B., Homaifar, A., and Esterline A.C., “A Statechart Framework for Agent Roles that Captures Expertise and Learns Improved Behavior," published in Lecture Notes in Computer Science, Formal Approaches to Agent-Based Systems, Springer Berlin / Heidelberg, Volume 2699 / 2002, 2006.
- [6] Hunt, P.G., T. Matheny, M. Poach, G. B. Reddy and K. Stone, “Denitrification in Marsh-Pond-Marsh Constructed Wetlands Treating Swine Wastewater at Different N Loading Rates”, Soil. Sci. Soc. Am. J. 70: 487-493, 2006.

### □ *Book Chapters:*

- [1] Esterline, A.C., BouSaba, C., Homaifar, A., and Piore, B., “Hierarchies, Holons, and Agent Coordination,” accepted by Springer-Verlag as a chapter in Volume 3825 of LNCS/LNAI (ed. Michael G Hinchey et al.), papers based on presentations at the 2<sup>nd</sup> Workshop on Radical Agent Concepts (WRAC II), Sept. 20–22, 2005, Greenbelt, MD
- [2] Esterline, A., Gandluri, B., and Sundaresan, M., “Characterizing Environmental Information for Monitoring Agents,” accepted by Springer-Verlag as a chapter in Volume 3825 of LNCS/LNAI (ed. Michael G Hinchey et al.), papers based on presentations at the 2<sup>nd</sup> Workshop on Radical Agent Concepts (WRAC II), Sept. 20–22, 2005, Greenbelt, MD
- [3] Esterline, A. C., “Using Statecharts and Modal Logics to Model Multiagent Plans and Transactions,” in Hinchey, M.G. et al. (eds.), Formal Approaches to Agent-Based Systems, Springer, LNAI/LNCS Vol. 2699, 2006.
- [4] Esterline, A., Rorie, T., and Homaifar, A., “A Process-Algebraic Agent Abstraction. In Rouff, C. A. et al. (eds.), Agent Technology from a Formal Perspective. Springer, 2006, pp. 99-137.

### □ *Conference Proceedings:*

- [1] Gab C. Jung, Seong M. Park and Jung H. Kim, “Efficient VLSI Architectures for Convolution and Lifting Based 2-D Discrete Wavelet Transform”, Proceedings of Tenth Asia-Pacific Computer Systems Architecture Conference (ACSAC05), Singapore, October 24-26, 2005.
- [2] Ji Lee, Sung Yoon, Jung H. Kim, Gi T. Hur, “Panoramic Image Coding for Virtual Reality Applications,” The 9<sup>th</sup> World Multi-Conference on Systemics, Cybernetics and Informatics, Orlando, FL, July 10-13, 2005.

- [3] Sung H. Yoon, Giyeon Park, Jihyun Lee, Jung H. Kim: “An Annealed Neural Network For Reliable Face Detection”, CISST 2006, Las Vegas, NV, June 27-39, 2006.
- [4] Sung H. Yoon, Gi-yeon Park, Gi T. Hur, Jung H. Kim: “Face detection and recognition using geometrical features and a neural network verifier”, SPIE Defense and Security 2005, Orlando, FL, April, 2006.
- [5] Sung H. Yoon, Gi T. Hur, Jung H. Kim, “Recurrent Neural Network Verifier for Face Detection and Tracking”, AIE/IEA 2006, France, June 23-30, 2006.
- [6] Tyesia Pompey, Cranos Williams, Jung Kim, and Winsor Alexander, “Analyzing the Reliability of ICA Estimates of Bacillus subtilis”, Proceedings of the IEEE International Workshop on Genomic Signal Processing and Statistics, 2005.
- [7] Ramsey Hourani, Ravi Jenkal, W. Rhett Davis and Winsor Alexander, “Tool Integration for Signal Processing Architectural Exploration”, Proceedings of the 2006 Electronic Design Processes Workshop, April 2006.
- [8] Ramsey Hourani, Winsor Alexander and T. Raithatha, “A Hardware Performance Analysis Framework for Architectural Exploration of DSP Systems”, Proceedings of the Global Signal Processing Expo (GSPx), October 2005.
- [9] William Edmonson, Ruchir Gupte, Jaya Ginchandandi, Senanu Ocloo and Winsor Alexander, “Interval Arithmetic Logic Unit for Signal Processing and Control Applications”, Workshop on Reliable Engineering Computing, February 2006.
- [10] Esterline, A., Gandhluri, B., and Sundaresan, M.J., “Development and Verification of Multiagent Systems for Vehicle Health Management,” Proceedings of the 2005 Integrated Systems Health Management (ISHM) Conference organized by AFRL, Cincinnati, OH, August 8-11, 2005.
- [11] Reddy, B.B.K., Homaifar, A., and Esterline, A.C., “Minimal Fuel Consumption of Electric Propulsion Space Vehicles for Deep Space Exploration”, IEEE Aerospace Conference, Big Sky, MT, March 2006.
- [12] Reddy, B.B.K., Esterline, A.C., and Homaifar, A., “Genetic Algorithms for Minimal Fuel Consumption of Electric Propulsion Space Vehicles”, International Conference on Computational Intelligence for Modelling, Control and Automation (CIMCA), Austria, November 2005.
- [13] BouSaba, B., Esterline, A.C., Homaifar, A., Rodgers, D. and Huber, J., “Statecharts in Non-Markov Environments For Multi-Agent Learning,” MARS 06 (Multi-Agent Robotic Systems Conf.), Setubal, Portugal, Aug. 2006 (Work in Progress).
- [14] BouSaba, B., Esterline, A.C., Homaifar, A., and Rodgers, D., “Formal, Holarchical Representation of Tactical Behaviors,” IEEE SMC 05 (Int. Conference on Systems, Man and Cybernetics), Hawaii, October 2005.
- [15] BouSaba, B., Esterline, A.C., Homaifar, A., and Rodgers, D., “Learning Coordinated Behavior: XCSs and Statecharts,” IEEE SMC 05 (Int. Conf. on Systems, Man and Cybernetics), Hawaii, October 2005.
- [16] Homaifar, A., H. Hawari, C. W. Bou-Saba, A. Esterline, A. Iran-Nejad, E. Tunstel, “Soft Computing for Agent-Based Decision Making Using the Biofunctional Theory of Knowledge,” IEEE SMC 05, special session on Soft Computing in Single & Multi-Agent Learning Systems, IEEE SMC 05 (Int. Conf. on Systems, Man and Cybernetics), Hawaii, October 2005.
- [17] Esterline, A., “A Rigorous Self-Configuring Multiagent Architecture for Vehicle Health Monitoring,” accepted for poster presentation at the Integrated Systems Health Management (ISHM) Conference, Cincinnati, OH, August 2006.

## **8. OUTREACH PROGRAM AND SUPPORTED STUDENTS**

About 3000 K-12 kids visit the A&T farm each spring. A part of their visit may include the AS. In terms of size, the A&T farm is the biggest laboratory in Greensboro.

## **9. FUTURE AND STRATEGIC DIRECTIONS**

We will make SNSL conducive to research and educational use during the second year and add more sensor nodes to the lab for education/research purposes.

- (a) Install a sensor networks with environmental sensors in the field at a weather station
- (b) Install network switches and sensor nodes.
- (c) Install additional sensor nodes in the field for large-scale experiments

# **RI: High-End Computing and Networking Research Testbed for Next Generation Data Driven, Interactive Applications**

**Proposal Number: RI#0403342, Project Year (Year 2)**

**The Ohio State University  
(Dept. of Computer Science and Engineering, Dept. of Biomedical Informatics and Ohio Supercomputer Center)**

**PIs: D. K. Panda, G. Agrawal, P. Sadayappan, J. Saltz and H.-W. Shen**

**Other Investigators: S. Ahalt, U. Catalyurek, H. Ferhatosmanoglu, H.-W. Jin, T. Kurc, M. Lauria, D. Lee, R. Machiraju, S. Parthasarathy, P. Sinha, D. Stredney, A. E. Stutz and P. Wyckoff**

**PIs E-mails: {panda,agrawal,saday,hwshen}@cse.ohio-state.edu,saltz-1@medctr.osu.edu**

## **1. Abstract**

This report describes the progress at the end of the second year of RI project at The Ohio State University. The overall goal of this project is to research the design and development of next generation data-intensive and interactive applications over a networked collection of storage, compute, memory, visualization, and wireless resources. During the second year of this project, major equipment of the infrastructure have been acquired, installed, and made operational. Research along the proposed directions has already produced a large number of publications.

## **2. Introduction/Background**

Over recent years, the field of high-end computing has been going through two important paradigm shifts. The first is an increasing focus on interactive data driven applications, away from the traditional batch-mode compute-intensive tasks. This is driven by a number of factors. First, scientific simulations and increasing numbers of high precision data collection instruments (e.g. sensors attached to satellites, medical imaging modalities) are creating very large datasets. With technology trends, our ability to store and share these datasets is also increasing. As a result, scientists and institutions can create repositories of large scientific datasets (of the order of 10-100 Terabytes) and make them available for use by other researchers. The ability to interact, synthesize, and visualize large datasets, potentially generated by researchers throughout the world, significantly increases the insights that can be achieved by a scientist. However, this imposes a need to co-schedule computing, networking, and visualization resources to achieve interactivity and meet real-time or near real-time constraints on the turn-around time.

The second is a clear shift from tightly coupled high-end computing systems to networked clusters of machines within an organization and across geographically dispersed organizations. This is driven by several developments: 1) high-speed Wide Area Networking (WAN) such as the National Light Rail (NLR) system providing 10Gbps connectivity across major cities/labs/research centers in the Nation, 2) the economics of creating cost-effective high-end computing environments with clusters (consisting of commodity PCs with multi GHz speed and cluster networks of less than 5.0 microsec latency and tens of Gbps bandwidth), and 3) the needs and characteristics of today's collaborative applications. A cluster environment can serve as a cost-effective and flexible mechanism for both hosting large data repositories and executing compute-intensive tasks.

In addition to the advances in computing and networking technologies, memory technology is also advancing. With the availability of dense memory modules, memory costs are declining rapidly. This is providing new opportunities in design and deploying large-scale memory-intensive clusters with commodity processors, memory, and interconnects. Moreover, on current generation systems, it is much faster (5.0-10.0 microseconds) for a node to access data from the memory of a neighboring node than to access (a few milliseconds) data from its own disk.

### **3. Goals, Objectives, Targeted Activities**

With the above application and technology trends, our overall goal in this project is to address a new set of research problems. These problems are:

- Can a large memory cluster be integrated into the organizational-level computing environment (compute cluster and clients) and a remote data repository in a seamless manner ?
- How will the various subsystems and components of the organization-level computing environment (such as communication, I/O, networking, QoS, programming models, compilers, scheduling, middleware, indexing, caching policies, etc.) need to be re-designed to take advantage of the capabilities provided by the memory cluster and address the challenges posed by the heterogeneous and distributed nature of the environment?
- How will the visualization and data mining algorithms need to be re-designed to harness the benefits of the new environment?
- How will the next generation data-intensive applications benefit from the new organizational-level computing environment and the above enhancements and by how much?

### **4. Infrastructure and Facilities**

We have procured and installed the following equipment: At the CSE department, a 64-node cluster (32 dual Intel EM64T nodes and 32 dual Opteron nodes) with InfiniBand DDR (20 Gbps) interconnect has been deployed. Eight of the Opteron nodes also have storage disks with a total storage capacity of 8 TBytes. All Opteron nodes have advanced Graphics adapters and they are connected to the existing video wall. At the BMI department, a 64-node dual Opteron cluster with InfiniBand SDR (10 Gbps) interconnect and 512 GBytes of physical memory has been installed. This cluster also has 6 storage nodes with a total storage capacity of 24 TBytes. Both these clusters are connected with dual 10Gbps (20 Gbps) WAN links and they are also connected to the Mass Storage system at Ohio Supercomputer Center. The CSE cluster is also connected with a wireless testbed. A set of laptops and PDAs for this wireless testbed has already been acquired and deployed. As a part of this facility, a full-time System administrator also has been hired. He is currently taking care of all hardware/software acquisitions, installation, maintenance, upgrades, etc. Our plans for the near future include the following: The purchase for wireless gears is through the purchasing process and they are planned to arrive by August 1<sup>st</sup> week. The identification for the second interconnect for the CSE cluster is currently being discussed. Since 10GigE is emerging, it is planned to have this as the second network instead of Quadrics (which is declining in the HPC cluster market). Several 10GigE adapters (such as Chelsio, NetEffect and Myrinet) are currently emerging in the market. We plan to carry out some preliminary experiments with some of these adapters about their cost-performance and then make a final decision on this.

### **5. Project Description**

To address the issues we listed in Section 3, and using the infrastructure described in the previous section, we are conducting research in four areas:

*Component I: Networking, Communication, QoS, and I/O:* Panda, Jin, Lee, Lauria, Sinha, and Wyckoff are conducting research on improving inter- and intra-cluster networking, communication, QoS and I/O. New designs using the emerging InfiniBand interconnect and its associated communication protocols for upper-layers (MPI, Datacenters, Virtual Machines, high performance sockets, I/O and File systems) and

efficient communication in wide area networks and protocols have been proposed and studied. Publications along these directions include: [16,17,18,19,21,22,25,27,29,30,31,32,34,35,36,37,39,20,23, 24,28,33,26,38].

*Component II: Programming Systems and Scheduling:* Saltz, Agrawal, Sadayappan, Kurc, Catalyurek, Ahalt, and Ferhatosmanoglu are addressing a number of challenges in improving application development on networked, heterogeneous clusters (a filter-stream programming based middleware for developing data intensive applications, tools for scheduling tasks within a single application, high-level language and compiler design for distributed environments, techniques for efficient execution of multiple applications in the same system, middleware for data mining tasks, and methods for indexing and declustering of data in distributed storage clusters). Publications along these directions include: [1,2,3,4,40,42].

*Component III: Data Intensive Algorithms:* To effectively use a networked collection of storage, compute, memory and visualization resources for an application, novel algorithms are required. Shen, Agrawal, Parthasarathy, and Machiraju have been exploring issues related to visualization, data mining, and scientific data analysis. Algorithms through stand-alone implementations, as well as by implementing them on top of the middleware systems DataCutter and FREERIDE are being studied. Publications along these directions include: [5,6,7,8,9,12,13,14,15].

*Component IV: Data Intensive Applications:* The computer science researchers at the Ohio State University are involved in a number of collaborative projects that focus on challenging end applications. These applications are data intensive, data driven, interactive, and collaborative in nature. Publications along these directions include: [10,11,41,43].

## **6. Indicators of Success and Major Accomplishments**

During the second year of the project, we have carried out research along several directions using the proposed infrastructure. The publications are listed in the following section. We plan to continue on these initial results and design the complete solutions for next-generation interactive and data-driven applications.

## **7. Publications**

- [1]. T. Apaydin, G. Canahuate, H. Ferhatosmanoglu, A. Tosun., Approximate Encoding for Direct Access and Query Processing over Compressed Bitmaps. 32nd International Conference on Very Large Data Bases (VLDB '06), Seoul, Korea, September, 2006, to appear.
- [2]. H. Ferhatosmanoglu, A. Tosun, G. Canahuate, A. Ramachandran, Efficient Parallel Processing of Range Queries through Replicated Declustering. Distributed and Parallel Databases Journal, to appear.
- [3]. G. Canahuate, M. Gibas, H. Ferhatosmanoglu, Indexing Incomplete Databases. 10th International Conference on Extending Database Technology (EDBT '06). Munich, Germany, March, 2006, pp. 884-901.
- [4]. T. Apaydin and H. Ferhatosmanoglu, Access Structures for Angular Similarity Queries. IEEE Transactions on Knowledge and Data Engineering (TKDE), to appear.
- [5]. G. Buehrer, S. Parthasarathy, and A. Ghoting, "Efficient Out-of-core Frequent Itemset Mining on a Commodity PC", In Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), 2006.
- [6]. A. Ghoting, G. Buehrer, S. Parthasarathy, D. Kim, A. Nguyen, Y.Chen, and P. Dubey, "Cache-Conscious Frequent Pattern Mining on Modern and Emerging Processors", In International Journal on Very Large Data Bases (VLDBJ), 2006.
- [7]. G. Buehrer, A. Ghoting, X. Zhang, S. Tatikonda, S. Parthasarathy, T. Kurc, and J. Saltz, "I/O Conscious Algorithm Design and Systems Support for Data Analysis on Emerging Architectures", In Proceedings of IPDPS Workshop on Next Generation Software, 2006.

- [8]. A. Ghoting, G. Buehrer, S. Parthasarathy, D. Kim, A. Nguyen, Y. Chen, and P. Dubey, "Cache-Conscious Frequent Pattern Mining on a Modern Processor", In Proceedings of the International Conference on Very Large Data Bases (VLDB), 2005.
- [9]. A. Ghoting, G. Buehrer, S. Parthasarathy, D. Kim, A. Nguyen, Y. Chen, P. Dubey, "A Characterization of Data Mining Algorithms on a Modern Processor ", In Proceedings of the SIGMOD Workshop on Data Management on New Hardware (DaMoN), 2005.
- [10]. H. Wang, S. Parthasarathy, A. Ghoting, S. Tatikonda, G. Buehrer, T. Kurc, and J. Saltz, "Design of a Next Generation Sampling Service for Large Scale Data Analysis Applications", In Proceedings of the ACM International Conference on Supercomputing (ICS), 2005.
- [11]. H. Wang, A. Ghoting, G. Buehrer, S. Tatikonda, S. Parthasarathy, T. Kurc, and J. Saltz, "A Services Oriented Framework for Next Generation Data Analysis Centers ", In Proceedings of the IPDPS Workshop on Next Generation Software, 2005.
- [12]. T. Kerwin, H.-W. Shen, and D. Stredney, Capture and review of interactive volumetric manipulations for surgical training, In Proceedings of International Workshop on Volume Graphics 2006 (to appear)
- [13]. C. Wang and H.-W. Shen, LOD Map - A Visual Interface for Navigating Multiresolution Volume, Visualization, In Proceedings of IEEE Visualization 2006 (to appear)
- [14]. G. Ji and H.-W. Shen, Dynamic View Selection for Time-Varying Volumes, In Proceedings of IEEE Visualization 2006 (to appear)
- [15]. J. Woodring and H.-W. Shen, Multi-variate, Time Varying, and Comparative Visualization with Contextual Cues, In Proceedings of IEEE Visualization 2006 (to appear)
- [16]. G.-Q. Shu and D. Lee, "Message Confidentiality Testing of Security Protocols - Passive Monitoring and Active Checking", IFIP TestCom May 2006
- [17]. G.-Q. Shu, D. Lee and M. Yannakakis, "Key Management for Large Scale Emergency Alert Systems", SSN May 2006.
- [18]. G.-Q. Shu and D. Lee, "Network Protocol System Fingerprinting – A Formal Approach", IEEE Infocom, May 2006
- [19]. Z. Liu, N. Li, G. Shu, W. Lin and D. Lee, "Instant Messaging Spams and Worms - their Detection and Filtering", Proc. ICNP NSEC, Nov. 2005
- [20]. P. Balaji, K. Vaidyanathan, S. Narravula, H.-W. Jin, and D. K. Panda, Designing Next-Generation Data-Centers with Advanced Communication Protocols and Systems Services, Workshop on NSF Next Generation Software(NGS) Program; held in conjunction with IPDPS, Greece, 2006.
- [21]. W. Huang, G. Santhanaraman, H.-W. Jin, Q. Gao and D. K. Panda, Design and Implementation of High Performance MVAPICH2: MPI2 over InfiniBand, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Singapore, May 2006
- [22]. L. Chai, R. Noronha and D. K. Panda, MPI over uDAPL: Can High Performance and Portability Exist Across Architectures? Int'l Symposium on Cluster Computing and the Grid (CCGrid), Singapore, May 2006
- [23]. S. Narravul, H.-W. Jin, K. Vaidyanathan and D. K. Panda, Designing Efficient Cooperative Caching Schemes for Multi-Tier Data-Centers over RDMA-enabled Networks, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Singapore, May 2006
- [24]. J. Liu, W. Huang, B. Abali and D. K. Panda, High Performance VMM-Bypass I/O in Virtual Machines, USENIX Annual Technical Conference, June 2006.
- [25]. Q. Gao, W. Yu, W. Huang and D. K. Panda, Application-Transparent Checkpoint/Restart for MPI Programs over InfiniBand, Int'l Conference on Parallel Processing (ICPP), August 2006.
- [26]. S. Liang, W. Yu and D. K. Panda, High Performance Block I/O for Global File System (GFS) with InfiniBand RDMA, Int'l Conference on Parallel Processing (ICPP), August 2006.
- [27]. M. Koop, W. Huang, A. Vishnu and D. K. Panda, Memory Scalability Evaluation of the Next-Generation Intel Bensley Platform with InfiniBand, Int'l Symposium on Hot Interconnect (HotI), August 2006.
- [28]. W. Huang, J. Liu, B. Abali, and D. K. Panda, A Case for High Performance Computing with Virtual Machines, The 20th ACM International Conference on Supercomputing (ICS'06).
- [29]. S. Sur, L. Chai, H.-W. Jin and D. K. Panda, Shared Receive Queue based Scalable MPI Design for InfiniBand Clusters, Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006, Rhode Island, Greece.
- [30]. W. Yu, Q. Gao and D. K. Panda, Adaptive Connection Management for Scalable MPI over InfiniBand, International Parallel and Distributed Processing Symposium (IPDPS 2006).

- [31]. A. Mamidala, L. Chai, H.-W. Jin and D. K. Panda, Efficient SMP-Aware MPI-Level Broadcast over InfiniBand's Hardware Multicast, Communication Architecture for Clusters (CAC) Workshop, to be held in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006,
- [32]. S. Sur, H.-W. Jin, L. Chai and D. K. Panda, RDMA Read Based Rendezvous Protocol for MPI over InfiniBand: Design Alternatives and Benefits, Symposium on Principles and Practice of Parallel Programming (PPoPP), March 2006.
- [33]. P. Balaji, S. Bhagvat, H. -W. Jin and D. K. Panda Asynchronous Zero-copy Communication for Synchronous Sockets in the Sockets Direct Protocol (SDP) over InfiniBand, Communication Architecture for Clusters (CAC'06).
- [34]. S. Sur, U. Bondhugula, A. Mamidala, H.-W. Jin, and D. K. Panda, High Performance RDMA Based All-to-all Broadcast for InfiniBand Clusters, International Conference on High Performance Computing (HiPC 2005), December 18-21, 2005, Goa, India.
- [35]. A. Vishnu, G. Santhanaraman, W. Huang, H. -W. Jin and D. K. Panda, Supporting MPI-2 One Sided Communication on Multi-Rail InfiniBand Clusters: Design Challenges and Performance Benefits , International Conference on High Performance Computing (HiPC 2005), December 18-21, 2005, Goa, India.
- [36]. L. Chai, R. Noronha, P. Gupta, G. Brown, and D. K. Panda, Designing a Portable MPI-2 over Modern Interconnects Using uDAPL Interface, EuroPVM/MPI 2005, Sept. 2005.
- [37]. A. Mamidala, H. -W. Jin, and D. K. Panda, Efficient Hardware Multicast Group Management for Multiple MPI Communicators over InfiniBand, EuroPVM/MPI 2005, Sept. 2005.
- [38]. W. Yu and D. K. Panda, Benefits of Quadrics Scatter/Gather to PVFS2 Noncontiguous I/O, International Workshop on Storage Network Architecture and Parallel I/Os (SNAPI) 2005.
- [39]. S. Sur, A. Vishnu, H. -W. Jin, W. Huang, and D. K. Panda, Can Memory-Less Network Adapters Benefit Next-Generation InfiniBand Systems?, Hot Interconnect (HOTI 05), August, 2005.
- [40]. G. Khanna, U. Catalyurek, T. Kurc, P. Sadayappan, and J. Saltz. "A Data Locality Aware Online Scheduling Approach for I/O-Intensive Jobs with File Sharing.", 12th Workshop on Job Scheduling Strategies for Parallel Processing, In Conjunction with SIGMETRICS 2006, Saint-Malo, France, June 2006, to appear.
- [41]. N. Vydyanathan, S. Krishnamoorthy, G. Sabin, U. Catalyurek, T. Kurc, P. Sadayappan, and J. Saltz. "An Integrated Approach for Processor Allocation and Scheduling of Mixed-Parallel Applications.", International Conference on Parallel Processing (ICPP-06), Columbus, OH, August 2006, to be presented.
- [42]. L. Weng, U. Catalyurek, T. Kurc, G. Agrawal, J. Saltz. "Using Space and Attribute Partitioned Partial Replicas for Data Subsetting and Aggregation Queries.", International Conference on Parallel Processing (ICPP-06), Columbus, OH, August 2006, to be presented.
- [43]. S. Narayanan, T. Kurc, U. Catalyurek, J. Saltz. "On Creating Efficient Object-relational Views of Scientific Datasets.", Accepted for publication in the The 2006 International Conference on Parallel Processing (ICPP-06), Columbus, OH, August 2006.

## 8. Outreach Programs and Supported Students

As a part of this RI program, several investigators (Panda, Sadayappan, Agrawal, Lauria, Parthasarathy and Stredney) have participated in the summer parallel computing institutes, held annually by Central State University (CSU), one of the HBCU's in the State of Ohio. During Summer '05, one of the undergraduate students of CSU also visited Panda's lab and spent the summer in having hands-on experience with modern networks for clusters. A large number of graduate students (being supported by other grants from NSF, DOE, NIH and industry) are currently using the RI infrastructure for their research and dissertations. These students include: T. Apaydin, G. Canahuate, A. Tosun, M. Gibas, G. Buehrer, A. Ghoting, X. Zhang, S. Tatikonda, T. Kerwin, C. Wang, G. Ji, J. Woodring, G.-Q. Shu, Z. Liu, N. Li, P. Balaji, K. Vaidyanathan, S. Narravul, W. Huang, G. Santhanaraman, Q. Gao, L. Chai, R. Noronha, W. Huang, S. Liang, M. Koop, S. Sur, W. Yu, A. Mamidala, S. Bhagvat, A. Vishnu, P. Gupta, G. Brown, G. Khanna, N. Vydyanathan, S. Krishnamoorthy, G. Sabin, L. Weng and S. Narayanan

## 9. Future and Strategic Directions

With our infrastructure being completely operational in Year2, we have carried out initial studies along proposed research directions. Much of our proposed work still needs to be carried out. Our plans in the next 1 year (third year of the grant) will focus on larger-scale evaluation studies of the various new tools and techniques we have been developing. Starting from this year, we will also be emphasizing integration of various components and testing with end-applications.

# RI: I<sup>3</sup>C: An Infrastructure for Innovation in Information Computing

**Award Number: CISE/EIA-0202007**

**Institution: Pennsylvania State University**

**Principal Investigators: Chita R. Das, Raj Acharya, C. Lee Giles,  
Mary J. Irwin, and Padma Raghavan**

**Email Address: {das, acharya, giles, irwin, raghavan}@cse.psu.edu**

**Project Website: <http://www.cse.psu.edu/i3c>**

## Abstract

The objective of this research infrastructure, called I<sup>3</sup>C, is to support multidisciplinary research spanning across several engineering and science disciplines, while focusing on three core areas of computer science that are essential to advance the state-of-the-art in cluster computing. The infrastructure primarily consists of three cluster platforms, a storage area network, and a multimedia laboratory for conducting research in *applications*, *system software*, and *architecture*. The application-based research in turn covers three specific domains: *Computational Science*, *Digital Immortality and Bioinformatics*.

## 1. Introduction

The I<sup>3</sup>C research infrastructure is aimed at providing a unique experimental facility for collaborative research involving more than twenty researchers from three colleges: College of Engineering, School of Information Sciences and Technology (IST), and College of Science. The infrastructure is a cluster computing platform with large storage for supporting data intensive applications. It consists of two Linux clusters (96-node and 64-node) and a 28-node Solaris cluster that will be used for satisfying the computational needs of a broad spectrum of applications, a 15-TB network attached storage (NAS), a multimedia laboratory for visualization and HCI related projects, and required networking equipment for integrating all the components as a shared resource.

The infrastructure supports interdisciplinary research primarily in three areas. The first area is *applications*, which is the driving force for the infrastructure. The application-based research in turn covers three specific domains: *Computational Science*, *Digital Immortality and Bioinformatics*. All these applications are not only computationally intensive, but also need substantial storage, where terabytes of data will be generated and archived. The second area of research, *System Software*, investigates how various resources can be effectively used for improving the delivered performance. The third area of research, *Architecture*, focuses on low-level design details to provide high and assured performance. In summary, the infrastructure will facilitate research activities, which are currently supported by many NSF and other federal and industry awards. This report summarizes the status of the project to date.

## 2. Goals, Objectives and Targeted Activities

The I<sup>3</sup>C project is likely to have a significant positive impact on the overall high performance computing research. It will be extensively used in conducting fundamental and applied research involving researchers from three colleges at Penn State. The outcome of the research will advance the state-of-the-art in cluster computing by demonstrating how both commercial and scientific applications with varying performance and QoS requirements can be executed effectively on large platforms. Locally, the infrastructure will be an invaluable asset for both undergraduate and graduate teaching and research. We plan to involve both undergraduate and graduate students in this project making every effort to make sure that students from underrepresented groups are included. The experimental platform will be a catalyst to engage in many other collaborative research activities at Penn State, and help us in obtaining additional research funding from other sources. The successful deployment of various projects and the tools and techniques developed through this cohesive research is likely to have a much broader impact in many areas of information processing.

### 3. Infrastructure and Facilities

The major equipments, excluding the workstations and laptops, purchased to date through this grant are summarized below.

- We have procured three clusters for meeting our computational demand. In the first year, we have procured a 16-node Myrinet-connected Linux cluster to initiate research in all the three main areas. In the second year, our compute engine has been expanded to include a 28-node Solaris cluster. Each compute node has dual UltraSPARC IIIi processors clocked at 1 GHz and 4 GB of RAM. In the following year, we installed a 96-node IBA-connected Linux cluster. Each node in the cluster has two 64-bit processors and 4GB of RAM. We are in the process of buying another 64-node IBA-connected cluster. This cluster will include voltage scaling facility to conduct research on power saving techniques.
- For the storage system, we have purchased 4TB of FibreChannel storage and 11.5TB of SATA storage on the Bluearc array for a total of 15.5TB of disk space. We have also purchased two additional arrays from Maxtronic containing a total of 4TB of disk space. Last year, we procured a 50TB tape robot from IBM to backup the storage network.
- For the multimedia laboratory, we have purchased a 50-inch plasma display, cameras, and other photographic equipments.

### 4. Project Description

The research activities that have extensively utilized the I<sup>3</sup>C infrastructure are summarized in three areas: applications, system software and architecture.

#### 4.1 Application Research

The three application areas are characterized by high computational and storage requirements. They have utilized the Linux and Solaris cluster and the NAS extensively. Highlights from the applications research from the various areas are given below.

**4.1.1 Computational Science:** The computational science research focuses on developing algorithm framework, designing scalable parallel algorithms, and implementing these algorithms for a number of applications such as many-body problems and sparse matrix solvers. For example, Paul Plassmann, Daniel Haworth and their graduate students have developed new algorithms and software for a scientific, on-line database for function approximation. The software implementation of these algorithms is called DOLFA and is freely available for use by research scientists. This database software can significantly reduce the computational demands for problems requiring the frequent evaluation of computationally expensive functions and has been extensively tested and used in complex computational chemistry calculations.

DOLFA can be of great importance as scientists incorporate more sophisticated models into simulations because the underlying computational cost of these models can become substantial. Furthermore, calculations that involve detailed chemistry are typically so complex and demanding that they often exceed computational resources of today's computers.

The second project, called Grid-Enabled On-Line Laboratory for Multiscale Multicomponent Materials Design, focuses on the development of an extensible software system for enabling domain-specific, automated design space exploration. This project focuses specifically on the design of the technologically important alloys and the enabling software that can be utilized in other applications, such as for the design of bio-nano materials. This project concerns the development of a system to automate investigations that predict macroscopic properties (such as the mechanical response) by combining a four stage multi-scale, multi-physics computational process with empirically obtained material properties. Dr. Padma Raghavan leads the computational aspect of this project, which includes generation of microstructures using phase-field simulations, and performing finite-element analysis on the simulated microstructures to determine properties such as the mechanical response.

**4.1.2 Digital Immortality:** The digital immortality research focuses on a broad spectrum of interactive and archival information processing issues such as design of scalable and efficient niche search engines, multiagent architectures, and multimedia servers. Here we summarize two projects CiteSeer and ALIP.

For computer scientists and related disciplines CiteSeer, a computer science document search engine and digital library, has become over the last three years the document search engine of choice. The goal of this research is to explore methods for making CiteSeer a permanent research fixture for the computer science community. Methods of archive and URL permanence plus methods for enhanced archive mirroring and archive scaling are currently investigated. Availability of CiteSeer resources will be promoted through exploratory APIs (application programming interfaces) for researcher data access. In addition new linked information such as organizations and individuals in acknowledgments will be correlated with citation rankings. This research will give computer scientists enhanced access to the computer science literature and promote the addition of new additions to the document database by investigating new document crawling algorithms and procedures.

James Wang and his research group in the Intelligent Information Systems Research Laboratory have been developing novel theories and techniques for interpreting the semantic content of images. Since 2003, the team have benefited from the computing infrastructure made available by the I<sup>3</sup>C grant. Unlike conventional approaches to image understanding, the team attempt to train computers with hundreds of concepts. With advanced cluster computers, the research team can now conduct a full-scale training experiment of 600 different concepts within a day or two. Such experiments could take months of computer time on a single-chip computer. Their ALIP (Automatic Linguistic Indexing of Pictures) system has become the first computer system that can annotate photographic images with as many as 600 concepts.

The cluster computers, massive storage, and the imaging equipment obtained through this grant have also enabled the team to work on several related research problems, including computational aesthetics, image-based computer security, story picturing engine, biomedical informatics, and machine-assisted analysis of ancient paintings. They treat the challenge of automatically inferring aesthetic quality of pictures using their visual content as a machine learning problem, with a peer-rated online photo sharing Website as data source. Certain visual features are extracted based on the intuition that they can discriminate between aesthetically pleasing and displeasing images. Automated classifiers are built using support vector machines and classification trees. Linear regression on polynomial terms of the features is also applied to infer numerical aesthetics ratings. The work attempts to explore the relationship between emotions which pictures arouse in people, and their low-level content. Potlucyential applications include content-based image retrieval and digital photography.

**4.1.3 Bioinformatics:** The bioinformatics research, another interdisciplinary collaboration, focuses on development of information theoretic methods for temporal analysis of gene expressions, analysis of genes and gene networks for protein sequence, biodiversity and study of genomic function by the analysis of DNA replication and transcription sites.

As an example, Raj Acharya's group is working on identification and evaluation of Discriminative Motifs for Gene Clusters. With the availability of large amounts of high-throughput experimental and genomic sequence data, several approaches are being taken to tackle the challenging task of accurately predicting cis-regulatory elements through computational methods. Only a few of the cis-acting elements that control gene expression levels are well known, but their identification genome-wide will help in understanding the mechanisms that underlie coordinated gene regulation. Putative regulatory elements can be identified by computational methods by searching for shared motifs within non-coding sequences of genes believed to be co-regulated, but current success rates are modest. An enumeration-based algorithm to identify statistically significant motifs has been developed. The motifs discriminate between multiple gene clusters, for example, obtained by expression clustering. The algorithm incorporates use of multi-species sequence conservation information. Preliminary results indicate significant improvement in classification performance by the addition of conservation. The algorithm is designed to perform exploratory analysis and unlike other motif finding methods, handles very large amounts of data required for discovering motifs in higher eukaryotes. The putative motifs are assessed using benchmark data with known binding sites and on real mouse genes that are up- or downregulated in Gata1-null proerythroblasts (G1E cells) after restoration of Gata1 activity.

In another project, called GALA, we have developed a relational database to contain whole genome sequence alignments between human and mouse with extensive annotations of the human sequence. Complex queries are supported on recorded features, both directly and on proximity among them. Searches can reveal a wide variety of relationships such as finding all genes expressed in a designated tissue that have a highly conserved non-coding sequence 5' to the start site. Other examples are finding single nucleotide polymorphisms that occur in conserved non-coding regions of upstream of genes and identifying CpG islands that overlap the 5' ends of divergently transcribed genes.

## 4.2 System Software Research

The software research primarily focuses on developing efficient scheduling techniques for clusters to improve application performance. Although several coscheduling mechanisms have been proposed for clusters, so far no work has matured into a practically deployable implementation. Keeping this in view, we have designed a generic framework that facilitates fast development of customized coscheduling heuristics. The entire framework has been implemented on the 16-node Linux cluster that uses the industry standard VIA and GM as the communication paradigm. This is the first implementation that provides opportunities to bridge the gap between a network and the OS, with respect to end-to-end QoS, and also has the flexibility to implement adaptive co-scheduling schemes depending on the workload characteristics.

Then, we have proposed two new co-scheduling algorithms, Coordinated Coscheduling (CC) and HYBRID, which make use of the proposed framework. Performance evaluation using NAS Benchmarks indicates that our proposed CC scheme can outperform the previously proposed coscheduling algorithms (DCS, SB and PB) and HYBRID technique outperforms CC, batch scheduling (PBS) and gang scheduling (SCORE) techniques in terms of average response time. We plan to use the workload characteristics of the above applications to analyze the effectiveness of the proposed algorithms on the 96-node cluster.

Next, we conduct research designing efficient and dependable data centers in a cluster environment. We mainly focus on a three-tier data center architecture, which includes a cluster of Web servers at the front-end tier, a cluster of Java-based application servers at the mid-tier and a cluster of DB servers at the backend-tier. Initially, we have implemented a prototype data center on a 32 node IBA-connected Linux cluster for performance analysis. Our comparative analysis has shown that using IBA instead of the traditional Ethernet can improve the throughput of a data center. In addition, we have proposed and implemented a load-balancing mechanism between the application server nodes in the mid-tier, which improved the throughput further in our experimental analysis.

We are currently working on a workload characterization study in order to better understand the effects of different parameters in a data center design. In the future, we plan to extend our study to increase the dependability of data centers. This will be done by injecting several types of faults and developing a comprehensive fault-tolerant technique. Furthermore, we will use the cluster platform to investigate effective virtualization techniques.

## 4.3 Architecture Research

The architecture research focuses on three main areas: Design and analysis of high performance and QoS Capable Clusters, Network-on-Chip (NoC) architectures, and 3D cache design. Most of these studies have used the Linux and Solaris clusters for extensive simulation.

**Design and Analysis of High Performance and QoS capable Clusters:** In this research, we explore various design alternatives to provide high and predictable performance in cluster interconnects. Based on the recently released InfiniBand Architecture (IBA) specification, we have designed a router and a network interface card (NIC) to design any IBA-style cluster interconnect for providing high and predictable performance. The IBA simulator will be made publicly available through our website.

Since minimizing power consumption in clusters has become a critical issue, we have developed an energy model for estimating the power consumption in a cluster interconnect. Currently, we are developing various techniques to minimize power consumption in different components of an interconnect.

**Network-on-Chip (NoC) Architectures:** Design of high performance, energy-efficient and fault-tolerant NoCs has become a recent research thrust for supporting the System-on-Chip (SoC) design paradigm. We are currently developing a comprehensive NoC simulator to estimate the performance, energy consumption and fault-tolerance of various on-chip interconnects. We have proposed two router architectures that are quite effective in terms of performance and energy consumption. We will investigate the impact of soft errors on these architectures.

As technology scales, power consumption and thermal effects have become challenges for system-on-chip designers. The rising on-chip temperatures can have negative impacts on SoC performance, power, and reliability. In view of this, we present a hybrid optimization approach which aims at temperature reduction and hot spot elimination. We demonstrate that considerable improvement in the thermal distribution of a design can be achieved through careful voltage island partitioning, voltage level assignment, and voltage island floor planning. The experimental results on MCNC benchmarks show significant improvement on the thermal profiles. To the best of our knowledge, this is the first work to explore the thermal impacts of voltage islands.

**3D Cache Design:** As technology scales, interconnect dominate the performance and power behavior of deep submicron designs. Three-dimensional integrated circuits (3D ICs) have been proposed as a way to mitigate the interconnect challenges. We are exploring the architectural design of cache memories using 3D circuits. We have developed a delay and energy model, 3DCacti, to explore different 3D design options of partitioning a cache. The tool allows partitioning of the cache across different device layers at various levels of granularity. The tool has been validated by comparing its results with those obtained from circuit simulation of custom 3D layouts. We also explore the effects of various cache partitioning parameters and 3D technology parameters on delay and energy to demonstrate the utility of the tool.

## 5. Indicators of Success and Major Accomplishments

The I<sup>3</sup>C infrastructure has made a dramatic impact in enabling areas of multidisciplinary research that required integrated computational resources. Major indicators of success include development of a shared resource infrastructure for advancing the state-of-the art in high performance computing, significant increase in collaborative research in established and new areas, increased level of joint research funding from NSF as well as other sources, and development of cross-disciplinary graduate courses.

## 6. Publications

Due to space limitations, we do not include the publications here. Key publications related to this projects can be found at our project website (<http://www.cse.psu.edu/i3c>) and in the NSF annual report.

## 7. Outreach Programs and Supported Students

The clusters are extensively used by a number of graduate and undergraduate students for their research in the three core areas. In addition to student training, the major outreach activities are the following:

- The I<sup>3</sup>C infrastructure currently supports the CiteSeer website. This website is the largest full-text index of scientific literature covering more than 500,000 documents. It needs at least 2 TB of storage and high computing power, which can be provided by our infrastructure. When fully operational, it is expected to support thousands of accesses daily form all over the world.
- The Genome Alignment and Annotation Database (GALA) is a comprehensive database that incorporates information about genes, SNPs, alignments, and disease association. It is also supported by our infrastructure and it is expected to be one of the most accessed website for Bioinformatics research.
- Our current partners in this project are SUN Microsystems, Lawrence Livermore National Lab (LLNL), and Intel. We are actively seeking outside collaboration for various research activities.
- We have started a graduate seminar course to discuss various multidisciplinary research issues. This seminar should open new research directions and should attract new students to work on research enabled through the I<sup>3</sup>C project.
- PSU CSE has joined forces with six cancer centers of Pennsylvania in order to form the Pennsylvania Cancer Alliance Bioinformatics Consortium (PCABC). The research infrastructure hosts a public domain data warehouse containing cancer biomarkers from these cancer centers. The data warehouse can also aid researchers to perform complex analysis on these biomarkers using the tools of information fusion.

## 8. Future and Strategic Directions

We have high expectations that this infrastructure will continue to enable new avenues of research well beyond the project duration. We plan to continue our interdisciplinary research activities in many of the key/upcoming areas such as bioinformatics, biologically inspired computing, computational science, material science, and nanotechnology. Moreover, we plan to apply for an NSF IGERT grant that would utilize the I<sup>3</sup>C infrastructure.

# RI: Pervasive Computing: Applications and Systems

(0101247)

Perry Cook, Kai Li, Margaret Martonosi, Vivek Pai, and Larry Peterson  
Princeton University

## Abstract

This project is broadly focused on networked systems that pervade (and are embedded in) the physical world, and the applications that run on top of them. Much of our work this past year has involved the development of novel networking capabilities that support pervasive applications, ranging from sensor networks used to track animal migrations, to postal-based network for weakly-connected parts of the world, to global-scale overlay networks used to deploy new wide-area services.

## Project Description

### ZebraNet

This part of the project focused on mobile, pervasive, and energy-aware computing. In particular, one aspect of our work has been on experimental and analytic techniques for dynamically estimating and controlling the energy consumption of computing devices, particularly mobile ones. Our work recognizes that in many domains, the central problem is not to minimize power consumption at all costs, but rather to create systems that responsively trade-off performance and power in ways that improve performance while keeping to the system's power budget. Our work has spanned the microarchitecture, architecture, compiler and OS levels of computer systems. We have developed on-the-fly techniques for power phase analysis, and applied these techniques to dynamic adaptation mechanisms such as voltage/frequency scaling and thread migration.

A second key use of the infrastructure equipment has been to support deployments of pervasive technology that seek to demonstrate the real-world viability of energy-efficient hardware and software techniques we have proposed in research. In particular, we have focused much of our efforts on the ZebraNet project (supported by the ITR program via NSF CCR-0205214 "ZebraNet: Position-aware Power-aware Wireless Computing for Wildlife Tracking"). The ZebraNet project mixes computer systems and biology research thrusts in a fully interdisciplinary project aimed at using mobile ad hoc wireless networks for detailed wildlife tracking. ZebraNet allows wildlife researchers to pose and to answer important long-standing questions about long-range migration, inter-species interactions, and nocturnal behavior. As a computer systems research problem, ZebraNet is compelling because the needs of the biological researchers are stringent enough to require real breakthroughs in wireless protocols and in low-power computer systems design and computer systems power management. These breakthroughs can be leveraged into other (non-wildlife-oriented) fields of research. ZebraNet is a power-aware wireless ad hoc sensor network, but with more serious bandwidth and compute needs than most prior sensor networks research problems.

### Scalable Displays: Visualizing Biological Data

Science disciplines are experiencing a data avalanche. As a result, scientific research is limited by data analysis and visualization capabilities. This project is to explore how to build and use scalable display systems to substantially improve scientists' ability to visualize data and to collaborate through space and time. During the past year, we have been working closely with Genomic and Plasma Physics researchers on effective data visualization software tools.

During the past year, we have focused on research to develop software tools for high-resolution display walls to alleviate the current limitation on visualization resolution and single-user window system.

### PlanetLab

The key contribution of PlanetLab is its support of distributed virtualization—allocating a widely dispersed set of virtual machines to a user or application. Distributed virtualization enables broad-coverage services that benefit from having multiple points-of-presence on the network. The main technical challenge of PlanetLab is managing distributed virtualization in the wide-area.

When a single organization owns and manages a wide-area platform that supports distributed virtualization, control can be logically centralized, even if the underlying mechanisms are distributed. This is the model used by companies that offer customers the ability to run "business logic" on their distributed nodes (e.g., Akamai). However,

distributing a network application across a federation of such platforms—perhaps acquiring a point-of-presence at hundreds or thousands of autonomous Internet sites—requires that distributed virtualization be managed in a decentralized fashion. A key challenge is to decentralize control while preserving the necessary trust relationships.

We characterize the problem as follows. First, there is a set of autonomous owners, each of which provides a set of nodes capable of hosting virtual machines. A single owner's nodes may be located at one site, or distributed across multiple sites. Second, there exist a set of service providers that implement and deploy network services across a set of virtual machines. These virtual machines may run on nodes that belong to more than one owner. The central question is how to manage the trust relationships between the owners and service providers, so as to provide effective management of the system as a whole.

A key insight to reducing such an NxN problem into an order N problem is to use trusted intermediaries. The PlanetLab Consortium is one such trusted intermediary: node owners trust it to manage the behavior of VMs that run on their nodes, and service providers trust it to provide access to a set of nodes that are capable of hosting their services.

### **CoDeeN: Content Distribution**

The overall CoDeeN project includes several components that either handle or generate large volumes of data. The main content distribution network component now handles over 25 million requests per day. On top of this component, we have built a mechanism for efficiently serving large files, named CoBlitz. This system is used by CiteSeer to serve documents, by the Fedora project to distribute their Linux images on CD & DVD, and by a number of other PlanetLab projects, including the Stork system, which uses it to distribute software packages for PlanetLab researchers. We are also in the process of building a publically-accessible sensor, modeled after the private ones we developed for the PlanetSeer project, that would allow other researchers to view summarized details about the traffic that traverses these projects. On the measurement front, we also developed and maintain the CoMon system, which monitors PlanetLab nodes and many aspects of the experiments that run on them.

All of these projects generate large volumes of traffic, with the most currently coming from CoDeeN's content distribution network. All requests are logged, and we record certain timing information along with summarized request history. Even at high compression rates, these logs are growing at over 5 GB/day. The CoMon logs are archived, and serve as a public history of PlanetLab activity. These logs are also compressed, and are growing at 5GB/month. We expect that when our sensor becomes public, its logs will also grow at the rate of a few GB per month. While the CDN logs have not been made public due to privacy issues, we have worked with other researchers to utilize them. The CoMon logs and sensors have been used by a number of other projects, and have also appeared in papers by other researchers.

### **Sound Lab**

Research was conducted in algorithms for analyzing and simulating audio, emphasizing real-time analysis for machine "understanding." Applications range from musical genre classification, song thumbnailing, audio scene identification (outdoors, bus station, board meeting, etc.), speaker identification, word spotting, etc. Flexible audio synthesis and resynthesis systems and software enable efficient coding and compression, using audio as a display component in interactive systems, and transmission/control of audio over networks.

## **Indicators of Success and Major Accomplishments**

ZebraNet:

- The ZebraNet efforts have included 2 major real-life deployments of the GPS-based tracking collars and software we designed. Both deployments occurred on Burchell's Zebras at Sweetwaters Game Reserve in Nanyuki, Kenya. The first deployment, in January 2004, and a subsequent deployment in June, 2005, collectively led to the gathering of thousands of sensed datapoints on zebra movements. These sets of position information have offered valuable new insights into the social behavior of these animals, and into their ability to use the landscape and compete for natural resources.
- The ZebraNet Project has also been selected to be featured in an NSF Special Report on Animal Tracking Research.
- Best Paper award. [12]

- Paper selected for Annual “Top Picks” in Computer Architecture issue of IEEE Micro. [7]
- Best Poster award. 2005 Richard Tapia Celebration of Diversity in Computing Conference.
- In addition to the sensor networks and biology research it has supported, our research has also been quite successful in involving undergraduates and under-represented minority groups as well.

#### Scalable Display Wall:

- Developed a novel data visualization tools for genomic data visualization that is dynamic and scale free [24, 25].
- Developed a multi-cursor window system for shared data visualization that allows multiple users to use a shared display wall simultaneously for collaborative [27, 25]. User can move and use X11 windows on a shared display wall simultaneously.
- Developed an improved VNC-based shared display system that allows multiple users to share information from Windows, Linux and Mac platforms independent of their data analysis applications. We have deployed the software tools to three Plasma physics lab control rooms [23]
- Develop an automatic alignment method to align a small-scale display wall system for shared information visualization [26].

#### PlanetLab:

- PlanetLab currently spans 670 machines at 327 sites in over 35 countries, where each of over 600 projects gets access to an isolated “slice” of PlanetLab’s global resources.
- PlanetLab has significantly lowered the barriers to deploying overlays. We continue to investigate how to lower them further by introducing a general data plane for overlay networks [28].
- To help the research community that uses PlanetLab, we identify some perceptions, discuss whether they are myth or reality, and describe best practices that can be applied to make PlanetLab as effective a research platform as possible [29].
- PlanetFlow, a network auditing service, is implemented to maintain comprehensive, permanent accountability for all traffic generated by PlanetLab services, in accordance with common Internet practice and the terms of the PlanetLab Acceptable Use Policy. PlanetFlow audits the usage of PlanetLab network resources in order to facilitate the resolution of complaints, limit liability, and minimize problematic behavior [30].

#### CoDeeN:

- The CoDeeN project receives over 25 million requests every day, from a user population exceeding 50,000 users [1].
- The traffic from this project has been used to develop new mechanisms to automatically detect malicious users, and mitigate their impact [2].
- The CDN has served as the starting point for a large-file transfer service, which can easily handle documents in the range of megabytes to several gigabytes, and which has been used to support a range of research and technical projects, including CiteSeer, the Fedora Core Linux distribution, and various PlanetLab researchers [3].
- The CoMon monitoring service has helped to monitor PlanetLab itself, as well as a number of the projects running on it. Researchers can use it to observe their own projects, as well as any impact from other projects that may be behaving abnormally [4].

#### Sound Lab:

- MARSYAS Open-Source software and systems for audio analysis. This software is now the starting point for nearly 80% of new projects in Music Information Retrieval, and is used by countless other research projects in general audio "machine listening." [19][20] Systems have been constructed for a wide variety of applications in machine understanding of audio.
- Chuck open-source software for real-time audio synthesis and processing. This software, which won the ACM Multimedia open-software competition in 2004, has a large and growing user and developer base. It is also being used to teach beginning programming at a number of institutions.
- GigaPopR: high-speed network software for real-time musical performances, providing the most strict tests of quality and latency. A concert held between Princeton University and McGill University used Internet2 and Ca2Net to connect multiple players at both ends in real time. [21]

- Publication of a textbook on audio synthesis for interactive applications [22]

## Publications/Patents

- 1) L. Wang, K. Park, R. Pang, V. Pai, L. Peterson. Reliability and Security in the CoDeeN Content Distribution Network. In *Proceedings of the USENIX 2004 Annual Technical Conference*, Boston, MA, June 2004.
- 2) K. Park, V. Pai, K. Lee, S. Calo. Securing Web Service by Automatic Robot Detection. In *Proceedings of the 2006 Usenix Annual Technical Conference*, Boston, MA, June 2006.
- 3) K. Park, V. Pai. Securing Web Service by Automatic Robot Detection. In *Proceedings of the Third Symposium on Networked Systems Design and Implementation (NSDI 2006)*, San Jose, CA, May 2006.
- 4) K. Park, V. Pai. CoMon: A Mostly-Scalable Monitoring System for PlanetLab. *Operating Systems Review*, 40(1), 2006.
- 5) B. Bisbeborn, M. Golightly, K. Park, V. Pai. ReDesign considerations for scalable large-file content distribution. In *Proceedings of Second Workshop on Real, Large Distributed Systems (WORLDS 2005)*, San Francisco, CA, Dec 2005.
- 6) Gilberto Contreras, Margaret Martonosi, Jinzhang Peng, Guei-Yuan Lueh, and Roy Ju. The XTREM Power and Performance Simulator for the Intel XScale Core: Design and Experiences. *ACM Transactions on Embedded Computing Systems*. 2006.
- 7) Q. Wu, M. Martonosi, D. W. Clark, V.J. Reddi, D. Connors, Y. Wu, J. Lee, and D. Brooks, "Dynamic Compiler Driven Control for Microprocessor Energy and Performance", in *IEEE Micro Special Issue: Top Picks from Computer Architecture Conferences*, Vol. 26, No. 1, February, 2006
- 8) Qiang Wu, Philo Juang, Margaret Martonosi, Li-Shiuan Peh, and Douglas W. Clark. Formal Control Techniques for Power-Performance Management in High-Performance Processors. *IEEE Micro. Special issue on Low-Power Processors and Technology (Sep/Oct 2005)*
- 9) Canturk Isci, Margaret Martonosi, Alper Buyuktosonoglu. Long-term Workload Phases: Duration Predictions and Applications to DVFS. *IEEE Micro. Special issue on Low-Power Processors and Technology (Sep/Oct 2005)*
- 10) James Donald and Margaret Martonosi. Techniques for Multicore Thermal Management: Classification and New Exploration. *33<sup>rd</sup> International Symposium on Computer Architecture*. Boston, MA. June, 2006.
- 11) Canturk Isci and Margaret Martonosi. Phase Characterization for Power: Evaluating Control-Flow-Based and Event-Counter-Based Techniques. *12th International Symposium on High-Performance Computer Architecture (HPCA)*. February, 2006.
- 12) Qiang Wu, Vijay J. Reddi, Youfeng Wu, Jin Lee, Daniel Connors, David Brooks, Margaret Martonosi, Douglas W. Clark. A Dynamic Compilation Framework for Controlling Microprocessor Energy and Performance. *38th Annual International Symposium on Microarchitecture*. Barcelona, Spain. Nov, 2005.
- 13) Canturk Isci and Margaret Martonosi. Detecting Recurrent Phase Behavior under Real-System Variability. *2005 IEEE International Symposium on Workload Characterization*. October, 2005.
- 14) Fen Xie, Margaret Martonosi, Sharad Malik. Efficient Behavior-driven Runtime Dynamic Voltage Scaling Policies. *International Conference on Hardware/Software Codesign and System Synthesis. (CODES+ISS)*, September, 2005.
- 15) Fen Xie, Margaret Martonosi, Sharad Malik. Bounds on Power Savings Using Runtime Dynamic Voltage Scaling: An Exact Algorithm and A Linear-time Heuristic Approximation. *10th International Symposium on Low Power Electronics and Design (ISLPED)*. August, 2005.
- 16) Gilberto Contreras, Margaret Martonosi. Power Prediction of Intel Xscale Processors Using Performance Monitoring Unit Events. *10th International Symposium on Low Power Electronics and Design (ISLPED)*. August, 2005.
- 17) Philo Juang, Qiang Wu, Li-Shiuan Peh, Margaret Martonosi, D. W. Clark. Coordinated, Distributed, Formal Energy Management of Chip Multiprocessors. *10th International Symposium on Low Power Electronics and Design (ISLPED)*. August, 2005.
- 18) Qiang Wu, Philo Juang, Margaret Martonosi, Douglas W. Clark. Voltage and Frequency Control with Adaptive Reaction Time in Multiple-Clock-Domain Processors. *Eleventh International Symposium on High-Performance Computer Architecture (HPCA)*. February, 2005. San Francisco, CA.
- 19) G. Tzanetakis and P. Cook "MARSYAS: a framework for audio analysis," *Organized Sound* 4(3), Cambridge University Press, 2000.

- 20) Tzanetakis, George and Cook, Perry "Music Analysis and Retrieval Systems", *Journal of American Society for Information Science and Technology* Volume 55, No. 12, pp. 1077-1083, 2004.
- 21) A. Kapur, G. Wang, P. Davidson, and P. Cook, Interactive Network Performance: A Dream Worth Dreaming? *Organised Sound*, 10(3), 2005.
- 22) Perry Cook, "Real Sound Synthesis for Interactive Applications," *AK Peters Press*, 2002.
- 23) Grant Wallace, Otto Anshus, Peng Bi, Han Chen, Yuqun Chen, Perry Cook, Adam Finkelstein, Thomas Funkhouser, Anoop Gupta, Matthew Hibbs, Kai Li, Zhiyan Liu, Rudrajit Samanta, Rahul Sukthankar, and Olga Troyanskaya. Tools and Applications for Large-Scale Display Walls. *IEEE Computer Graphics & Applications, Special Issue on Large Displays*. 25(4):24-33, July/August 2005.
- 24) Hibbs, MA, Dirksen NC, Li K, Troyanskaya OG. Visualization Methods for Statistical Analysis of Microarray Clusters. *BMC Bioinformatics*, 6:115, 2005.
- 25) Kai Li, Matthew Hibbs, Grant Wallace and Olga Troyanskaya. Dynamic Scalable Visualization for Collaborative Scientific Applications. Proceedings of *The Next Generation Software Workshop*. Denver, Colorado, USA, April 2005.
- 26) Grant Wallace, Han Chen and Kai Li. Automatic Alignment of Tiled Displays for a Desktop Environment. *Journal of Software*. 15(12): 1776-1786. December 2004.
- 27) Grant Wallace, Peng Bi, Kai Li, Otto Anshus. A MultiCursor X Window Manager Supporting Control Room Collaboration. Princeton University, Computer Science, *Technical Report TR-707-04*, July 2004.
- 28) Andy Bavier, Mark Huang, and Larry Peterson. An Overlay Data Plane for PlanetLab. Proceedings of *the Advanced Industrial Conference on Telecommunications*, July 2005.
- 29) N. Spring, L. Peterson, A. Bavier, and V. Pai. "Using PlanetLab for Network Research: Myths, Realities, and Best Practices". Proceedings of *the Second Workshop on Real, Large Distributed Systems (WORLDS '05)*, December 2005.
- 30) Mark Huang, Andy Bavier, and Larry Peterson. PlanetFlow: Maintaining Accountability for Network Services. *ACM SIGOPS Operating System Review*, Vol. 40, Issue 1, January 2006.

## Outreach Programs and Supported Students

### Supported Students:

- Michael Golightly, summer intern, officer of MAES at UCI. His work was involved in a workshop publication [5] (CoDeeN)
- Mentor for 2 students, CRA-W Distributed Mentoring Program. (ZebraNet)
- Ting Liu, Christopher M. Sadler, Yong Wang, Pei Zhang. (ZebraNet)
- George Tzanetakis (PhD Completed 2002) "Manipulation, Analysis and Retrieval Systems for Audio Signals", John Hainsworth (PhD Completed 2006), "CASTER: Systems for Collaborative Writing", Ge Wang (5th year PhD), "ChucK, an On-The-Fly Multimedia Programming Language"
- Han Chen, Peng Bi (female), and Matt Hibbs. (Scalable Display Wall)

### Outreach Programs:

- Publication, presentations (workshops), and course integration of open-source software (MARSYAS and ChucK).
- Courses offered included "HCI Technology" (Cook), "Audio DSP" (Cook), and "Pervasive Information Systems" (Cook and Wolf).
- Shared displays are becoming more prevalent in government labs and industry. Our aim is to help researchers use these displays more effectively for collaboration.

# CRI: PlanetLab: A Community Resource Development Program

## Proposal CNS-0454278, Project Year 2005

Princeton University

PI: Larry Peterson

llp@cs.princeton.edu

<http://www.planet-lab.org>

**Abstract.** Overlay networks have recently emerged as a promising technique for deploying new network services, and more generally, for introducing disruptive technology into the Internet. However, there is a pressing need for a testbed that lets the research community experiment with overlay services at scale, and under real-world conditions. We are building-out and enhancing a global overlay network, called PlanetLab, that provides just such a capability. At present, PlanetLab consists of 670 machines hosted at 325 sites in over 30 countries. The opportunity PlanetLab provides for real-world experimentation at an unprecedented scale has sparked an explosion of research activity, with 400 active projects transmitting an average of 4 terabyte of data a day on the current system. PlanetLab's community also continues to grow dramatically: since its inception in mid-2002 an average of a half-dozen Universities worldwide join PlanetLab per month. This grant funds our efforts to scale PlanetLab to over 1,000 nodes, and to expand our user base to literally every researcher or student who wants to conduct a wide-area experiment.

## 1 Introduction

The last few years have seen the emergence of a new class of network services, including file sharing and network-embedded storage [7, 9, 15], content distribution networks [8, 19], overlay services for routing, multicast, and QoS [1, 5, 18], scalable object location services [2, 13, 14, 17], anomaly detection mechanisms [6], fast and reliable DNS services [11, 12], network monitoring, characterization and monitoring tools [10, 20, 16]. One thing these services have in common is that they benefit from being widely distributed over the Internet—from having multiple vantage points from which applications can observe and react to the network's behavior, and from being in close proximity to many data sources and data sinks. The other thing they have in common is that they all add value to the network—they make applications more robust, responsive, and secure; they allow the network as a whole to support higher levels of throughput; they provide a richer set of capabilities than the current Internet; and they enable the network measurements that are essential to the research community.

To support the design and evaluation of such network services, we are expanding and enhancing a global overlay network called PlanetLab. PlanetLab is the result of a grass-roots effort to provide a shared overlay infrastructure to support research into network services. With seed funding from Intel and follow-on support from NSF, we began assembling PlanetLab in mid-2002, starting with 100 identical nodes distributed across 42 sites. New sites join by dedicating a set of x86 based machines connected to the Internet. Our installation and remote management software reduces a site's setup involvement to filling out a web form and then booting the machines using our PlanetLab BootCD. This has let us grow to 670 machines distributed over 325 sites across 30 countries since 2002.

The opportunity for real-world experimentation at an unprecedented scale has sparked an explosion of research activity, which will continue to grow as the testbed becomes more widely deployed. PlanetLab currently supports over 400 research projects. In the week before a recent conference deadline, those projects used PlanetLab to transmit several terabytes of data per day, while on average 4 TB of data is sent daily. Over seventy papers and reports have been published in major conferences and workshops sponsored by the ACM, the IEEE, and USENIX since 2003.

While PlanetLab has been a huge success by many measures, it continues to be limited by its size, the number of users it reaches, and the diversity of the links by which nodes are connected. We are working to scale PlanetLab to over 1,000 nodes and to expand our user base to literally every researcher or student who wants to conduct a wide-area experiment.

## 2 Goals, Objectives, and Targeted Activities

PlanetLab's key hardware infrastructure consists of commercial off-the-shelf (COTS) x86 server hardware, called *nodes*, of which there currently are 670 deployed at 325 sites. As mentioned in the Introduction, we are expanding

PlanetLab to over 1,000 nodes, effectively doubling PlanetLab's current size. Our long-term goal is to expand to 1,000 widely distributed sites with 2–3 nodes per site.

Our current model to foster PlanetLab's growth is that most of the hardware and network costs will be provided by the participating institutions. This model has worked well for research-oriented universities who join to gain access to the experimental facility provided by PlanetLab. Based on the average rate that universities have joined (six per month) since we opened PlanetLab to the research community, and the size of the "market" (i.e., academic research groups world-wide), we expect this stage to roughly double our current size over the next five years—from 325 to 600 sites.

While growth to 600 sites would be a great success in of itself, one criticism of PlanetLab in its current form is its lack of diversity: most nodes are hosted by academic institutions. To date more than 80% of PlanetLab nodes are hosted by academic research institutions that are connected to the Global Research and Education Network (GREN) — an interconnection network of high-speed networks such as Internet2 in the USA and Dante in Europe. A recent report [3] found that 70% of end-to-end measurements between PlanetLab node pairs represent measurements of the GREN rather than the "commercial" Internet as regular users see it. The path characteristics between these two networks are different: GREN's transmit and receive path is far more symmetric. To ensure that PlanetLab more closely resembles the Internet for research purposes, it is important to foster PlanetLab's diversity, which we cannot obtain using our current growth model. This requires seeding machines at strategic Internet sites.

This proposal lets us purchase roughly 400 new x86 based servers as nodes to drastically diversify PlanetLab's site composition and meet the resource requirements needed by our growing community of members. Note that 400 new nodes will let us add roughly 200 new sites, thereby making significant progress towards our long-term goal of 1,000 widely distributed sites while also diversifying our user community.

We are executing on our plan to foster PlanetLab's diversity and growth as follows:

- **VINI.** We are currently deploying hardware for an extension of PlanetLab called *VINI* [4], a Virtual Network Infrastructure that lets network researchers evaluate their protocols and services in a realistic environment while also providing a high degree of control over network conditions. VINI enables researchers to deploy and evaluate their ideas with real routing software, traffic loads, and network events. To provide researchers flexibility in designing their experiments, VINI will support simultaneous experiments with arbitrary network topologies on a shared physical infrastructure. Our initial VINI deployment involves placing 34 nodes at seventeen PoPs within Internet2 and the National Lambda Rail (NLR) networks.
- **EPScoR educational institutions.** We are deploying nodes to US-based educational institutions such as colleges, technical institutes and high schools. To date we have used the grant to purchase and ship machines five recognized Hispanic serving institutions. We are also refreshing hardware donated by Intel to various active minority-serving institutions and womens colleges.

### 3 Project Description: Community Resource Development

A critical element to PlanetLab's growth and diversity is our effort to build a community that encourages interaction between companies, researchers, students, and educators. We envision technology transfer interactions between companies and researchers at academic institutions, and a transfer of educational material between researchers at Universities and educators at non-research educational institutions. This section describes our plans for fostering community participation, exploiting unique educational opportunities exposed by PlanetLab, and encouraging international and industrial participation in the PlanetLab community.

PlanetLab is both a community resource, and a community project. That is, our larger goal is to engage the research community in the technical and architectural discussions about how PlanetLab evolves. We have already begun, and will continue to take, the following steps toward this end. Many of the items outlined below are already available on-line at [www.planet-lab.org](http://www.planet-lab.org).

- We provide the "boot software" that lets researchers add their own nodes to PlanetLab, thereby both increasing PlanetLab's size, and giving researchers access to a slice of PlanetLab in which they can run their own experimental services.
- We provide a web-based, source-code repository for contributing and accessing the research tools and network services developed on PlanetLab. In addition to supporting access to the software, some of these services will run continuously on PlanetLab, and made available to all users (not just network researchers) at sites that join PlanetLab.

- Software developed that becomes part of the PlanetLab core is released with under a liberal BSD-like license. We will also encourage the research community to fold useful services they develop into the PlanetLab core under the same terms.
- We manage interest-based mailing lists focused on various technical aspects of the project. These are an outgrowth of the working groups already forming around PlanetLab.
- We support a document series that provides a forum for debating and defining PlanetLab's architecture, as well as cataloging any architectural decisions that are made. This series, called PlanetLab Design Notes (PDN) is patterned after the Internet's Request For Comments (RFC) series.
- We will continue to hold technical workshops, co-located with major conferences like SIGCOMM, OSDI, and NSDI, to engage the larger networking community in a discussion about architectural issues. These meetings provide feedback to the steering committee, and help ensure that PlanetLab continues to meet the needs of the research community.
- Education and Outreach: In addition to encouraging the involvement of those researchers currently connected to PlanetLab, this proposal will let us expand the core community by exploiting the unique educational opportunities made possible by PlanetLab.
- International Participation: We are encouraging international involvement for the greater good of all PlanetLab participants. Note that machines purchased under this grant will *not* be given to foreign organizations.
- Industry Participation: Due to significant interest in PlanetLab throughout the networking and computing industry, we have formed a Consortium that both academic and commercial organizations can join. The expectation is that, like academic institutions, companies will contribute nodes, with the advantage of being spread over multiple sites, including international locations. Corporate membership fees provide funds that support network operations, although we are finding it much easier for companies to provide in-kind aid, particularly network bandwidth.

## 4 Indicators of Success and Major Accomplishments

We have three areas to report on: VINI, EPSCoR deployment, and incentive based growth/purchases.

**Initial VINI Deployment.** We are currently deploying 34 nodes at seventeen PoPs within the National Lambda Rail (NLR) and within Internet2's Abilene network as an initial deployment of VINI [4] (Virtual Network Infrastructure). VINI will enable controlled, realistic experiments with new network protocols and services. We are working with the NLR and Abilene Internet2 backbones to deploy VINI nodes that have direct connections to the routers in these networks and dedicated bandwidth between the sites. VINI will have its own globally visible IP address blocks, and it will participate in routing with neighboring domains. Our goal is for VINI to become a shared infrastructure that enables researchers to simultaneously evaluate new protocols and services using real traffic from distributed services that are also sharing VINI resources. The nodes at each site will initially be high-end PlanetLab servers, but may eventually be programmable hardware devices that can better handle a large number of simultaneous experiments carrying a large volume of real traffic and many simultaneously running protocols.

In addition to enabling new networking experiments that are currently not possible on PlanetLab, these nodes will also provide current PlanetLab researchers with access to the NLR facility, for example as a high-speed backbone network for the distributed services that PlanetLab supports. This has at least three synergistic benefits: (1) it allows existing PlanetLab services to explore new designs that include processing and storage "inside" the network backbone; (2) it gives researchers working on new network architectures an opportunity to carry real traffic generated by innovative distributed services; and (3) it gives current PlanetLab slices access to additional capacity to the commodity Internet.

**EPSCoR Deployment.** We are strategically seeding nodes at EPSCoR organizations to increase the diversity of the PlanetLab community. Specifically, we have used this grant (to date) to purchase and ship machines to the University of Puerto Rico (Mayaguez and Rio Piedras), University of Texas (San Antonio and El Paso), and University of New Mexico, which are recognized Hispanic serving institutions. For each year of this grant we plan to ship roughly 16 machines to universities/colleges either serving underrepresented groups or in EPSCoR states. We recognize the importance keeping in touch with these institutions. Dr. Fiuczynski keeps in touch with PIs at these institutions, ensuring that both the machines remain operative and are of use to them, as well as seeking potential collaborative opportunities.

**Incentive based Deployment.** The existing hardware infrastructure is currently underprovisioned relative to the number of researchers wishing to deploy services on PlanetLab. For this reason, it behoves us to provide the appropriate incentives to both existing sites and new sites to purchase hardware that far exceeds our current minimum hardware requirements. The NSF CRI grant has enabled us to do so in two ways:

- We negotiated a special price for HP DL320g4 servers for all US-based PlanetLab members that is roughly 40% below the retail price. These servers exceed our minimum hardware requirements by more than 3x. PlanetLab members may purchase machines via a PlanetLab specific HP eBusiness portal (<http://www.hp.com/buy/planetlab>). To date more than six US sites have upgraded their hardware.
- We are encouraging sites to purchase at least one HP DL320g4 (or equivalent) machine with a “buy one get one free” incentive. We have done this in a targeted manner with a few sites. However, to avoid alienating sites who have not been targeted, we are going to publically raffle a number of machines off on a monthly basis—i.e., we will give one node to sites (on a first-come first-serve basis) who can prove that they have within the timeframe of the raffle purchased an equivalent machine from HP or Dell.

## 5 Future and Strategic Directions

Extending the PlanetLab software and hardware base to support the VINI vision is a key focus of our current work. Our short-term goal with VINI is to build relationships with NLR and other high-bandwidth research networks, and to facilitate access to these resources by the networking and distributed systems research community. Our medium-term goal is to enable experimental validation of existing and new networking research ideas in a realistic setting. In the long-term, our goal is to revolutionize the research and teaching process by supporting the seamless migration of networking architectures from early prototype, through multiple design iterations, to a continuously running experiment that generates meaningful results by carrying real traffic on behalf of a large end-client community. In the end, we envision VINI serving as a microcosm for the next generation Internet.

In addition to this CRI grant, we have recently submitted an NSF MRI proposal for the VINI project. If funded, the MRI grant will pay for people to develop the VINI software as well as grow the relationships with VINI users. In contrast, this CRI grant primarily provides the necessary hardware budget for VINI.

Finally, PlanetLab’s founding charter member, Intel, donated machines to about a number of minority-serving institutions in 2003. We plan to refresh all of those machines located at these institutions by the end of 2006.

## References

- [1] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, pages 131–145, Chateau Lake Louise, Banff, Alberta, Canada, October 2001.
- [2] M. Balazinska, H. Balakrishnan, and D. Karger. INS/Twine: A Scalable Peer-to-Peer Architecture for Intentional Resource Discovery. In *Proceedings of the 1st International Conference on Pervasive Computing (Pervasive 2002)*, pages 195–210, Zurich, Switzerland, August 2002.
- [3] S. Banerjee, T. G. Griffin, and M. Pias. The Interdomain Connectivity of PlanetLab Nodes. In *Proceedings of the Passive and Active Measurement Workshop (PAM2004)*, Antibes Juan-les-Pins, France, April 2004.
- [4] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In VINI Veritas: Realistic and Controlled Network Experimentation. In *Proceedings ACM SIGCOMM 2006 Conference*, September 2006. To appear.
- [5] Y. Chu, S. Rao, and H. Zhang. A Case For End System Multicast. In *Proceedings of the ACM SIGMETRICS 2000 Conference*, pages 1–12, Santa Clara, California, June 2000.
- [6] B. Chun, J. Lee, and H. Weatherspoon. Netbait: a Distributed Worm Detection Service, 2002. <http://netbait.planet-lab.org/>.
- [7] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, pages 202–215, Chateau Lake Louise, Banff, Alberta, Canada, October 2001.
- [8] M. Freedman, E. Freudenthal, and D. Mazieres. Democratizing Content Publication with Coral. In *Proceedings of the 1st Symposium on Networked Systems Design and Implementation (NSDI 2004)*, San Francisco, California, March 2004.

- [9] J. Kubiawicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. OceanStore: An Architecture for Global-Scale Persistent Storage. In *Proceedings of the 9th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 190–201, Cambridge, Massachusetts, November 2000.
- [10] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet Background Radiation. In *Proceedings of the 2004 Internet Measurement Conference*, Taormina, Sicily, Italy, October 2004. To appear.
- [11] K. Park, V. S. Pai, L. Peterson, and Z. Wang. CoDNS: Improving DNS Performance and Reliability via Cooperative Lookups. In *Proceedings of the Sixth USENIX Symposium on Operating System Design and Implementation (OSDI)*, San Francisco, California, December 2004. To appear.
- [12] V. Ramasubramanian and E. G. Sirer. The Design and Implementation of a Next Generation Name Service for the Internet. In *Proceedings ACM SIGCOMM 2004 Conference*, Portland, Oregon, August 2004.
- [13] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Topologically-Aware Overlay Construction and Server Selection. In *Proceedings of the IEEE INFOCOM 2002 Conference*, pages 1190–1199, New York City, June 2002.
- [14] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proceedings of the 18th IFIP/ACM International Conference on Distributed Systems Platforms (Middleware 2001)*, pages 329–350, Heidelberg, Germany, November 2001.
- [15] A. Rowstron and P. Druschel. Storage Management and Caching in PAST, A Large-Scale Persistent Peer-to-Peer Storage Utility. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, pages 188–201, Chateau Lake Louise, Banff, Alberta, Canada, October 2001.
- [16] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A facility for distributed internet measurement. In *Proceedings of the 4th USENIX Symposium on Internet Technologies (USITS '03)*, Seattle, Washington, March 2003.
- [17] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: A Peer-to-Peer Lookup Service for Internet Applications. In *Proceedings ACM SIGCOMM 2001 Conference*, pages 149–160, San Diego, California, August 2001.
- [18] L. Subramanian, I. Stoica, H. Balakrishnan, and R. Katz. OverQoS: Offering Internet QoS Using Overlays. In *Proceedings of HotNets-I*, Princeton, New Jersey, October 2002.
- [19] L. Wang, V. Pai, and L. Peterson. The Effectiveness of Request Redirection on CDN Robustness. In *Proceedings of the Fifth USENIX Symposium on Operating System Design and Implementation (OSDI)*, pages 345–360, Boston, Massachusetts, December 2002.
- [20] M. Zhang, C. Zhang, V. S. Pai, L. Peterson, and R. Wang. PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-Area Services. In *Proceedings of the Sixth USENIX Symposium on Operating System Design and Implementation (OSDI)*, San Francisco, California, December 2004. To appear.

# CRI: Wireless Open-Access Research Platform (WARP): A Scalable and Extensible Testbed for High-Performance Wireless Systems

Proposal # CNS-0551692

Project Year: 2006-10

Rice University

Ashu Sabharwal, Behnaam Aazhang, Joe Cavallaro, J. Patrick Frantz and Ed Knightly

{ashu,aaz,cavallar,jpfrantz,knightly}@rice.edu

<http://warp.rice.edu>

## 1. Abstract

Wireless communications has become the focus of intense multi-billion dollar research all over the world. For the USA to establish world leadership in theoretical and experimental research, and its exploitation for deployed wireless systems, it is critical that we initiate strong international collaborative research programs with a focused agenda to deliver the next generation of wireless networks faster and cheaper. We envision an environment which promotes a holistic, yet rapid approach to the design of wireless networks. At the center of our vision is the design of a common platform, the Wireless Open-Access Research Platform (WARP), used by a community of researchers to perform deep collaborative research. The platform is carefully designed to meet the needs of next generation concepts, and address all layers and components in a wireless network design.

## 2. Introduction

Wireless is poised to make headlines again, with new speed records, emerging markets, and innovative new applications. In fact, many nations are choosing wireless as the predominant technology for their communication infrastructures. However, establishing global leadership in high-performance wireless requires a coordinated multi-disciplinary collaboration across multiple organizations.

The success of large collaborative projects depends on establishing common design, development and experimental research methodologies. Collaborative research platforms require specialized design tools that provide seamless cross-layer integration to facilitate complete end-to-end systems. High performance wireless network design relies on the convergence of multiple research communities, including hardware designers, computer architects, communication engineers and network protocol researchers.

## 3. Goals, Objectives, Targeted Activities

In this project, we will design a scalable and extensible wireless platform from the ground up, leveraging our extensive experience in hardware and tool design to create a true community platform. The project will include the design, development and dissemination of the platform to enable high-performance research throughout our multi-faceted wireless community.

**Infrastructure:** The WARP infrastructure is a complete system with custom hardware, support packages, design tools and application libraries. The hardware uses FPGAs for DSP-optimized development, where the number of FPGAs can be scaled as necessary to deliver increased computational power as needed. Multiple daughtercard slots support a wide range of I/O devices (MIMO radios, video I/O etc) and/or other co-processors. Finally, WARP programming tools and support packages will ensure that software stays compatible with future hardware revisions.

**Open-Access Repository:** We will initiate, maintain and grow the WARP open-access repository, which will contain building blocks for a wide range of wireless systems. The open-access repository will allow researchers to integrate standard building blocks with custom components to assemble complete systems.

**Partnerships, Education and Dissemination:** We will use three methods to develop a community of WARP users. The first will be focused workshops, which will be co-sponsored by Xilinx, Inc., to educate attendees on programming and extending WARP. The second will be research partnerships with multiple university and industry groups. Third is a student exchange program, where students will spend extended periods of time developing their research on WARP.

A key enabler of the project will be our longstanding, close partnership with the DSP research group and University Program at Xilinx, Inc, which has identified wireless as a key growth area for their FPGA systems. Xilinx has committed to support manufacturing, distributing and supporting the WARP communication kits. The communication kits will contain WARP hardware, Platform Support Packages and Development Tools, all of which will be donated to academic institutions through the Xilinx University Program. Further, Xilinx will also support the education program by hosting international workshops and helping to establish laboratories at interested universities.

#### 4. Infrastructure


WARP is a complete platform, designed from the ground up to meet the demands of high performance wireless systems research. The platform architecture, depicted below, consists of four key components:

**WARP Hardware:** Custom designed hardware tailored to the needs of high-performance wireless communications. The hardware provides resources *capable* of high speed DSP-intensive algorithm implementations, *scalable* interconnects for allocating more processing power as needed, and *extensible* peripheral options for radios, specialized processors and other expansion interfaces.

**Platform Support Packages:** Design tools and low-level interface modules designed to enable seamless use of the hardware by researchers at all layers of wireless network design.

**Open-Access Repository:** Central archive for all source code, models and hardware design files associated with WARP. The full contents will be available under a BSD-like open-source license.

**Research Applications:** Implementation and evaluation of novel algorithms, enabled by WARP's custom hardware and platform support packages. Full systems incorporating novel algorithms can be rapidly built using standard building block modules provided in the repository.




WARP's modular and layered architecture ensures that the hardware components can benefit from advances in processor technology, following the exponential increases in performance famously described by Moore's Law. The approach also enables algorithm implementations to be used on future hardware revisions without being re-designed, maximizing the impact of researchers' implementation efforts.

## 5. Project Description

**WARP Hardware Design:** The first component of the WARP design is custom hardware which provides the processing, interface and peripheral resources required to meet the goals described above. The WARP hardware will be based on designs already completed for the NSF-funded Transit Access Point project. The WARP hardware will benefit greatly by leveraging the TAP hardware design experience and from the interest the TAP platform has generated.

*WARP FPGA Board :* The main board in the platform is the WARP FPGA board. At the heart of this design will be a Xilinx Virtex-4 FPGA, the latest and most powerful available. These FPGAs are very well suited for the kinds of DSP-intensive operations which the various applications for WARP will require. For example, the Virtex-4 provides dedicated DSP slices, hardware blocks designed specifically for high speed multiply-accumulate and other DSP operations. These slices are a feature not available in the FPGA used by the TAP platform and will play a major role in implementing computationally intensive wireless algorithms on WARP. Virtex-4 FPGAs also provide flexible and fast interconnect options for interfacing peripherals and creating multi-processor systems, two primary requirements of the platform design. Some versions of Virtex-4 also include embedded PowerPC processor cores, providing an ideal resource for implementing higher layer algorithms better suited for general purpose processors than programmable logic.

*Daughtercard Slots :* While the FPGA itself provides substantial processing power, its connections to other devices and boards will enable the variety of applications targeted by WARP. The WARP FPGA board will provide four daughtercard slots, each wired to a large number of dedicated FPGA I/O pins. These slots will house peripheral cards. The slots will be flexible enough to support a wide variety of future peripheral designs, including multimedia interfaces and specialized auxiliary processors. The four slots will be functionally identical, allowing users to mount the combination of peripheral cards that best suits their application. The slot interface will be backwards compatible with the TAP daughtercard specification. This interface will be documented in the WARP repository, allowing WARP users to design custom daughtercards.



*Figure:* Current TAP hardware platform on which the WARP designs will be based, including the Virtex-II Pro FPGA board and two existing peripheral daughtercards.

**WARP Peripheral Daughtercards:** The WARP FPGA board provides four daughtercard slots. Each slot is connected to a dedicated bank of I/O pins on the FPGA, providing a flexible, high-throughput interface. The daughtercard slots also provide a combined 30 watts of power isolated from the FPGA supplies. These daughtercard slots provide the opportunity to design a wide variety of custom daughtercards, ranging from peripheral interfaces like radios and Ethernet to auxiliary processing modules, built around DSPs or specialized processors. We have already designed three daughtercards for the TAP project which will be compatible with the first revision of the WARP hardware. The first is a radio interface, see figure above, built around a single chip direct-conversion radio transceiver from Maxim Integrated Products. The Maxim radio supports both the

full 2.4GHz and 5GHz ISM bands, including support for channels in these bands which are not available in the USA but are open in other countries. The radio is designed for MIMO applications, guaranteeing the phase coherency of carriers generated by chips which share a reference clock. The first revision of a radio board for the TAP platform has already been completed and is fully functional. The TAP radio daughtercard will be the first of many potential wireless interfaces (UWB, DSRC, DVB, WiMAX, etc.) which will be supported by the WARP hardware.

The second daughtercard provides analog video capture and playback capabilities (see Figure). The third daughtercard provides USB 1.1 and fast Ethernet interfaces via dedicated microcontrollers.

**Open Specification:** The specification for the daughtercard interface will be part of the WARP repository. We envision that a whole range of daughtercards will be designed over time, both at Rice and by other WARP users. Daughtercards with conventional processors, like a TI DSP, or ASIP processors, like Imagine or TTA, would provide a platform for advanced research in computer architectures and design partitioning. Cards with other peripheral interfaces, like HD video or digital audio, would enable sophisticated multimedia applications on WARP. In short, WARP's daughtercard interfaces provide immense *\emph{extensibility}*, future-proofing the platform as new applications and research interests emerge.

**Platform Support Packages:** WARP will require a number of tools and supporting packages in order to enable the functionality of all its interfaces and peripherals. These support packages will expose the platform's full capabilities to researchers working at various layers. These will include low-level support, logic-PowerPC support, peripherals support and board-to-board support.

## 6. Indicators of Success and Major Accomplishments

Our major indicators of success are (a) operational infrastructure, (b) adoption by both academic and industrial organization and (c) well-populated repository. Below, we highlight our major accomplishments to-date.

**Functional WARP Hardware:** The first version of WARP hardware, both FPGA and radio boards, have been successfully tested in the labs. The new WARP boards have a 10/100 ethernet access to the board which is being used to stream real traffic into the WARP. Furthermore, a clock circuit has also been incorporated, which allows clock synchronization of multiple radio boards to build advanced multiple antenna systems (all next generation wireless systems use multiple antenna technology).

**Operational Platform Support:** A beta version of the platform support for physical and medium access layers has been tested. Labeled as PHY driver, we can now connect any physical layer on the FPGA fabric to an arbitrary MAC protocol on the PowerPC. This version of the platform support is being tested and enhanced in Xilinx Labs, using two summer interns from Rice and two engineers from Xilinx Tools Division.

**Basic Research Applications:** A 12 Mbps physical layer based on OFDM systems is now functional for short distance communications, along with a random-access protocol based on Aloha. These will become basis for innovative wireless research as other groups adopt the platform.

**WARP Repository:** The warp repository is now live (<http://warp.rice.edu>) and contains all our work to-date.

## 7. Publications

1. P. Murphy, A. Sabharwal and B. Aazhang, "Design of WARP: A wireless open-access research platform," *to be presented at 14<sup>th</sup> European Signal Processing Conference*, Sept 2006, Italy.
2. M. Duarte, R. Rao, C. Dick and A. Sabharwal, "A high throughput beamforming architecture for MIMO systems," *to be presented at Asilomar Conf. on Signals, Systems and Computers*, Nov 2006.
3. C. Hunter, S. Gupta, P. Murphy, C. Dick, A. Sabharwal, "A Flexible Framework for Wireless Medium Access Protocols," *to be presented at Asilomar Conf. on Signals, Systems and Computers*, Nov 2006.

## 8. Outreach Programs and Supported Students

**Professor's Workshop:** Towards preparation for widespread distribution of the WARP hardware, we organized our first workshop at Rice University for Professors. The 17 participants represented experienced research groups (UCSD, Naval Postgraduate School, Worcester Polytechnic), newly initiated research groups (Minnesota Twin Cities, U Maine, Auburn, U. North Florida, Devry University) and predominantly-minority schools (Prairie-view A&M U, Montana State U, Texas A&M Kingsville).

**Undergraduate Research:** For last 4 years, CMC Labs (headed by PI Sabharwal) has included increasingly more undergraduate students in the research activities. In fact, it will be fair to say that almost 80% of the platform support packages (software glues together all the hardware components) has been developed by the undergraduate students. A large number of students join the research projects in their junior year and continue to work during summer under REU programs till they graduate. During the two years, they are given challenging open-ended problems and creative freedom to complete the tasks. The work environment is structured, with weekly online reports and group meetings, supervised by senior graduate students and faculty.

**International Collaborations:** To initiate international collaborations, two international workshops will be held later this year to train professors in India and Taiwan on WARP hardware. In addition, two international undergraduate students from India are working in CMC Labs, Rice University this summer.

**Supported Students:** Currently following students are partially supported by the CRI grant – Patrick Murphy, Yang Sun, Jingpu Shi and Gareth Middleton.

## 9. Future and Strategic Directions

The WARP platform has gained significant popularity in several communities – networking researchers, computer architecture, physical layer design and industrial research labs. Our future plans are to ensure large impact on the multiple research communities and will include the following major efforts:

**GENI:** WARP is strategically positioned to serve a key role in the planned large-scale networking infrastructure called GENI. The PIs are part of the GENI Wireless Working Group, and are helping shape the programmable wireless component of the GENI platform. Widespread distribution of the hardware and open-access research applications are directly aligned with GENI goals.

**Hardware Distribution:** We are planning to deliver hardware to 4 universities by end of 2006. The target universities include UC San Diego, RPI, UIUC and IIT Delhi (India) as part of beta development phase. This will ensure a global impact and exploit the varied expertise of these schools.

**Workshops:** Xilinx has already agreed to support two international workshops – in IIT Delhi (India) and NCTU (Taiwan), to train professors on WARP hardware and design flow. In addition to the international workshops, two domestic workshops are also being planned (potentially on east and west coasts).

**Growing the WARP Repository:** Rice, Xilinx and partnering universities will be contributing many software components for both research and education to WARP repository over next year to build a critical mass. We intend to make this process easier for all participants, and allow online collaboration across multiple organizations. In addition, we will also be posting packet capture from real wireless channels which can be used by researchers to perform research without the hardware.

# Collaborative Research: Planning Proposal: CRI: Community Resources for Research in Automated Authorship Attribution

Rutgers University  
Proposal # 0454126  
David Madigan  
dmadigan@rutgers.edu

Illinois Institute of Technology  
Proposal # 0454197  
Shlomo Argamon, David D. Lewis  
argamon@iit.edu,  
davelewis@DavidDLewis.com

<http://www.stat.rutgers.edu/~madigan/>

<http://lingcog.iit.edu/~argamon/>

<http://www.DavidDLewis.com>

## 1. Abstract

Homeland security and the criminal and civil justice systems increasingly require reliable and valid methods to identify the authors of anonymous documents. Further demand arises from fields as diverse as computer forensics and literary studies. Unfortunately, there are substantial disagreements in the authorship attribution community about the effectiveness, or even the legitimacy, of various techniques, and few comparable results to settle these debates. The scattering of researchers and publications across a wide range of fields compounds the confusion. Our one-year NSF planning grant is supporting the development of a detailed plan for test corpora, evaluation frameworks, and associated resources that should lead to a vast improvement in the quality and coherence of research on automated authorship attribution.

## 2. Introduction/Background

Individuals have distinctive ways of speaking and writing, and there is a long history of linguistic and stylistic investigation into authorship attribution. In recent years, practical applications for authorship attribution have grown in areas such as intelligence (linking intercepted messages to each other and to known terrorists), criminal law (identifying writers of ransom notes and harassing letters), civil law (copyright and estate disputes), and computer security (tracking authors of computer virus source code). This activity is part of a broader growth within computer science of identification technologies, including biometrics (retinal scanning, speaker recognition, etc.), cryptographic signatures, intrusion detection systems, and others. As with these other technologies, there is increasing interest in the accuracy and scientific validity of determinations of authorship by both human experts and automated methods.

Automating authorship attribution promises more accurate results and objective measures of reliability, both of which are critical for legal and security applications. Recent research has used techniques from machine learning, multivariate and cluster analysis, and natural language processing in authorship attribution. These techniques have also been applied to related problems such as genre analysis and author profiling (such as by gender or personality). This work is scattered among the literatures of several areas of computer science (machine learning, computational linguistics, computer forensics), as well as statistics, linguistics, and literary analysis. Methodological inconsistencies, unsettled controversies, and redundancy are rife.

In forensic text analysis, where quantifiable effectiveness is now a *sine qua non* (due to the *Daubert* decision), disputes among researchers (and practitioners) have sometimes become quite heated. These disputes concern the use of different feature sets, analysis techniques, and evaluation methodologies. There is a great deal of expert judgment involved in forensic text analysis, which would be acceptable if there were generally accepted standards in the field. However, there is very little agreement in the field over what level of control is necessary in confusion sets, what sorts of features are most reliable, what statistical tests to use, and most critically, how to determine the reliability of a proposed attribution. This state of affairs underscores the necessity both for a standard corpus and tasks, as well as for forums that bring together researchers to share results.

### 3. Goal/Objectives/Targeted Activities

The primary goal of the planning project is to identify those research resources whose development would best move the field of authorship attribution forward, by enabling standardized evaluation of research results and comparative evaluation of different techniques, as well as bringing together researchers from different fields working on similar problems in this area. Targeted research activities include the main area of developing more accurate and well-understood computational techniques for attributing the authorship of an unknown text, and also work on developing better theoretical models of how texts are constructed (from linguistic, cognitive, and social perspectives), so that aspects of word and syntax usage that relate to authorship may be distinguished from those deriving from a text's topic, genre, or social context. Targeted educational activities are closely linked to these targeted research activities – we expect that resources that are developed to support research on automated authorship attribution will also be used to support a variety of related educational activities in natural language processing and information retrieval.

### 4. Infrastructure and Facilities

As this is a planning project, we are developing a set of recommendations for future construction of appropriate research infrastructure and facilities, in the form of standardized test collections and an evaluation framework for the authorship attribution community. An interim draft of these recommendations is given in Section 6 below.

### 5. Project Description

The primary activities we have undertaken so far under the planning grant are:

- Begin assembling a bibliography and public resources for an eventual community website
- Forming a Working Group of leading researchers and practitioners in authorship attribution
- Leading email-based discussions on authorship attribution and resources to support progress in it
- Holding a two-day workshop for Working Group members to develop a detailed plan for community resources

The current members of the Working Group are:

- Shlomo Argamon (Illinois Institute of Technology, Computer Science)
- David Banks (Duke University, Statistics)
- Douglas Biber (Northern Arizona University, English)
- Carole Chaski (Institute for Linguistic Evidence, Georgetown, DE)
- Bruce Croft (University of Massachusetts at Amherst, Computer Science)
- Haym Hirsh (Rutgers University, Computer Science)
- David Holmes (College of New Jersey, Mathematics & Statistics)
- David Hoover (New York University, English)
- Patrick Juola (Duquesne University, Mathematics and Computer Science)
- Jussi Karlgren (Swedish Institute of Computer Science)
- Moshe Koppel (Bar-Ilan University, Computer Science)
- Dave Lewis (David D. Lewis Consulting, Chicago, IL)
- David Madigan (Rutgers University, Statistics)
- Gerald McMenamin (California State University at Fresno, Linguistics)
- Joseph Rudman (Carnegie Mellon University, English)
- Ian Soboroff (Information Access Division at NIST)
- Hans van Halteren (Radboud University Nijmegen, Language and Speech)

Our *Workshop on Developing Community Resources for Automated Authorship Attribution* was held June 8-9, 2006 at DIMACS, Rutgers University, Piscataway, NJ. Eleven members of the Working Group were able to attend.

The consensus of the workshop was that progress in authorship attribution, both intellectual and practical, has long been severely limited by the lack of appropriate community resources. Drawing on discussions on the Working Group mailing list, the Working Group members in attendance produced three major recommendations for community resources in authorship attribution, which we discuss below. Investment in these three areas would support rapid progress in both the intellectual development and practical use of authorship attribution.

*[Note: The details of the recommendations discussed below (Section 6) are tentative. The very short time between the end of our workshop and the preparation of this report has not allowed time for review and iteration with Working Group participants, particularly those who could not attend the workshop.]*

## **6. Indicators of Success and Major Accomplishments**

The major accomplishments of our project to date are the recommendations for resources summarized in the following sections.

### ***6.1. Recommendation: Data Sets for Authorship Attribution Research***

Working Group participants identified lack of high quality, shared data sets as the most serious factor limiting progress in authorship attribution. The challenges in data set construction for authorship attribution are unusually severe among language processing fields. Some characteristics of text that are of little importance in most areas of language processing (punctuation, capitalization, spelling variants and errors, editorial changes among editions, etc.) are of crucial importance in identifying authorship. Standard methods of preprocessing raw text may leave it unusable for authorship attribution work. Conversely, authorship attribution data often requires additional markup beyond that needed for other forms of language processing, for instance to distinguish quoted from authorial material. Unfortunately, authorship attribution has seen little of the investment in data sets that has benefited tasks such as speech recognition, information extraction, text retrieval, summarization, and question answering.

The Working Group discussed a range of challenges in producing data sets for authorship attribution including availability of materials with particular properties, cost, data input for printed materials, formatting, markup, privacy concerns, copyright, the role of blind datasets and availability of material on the web, and other factors. These issues will be discussed in detail in the final report of this project.

The broad range of research pursued in authorship attribution means that no single data set is appropriate for all researchers. Further, any given public data set will capture only some aspects of applied authorship attribution tasks, and multiple data sets are needed to explore the full range of issues. The workshop suggested three initial types of data sets of great value, one historical and two modern.

#### **6.1.1. Data Sets Based on Historical Material**

The majority of work on authorship attribution has until recently focused on the authorship of disputed literary and historical materials. Resolving authorship of such materials can make great contributions to scholarship, and sometimes captures the imagination of the general public. Further, historical materials have advantages over modern ones on some dimensions, such as detailed study of their context, and reduced legal concerns.

Since authorship attribution workers in the humanities are at least as much interested in particular authors, periods, or genres as in the development of authorship attribution technology, it is impossible to come up with a dataset that will be of interest to all of them. The proposal arrived at by the Working Group for the first data set was to focus on Victorian era English writers. The rationale here is that there is both a large quantity of known material (for testing methods), and a large body of anonymous fiction published in newspapers and magazines of the period. Historical records are available that show some of this fiction to be by important writers of the period (Dickens and Thackeray, for instance), but no authorship records survive for much of the material. Thus there is the opportunity both for pushing attribution technology, and applying it to cases of substantial intellectual and general interest.

## **6.1.2. Data Sets Based on Modern Materials**

In contrast to historical texts, modern texts have many of their own advantages: huge volumes of material in a much wider range of genres (and often in digital form), relatively clear ownership for materials under copyright, and the possibility of commissioning writings from living authors. Participants saw the need for (at least) two distinct datasets based on modern materials, supporting very different kinds of research.

### ***6.1.2.1. A Massive Attribution Dataset***

Traditional authorship attribution tasks, both in humanities and in forensic, have dealt with relatively small numbers of authors. An important thrust of research (of great interest for intelligence applications) is to expand the scope of authorship attribution to monitoring masses of data for large numbers of potential authors. Such tasks require advances in both the types of clues monitored and in the effectiveness and efficiency of the multivariate statistical methods applied. (Indeed, authorship attribution is one of the few machine learning tasks where the number of classes can be meaningfully increased essentially without bound.)

The explosion of networked public communication forums enables large scale authorship attribution experiments that were impossible until recently. Workshop participants reported studies they had carried out on blogs and listserv postings among others. Blog postings have some advantages with respect to (comparatively) clear authorship, and some are accompanied by demographic information useful for profiling experiments. A number of university projects are already gathering large blog archives, and software is available which can gather large data sets quickly. Preparing this data in a form suitable for authorship attribution will take substantial additional work to remove non-authorial material and inadvertent non-stylistic clues. Participants nonetheless felt that a large scale blog collection could be used as part of an initial evaluation run by early 2007.

### ***6.1.2.2. A Focused Dataset for Exploring Generalization***

One of the greatest difficulties in applied authorship attribution problems are mismatches along dimensions other than authorship, both among the texts of known authorship, and between those materials and the texts to be attributed. Factors such as genre, chronological period, topic, and intended audience are all known to affect some of the same stylistic characteristics used in determining authorship. Further, stylistic markers are not expressed randomly among authors, but correlate with gender, upbringing, education, and a variety of other demographic and life history variables. Whether the available evidence is sufficient to conclude a contested or anonymous text was written by a particular author, or just by someone with similar background, is a difficult question. The above issues have made authorship attribution testimony in court cases a highly controversial issue.

Research on how these factors impact the reliability of authorship attribution, and development of stylistic markers which are unaffected by such mismatches when possible, requires data sets that allow controlling some factors while systematically varying others. This requires having material for the same authors for many genres, topics, audiences, etc. Demographic and life history information should be gathered on the authors as well. Working with living authors and rights holders provides the best opportunity to obtain such materials. One may even be able to solicit authors to produce texts based on an experimental design, though the stability of stylistic markers in solicited texts must itself be studied.

## ***6.2. Recommendation: An Evaluation Framework for Authorship Attribution Research***

There was a consensus at the workshop on the desirability of a community evaluation framework for authorship attribution, analogous to language processing efforts such as DUC (<http://www-nlpir.nist.gov/projects/duc/>), MUC ([http://www-nlpir.nist.gov/related\\_projects/muc/](http://www-nlpir.nist.gov/related_projects/muc/)), TREC (<http://trec.nist.gov/>), CLEF (<http://clef.iei.pi.cnr.it/>), NTCIR (<http://research.nii.ac.jp/ntcir/>), and many others. While the details of these evaluations vary, activities commonly undertaken by the evaluation organizers include:

- Obtaining materials and negotiating license agreements
- Preprocessing and annotating data
- Specifying aspects of experimental design (e.g. training/test splits) that must be used by all participants

- Defining effectiveness measures
- Running a periodic conference, restricted to evaluation participants (and representatives of sponsors), at which to report results
- Obtaining agreements from participants not to use the results in ways which would be likely to discourage future participation (e.g. certain restrictions on the use of results in advertising)

Authorship attribution was seen as an area that would particularly benefit from such an evaluation. The diversity of materials studied, and the degree to which authorship attribution effectiveness is affected by details of preprocessing, has meant that published results have rarely been comparable. A large number of researchers analyzing shared datasets, produced under carefully described conditions, would aid greatly our understanding of current authorship attribution techniques. The selection of training and test materials by evaluation organizers, and the requirement that systems be frozen before test materials are released, help fight the unconscious biases when individual researchers choose materials and define stylistic markers. Finally joint evaluation frameworks and the associated conferences in the past have proven effective at bringing together researchers with common interests from distinct fields. Indeed, a small previous evaluation effort in authorship attribution, Juola's "Ad-hoc Authorship Attribution Competition" in 2004, provided some benefits of these sorts.

Many details of the proposed evaluation are still under discussion, and would be heavily influenced by what sponsorship could be found for this effort. There was, however, considerable interest in a relatively straightforward preliminary evaluation using easily available materials in rough form (perhaps blogs) with a timetable that would allow results to be discussed at the Digital Humanities conference at the University of Illinois Urbana-Champaign in June 2007.

### ***6.3. Recommendation: Web-based Library of Authorship Attribution Resources***

A third simple but important resource that is currently lacking is a web-based repository for bibliographies, data sets, and results in authorship attribution. A challenging issue here is to encourage researchers to submit enough detail on their data sets so that other researchers can both replicate their results, and evaluate the soundness of decisions (e.g. choices of which materials to include) that were made. If the submission process is made too cumbersome, then researchers will simply choose not to contribute to it. It was recognized that improving the standards of reporting on data preprocessing issues would require a community effort to, for instance, give better feedback when reviewing papers submitted to conferences and journals.

## **7. Publications/Patents**

No publications have been produced under this grant to date. We plan to seek the publication of the final report on our activities under this grant in a major forum appropriate for authorship attribution research.

## **8. Outreach Programs and Supported Students**

No specific outreach programs have been developed in the course of this project to date. One student at Rutgers has assisted with website management. Otherwise no students have been supported under this project; after the development of the draft recommendations by the Working Group, we will be involving one or two students in developing the first version of a Web-Based Library, as described in Section 6.3 above.

## **9. Future and Strategic Directions**

The publication of the complete recommendations of the Working Group will help to integrate the currently fragmented world of authorship attribution research. Immediate next steps will be the implementation of these recommendations, including development of standard testbeds and evaluation frameworks. Also key will be seeking appropriate funding from NSF and other sources to support these efforts.

# RI: MultiStore: A Research Infrastructure for Management, Analysis and Visualization of Large-Scale Multidimensional Data Sets

*INSTITUTION: State University of New York at Buffalo*

*Investigators: A. Zhang (PI), R. Acharya, A. Garg, D. Mark, R. Miller*

*Period: 9/1/01-8/31/06*

Emails: {azhang, agarg, [miller@cse.buffalo.edu](mailto:miller@cse.buffalo.edu), [dmark@geog.buffalo.edu](mailto:dmark@geog.buffalo.edu),  
[acharya@cse.psu.edu](mailto:acharya@cse.psu.edu)}

**Website:** <http://www.cse.buffalo.edu/DBGROUP/Infrastructure/index.htm>

**Abstract:** This project is to establish MultiStore, a massive storage system for supporting integrated research in specific targeted areas of Computer Science, including Multimedia, Visualization, Geographical Information Systems (GIS) and Bioinformatics. The research objective is to develop computational theories and algorithms for storing, managing, analyzing, querying and visualizing multi-dimensional data sets that are generated from the related fields.


**Introduction/Background:** With the advancement of new technology in many fields, large volumes of scientific data are being generated and archived. At SUNY-Buffalo, for instance, biomedical imaging routinely provides multi-dimensional images, ranging from microscopic biological structures to macroscopic human organs; the advance of spatial information technology results in geospatial data which is becoming available at an explosive rate; DNA microarrays now can be used to determine simultaneously the expression levels of the thousands of genes which determine the function of all cells; and new computational methods such as Shake-and-Bake (SnB), which is a major advancement in biological structure determination, generate extensive amounts of data before a solution is determined. The common nature of these data sets is that they are high dimensional vectors with numerical, functional or symbolic attributes. The archiving and management of these data sets require substantial amount of storage capability (in terabytes). No facility available at SUNY-Buffalo before this grant could provide such storage capability. In addition, the high dimensionality and enormous size of such data sets pose very challenging problems in management, analysis, retrieval and visualization of the data sets. Pre-existing data processing techniques were inadequate for solving these problems.

**Goals, Objectives, Targeted Activities:** The research objective is to develop computational theories and algorithms for storing, managing, analyzing, querying and visualizing multi-dimensional data sets that are generated from the related fields. The research components include: (1) Data storage and management. We develop approaches to manage large-scale multi-dimensional data sets. Particular research issues include: multi-dimensional data storage, indexing, and clustering. (2) Data visualization. We develop effective graphics and visualization techniques that can help the user in information processing tasks. Particular research topics addressed include graph visualization and detecting clusters in a multidimensional data set through visualization. The visualization tools will be used in biomedical image understanding and analysis. (3) Data analysis and querying. We focus on geographical image understanding, analysis and querying. The particular research issues include geographical metadata/knowledge extraction, geographical metadata/knowledge representation and management, and geographical metadata/knowledge querying. (4) Data mining and bioinformatics. We develop data mining techniques for determination of protein structures and detection of gene expression patterns. Through these research activities, the fundamental understanding and novel techniques will be provided to support the management of various large-scale multi-dimensional data sets.

**Infrastructure and Facilities:** The MultiStore infrastructure establishes a common platform to carry out the campus-wide research activities in Computer Science and Engineering, Geographical Information Systems and Bioinformatics. The MultiStore infrastructure consists of two main components: a massive storage system and ancillary networking facilities. The storage system is housed in the Computer Science and Engineering Research Center as the central data management resource, and the networking facilities is used to support high-speed connections between the central storage system and other research laboratories involved the project. The whole

system consists of three components: data storage facilities, compute/file servers, and database servers. This storage system provides a total storage capacity of roughly 25.0 terabytes (TB) for the research activities involved.

**Project Description:** This project establishes a massive storage system for supporting integrated research in Multimedia, visualization, geographical information systems (GIS) and bioinformatics. The goal is to develop computational theories and algorithms for storing, managing, analyzing, querying and visualizing multi-dimensional data sets. Our research provides powerful tools for accessing and analyzing large amounts of domain data. The research activities are led by 17 faculty members from seven departments and Institutes. With the new MultiStore infrastructure, we are able to extend the scope of the collaboration among the researchers. The MultiStore infrastructure brings together the researchers in this project to share their data, and create a system environment to promote the interaction among the researchers. The hierarchy of the overall structure of the projects shown below:



#### Indicators of Success and Major Accomplishments:

##### Ideas


Pattern analysis and visualization approaches are being developed for high-dimensional and large-scale data sets generated from geographic image features, gene expression data, and crystal structure parameters. In particular, we have developed a new model called EPD (empirical pattern detection), which includes a series of pattern measurements and an iterative adjustment approach to delineate empirical sample pattern quality with informative genes. The research is motivated by the needs of emerging high-dimensional microarray gene data analysis applications. Our approach is designed to improve the unsupervised empirical pattern detecting performance for microarray gene data sets. The model integrates techniques of statistics (pattern measurements), data mining (the iterative pattern adjustment algorithm), and machine learning (stimulated annealing technique) to detect significant patterns within samples while dynamically selecting informative genes which manifest the empirical interest. We have conducted experiments on various microarray data sets, including Leukemia, Multiple Sclerosis, Colon Cancer, Lymphoma, Fibroblasts, and Hereditary Breast Cancer data sets, to demonstrate the effectiveness of our approaches. Empirical evaluation shows that our approach outperforms the existing conventional clustering approaches, and our model appears to be a promising approach for unsupervised analysis of genomic data sets.

We have also developed a novel exploration approach to give users highly confident indications of the existence of coherent patterns in time series data. To derive a coherent pattern index, we devise an attraction tree structure that summarizes the coherence information between genes in the data set. The attraction tree effectively discloses the hierarchical relationship between data objects within one cluster (inner-cluster relationship) based on our novel similarity measure. This intermediate view of the hierarchy of cluster structures is highly innovative and critical for interactive exploration of time series data analysis.

We recently developed a novel essential component identification model, bridging centrality, based on information flow and topological locality in scale-free networks. Bridging centrality provides an entirely new way of scrutinizing network structures and measuring components' importance. We applied bridging centrality on real world networks, including one simulated network, two biological networks, two social networks, and one web network, and show that the nodes distinguished by bridging centrality are well located on the connecting positions between highly connected regions through analyzing the clustering coefficient and average path length of those networks. Bridging centrality can discriminate bridging nodes, the nodes with more information flowed through them and locations between highly connected regions, while other centrality measures can not.


## Tools.

We have developed VizStruct, a dynamic interactive visualization environment for effective and efficient pattern analysis of large-scale multi-dimensional data sets. Visualization enables us to find structures, features, patterns, and relationships in a dataset by presenting the data in various graphical forms with possible interactions. We developed a new paradigm for multidimensional multivariate information visualization, termed Fourier harmonic approach. This approach is an extension of radial coordinate visualization from a signal processing perspective. We offer a thorough investigation of the projections and provide a detailed proof of their properties which gives a profound theoretical basis for our approach. We realize the unique nature of gene expression datasets: asymmetric dimensionality and two-way analysis, and propose two visualization processes for analysis. In the explorative visualization process, we focus on sample space and present an interactive visualization clustering approach to classify different samples based on variations in gene expressions. In the confirmative visualization process, we concentrate on gene space and apply FHPs in conjunction with global visualization techniques to investigate the structure of the dataset, to find patterns, outliers, and possible clusters. Experimental results have demonstrated that the proposed models are superior to Sammon's mapping on sample classification and consistent with other clustering methods such as SOM.



## VizStruct

We have also developed GeneXplorer, a tool which provides interactive exploration of coherent gene expression patterns. It has the following features: (1) It presents a hierarchy of the coherent patterns and the corresponding co-expressed genes. (2) It uses the coherent pattern index graph as an exploration "compass" to illustrate the local clustering structure and guide the users in the exploration. (3) It provides the flexibility to plug-in various components such as various distance measures and different density definitions. By using various plug-in components, users can compare various approaches and select the one fitting the data set the best. (4) It contains a gene annotation panel. Using the panel, users can connect the derived groups of co-expressed genes with some public annotation databases. (5) It supports an integrated environment of multiple graphical views, such as parallel coordinates, to visualize the data set from different aspects and in different resolutions.



The figure above shows the hierarchy of co-expressed gene groups in the Iyer's data set generated by our GeneXplorer tool.

#### Publications/Patents:

##### Patents:

- (1) A. Zhang, G. Sheikholeslami and S. Chatterjee, WaveCluster: Wavelet-Based Clustering Method for Managing Spatial Data in Very Large Database, Patent number: 6,882,997, date issued: April 19, 2005.
- (2) A. Zhang, M. Ramanathan, Y. Cho, and W. Hwang, Bridging Centrality: A Concept and Formula To Identify Bridging Nodes In Scale-Free Networks, Provisional Patent Application, May 19, 2006.

##### Publications:

- (1) A. Zhang, Advanced Analysis of Gene Expression Microarray Data, World Scientific Publishing Co., Inc. 2006 (to appear).
- (2) K. Bhasi, L. Zhang, D. Brazeau, A. Zhang, and M. Ramanathan, VizStruct for Visualization of Genome-wide SNP Analyses, Bioinformatics, Oxford University Press.
- (3) W. Wang and A. Zhang, Extracting Semantic Concepts from Images: A Decisive Feature Pattern Mining Approach, ACM/Springer Multimedia Systems Journal, Vol. 11, No. 4, April 2006, pp. 352-366.
- (4) R. Aygun and A. Zhang, SynchRuler: A Rule-based Flexible Synchronization Model with Model Checking, IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 17, Num. 12, December 2005, pp. 1706-1720.
- (5) D. Jiang, J. Pei and A. Zhang, An Interactive Approach To Mining Gene Expression Data, IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 17, No. 10, October, 2005, pp. 1363-1378.
- (6) Y. Shi, Y. Song and A. Zhang, A Shrinking-Based Clustering Approach for Multi-Dimensional Data, IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 17, No. 10, October, 2005, pp. 1389-1403.
- (7) X. Xu and A. Zhang, Virtual Gene: Using Correlations Between Genes to Select Informative Genes on Microarray Datasets, Special issue of LNCS Transactions on Computational Systems Biology II, Springer, LNBI 3680, pp. 138-152, 2005.
- (8) L. Wang, A. Zhang and M. Ramanathan, BioStar Models of Clinical and Genomic Data for Biomedical Data Warehouse Design, International Journal of Bioinformatics Research and Applications (IJBRA), 2005, Vol. 1, No.1, pp. 63-80.
- (9) Y. Cho, W. Hwang, A. Zhang, and M. Ramanathan, Assessing Hierarchical Modularity in Protein Interaction Networks, in the Proceedings of the IEEE Computer Society Bioinformatics Conference (CSB2006), Stanford University, Stanford, CA, August 14-18, 2006.
- (10) X. Xu and A. Zhang, Selecting Informative Genes from Microarray Dataset by Incorporating Gene Ontology, the 5th IEEE Symposium on Bioinformatics and Bioengineering (BIBE05), Minneapolis, Minnesota, October 19-21, 2005.

- (11) P. Pei and A. Zhang, A Two Step Approach for Clustering Proteins based on Protein Interaction Profiles, the 5th IEEE Symposium on Bioinformatics and Bioengineering (BIBE05), Minneapolis, Minnesota, October 19-21, 2005, pp. 201-209.
- (12) J. Pei, D. Jiang, and A. Zhang, On Mining Cross-Graph Quasi-Cliques, The Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Chicago, 8/21/2005 - 8/24/2005.
- (13) P. Pei and A. Zhang, A Topological Measurement for Weighted Protein Interaction Network, in the Proceedings of the IEEE Computer Society Bioinformatics Conference (CSB2005), Stanford University, Stanford, CA, August 8-12, 2005.

**Outreach Programs and Supported Students:** The PIs have had meetings with researchers and physicians from Roswell Park Cancer Institute (Buffalo) and General Hospital (Buffalo) to discuss the application of the methods developed in this project. The PIs also gave colloquium talks of the research results developed from this project to other places.

A very informative web page at <http://www.cse.buffalo.edu/DBGROUP/Infrastructure/index.htm> has been set up which contains very useful information for graduate students and researchers.

Outreach activities have also been done through various demos to visitors and companies as well as presentations in conferences.

**Supported Students:** Graduate students working on their Masters and Doctoral degrees are able to conduct their experiments using the equipment. Students taking advanced graduate courses and seminars in various related fields will also benefit from this project. The MultiStore infrastructure will enable us to store vast amount of text, genomic, image, video, and audio data accessible over a high speed network. Students taking the multimedia, bioinformatics, and parallel computing courses will be able to conduct comprehensive analysis to evaluate and compare the new techniques in various systems. Recent Ph.D. students who have been benefited from the MultiStore system: Daxin Jiang, Yomg Shi, Chun Tang, Li Zhang, Young-rae Cho, Woo-chang Hwang, Chuan Lin, and Pritam Chanda.

**Future and Strategic Directions:** Through the research activities of this project, fundamental understanding and novel techniques are provided to support the management and visualization of large-scale multi-dimensional data sets. This theoretical foundation will find broad applications in other research areas which handle large-scale multi-dimensional data sets. The algorithms and tools developed in this project will be disseminated to other applications on campus which require handling of large-scale multi-dimensional data sets.

Management, analysis and visualization of large-scale multi-dimensional data sets have formed a new focus for the research of a significant portion of SUNY-Buffalo. Advanced graduate courses in various related fields will benefit from this project. Graduate students working on their Masters and Doctoral degrees in the related areas will be able to conduct their experiments using the proposed equipment.

# MII: Improving the Pipeline in Applied Computer Science

## NSF Proposal Number EIA-0330822

2005-2006, Project Year 3

Texas A&M University – Corpus Christi

PI

John D. Fernandez, Assistant Professor of Computer Science  
john.fernandez@tamucc.edu URL: <http://www.sci.tamucc.edu/~jfernand/>  
Carl Steidley, Professor of Computer Science (retired Jan 2006)

Co-PIs

Dulal Kar Associate Professor of Computer Science  
dulal.kar@tamucc.edu

Mario Garcia, Associate Professor of Computer Science  
mario.garcia@tamucc.edu URL: [www.sci.tamucc.edu/~garciam](http://www.sci.tamucc.edu/~garciam)

Scott King, Assistant Professor of Computer Science  
scott.king@tamucc.edu URL: <http://www.sci.tamucc.edu/~sking/>


Ray Bachnak, Professor Engineering Technology (replaced Steidley)  
ray.bachnak@tamucc.edu URL: <http://www.sci.tamucc.edu/~entc/BachnakPage.html>

### Abstract

The project, titled "Improving the Pipeline in Computer Science", seeks to improve A&M-CC's computer science program's ability to recruit, retain, educate, and graduate computer scientists. As the project's title indicates, this is a pipeline project so the initial activities involve a regional recruiting effort to increase the number of students entering the computer science program. The primary recruiting efforts are directed towards increasing the number of underrepresented students in the computer science program. To build the infrastructure to educate and retain students four labs were proposed: networking lab, visualization lab, computer engineering lab, and an artificial intelligence lab. Once students receive the benefit of hands-on lab work, students are more prepared to continue to work on a Master's degree and potentially a Ph.D.

### Introduction/Background

This project seeks to improve the Texas A&M University - Corpus Christi computer science program's ability to recruit, retain, educate, and graduate computer scientists. The project is composed of several components which all connect together to meet the project's goals (see figure below). Year One of the project focused on recruiting students into the pipeline and also laboratory development planning needed to enhance ongoing and new research efforts. The visualization lab was implemented first. Year two continued the efforts started in year one and included the implementation of the networking lab. Year three continued the efforts of the previous years and implemented the computer engineering lab.



The A&M-CC Education Pipeline

### Goals, Objectives, Targeted Activities

The goals of the project supported in part by the grant are:

NSF-CISE/CNS Computing Research Infrastructure Workshop  
Snowbird, Utah - June 23-25, 2006

1. improve recruitment and retention, particularly among minorities,
2. expand the current research activities, and
3. expand and improve the current computer science curriculum

The major activities of the grant include recruitment and retention and performing research in the areas of networking, visualization, information assurance, human computer interaction and software engineering.

### **Infrastructure and Facilities**

The Department of Computing and Mathematical Sciences (CAMS) has access to the full A&M-CC information technology infrastructure, and controls and administers 12 computer classroom and 4 laboratories. The computer laboratories are located in the Center for Instruction (CI), the same building where the PI maintains his office. The computer laboratories (CI 226 and CI 228) are available seven days a week for all students to use for research and programming assignments. Additional instructional laboratories are located in CI 229 and CI 230, both of which require an instructor present during the open laboratory sessions. Classroom facilities for instruction and demonstrations are located in several buildings on campus.

The equipment and software in the specialized computer classroom laboratories are aimed primarily at instruction in the computer science, geographic information science, and engineering technology programs. The lab in CI 226 has 26 Pentium IV workstations (and one Pentium III) running a combination of Linux and Windows XP professional. The lab in CI 228, added in Fall 2001, was constructed to meet freshmen enrollment demands and houses 30 Pentium IV PCs. The two other classroom laboratories (CI 229 and CI 230), administered by CAMS, contain a total of 53(21/P4 in ci229 + 31/P3 in ci230) Pentium III or Pentium IV based workstations running, dual-boot, Windows XP/Linux. All of the workstations in these laboratories are equipped to support 100 Mbps connections. The real-time laboratory (CI 345), partially equipped and developed with an NSF CCLI grant and University matching funds, is equipped with 6 Linux-based PCs, 2 Pioneer mobile robots, 10 Lego Mindstorms robot systems, and an SGI Onyz 3400 machine that is used as the main platform for computational intensive work.

The visualization laboratory (S&T 209) has 6 Dual-Pentium IV systems and 3 SGI Graphics Workstations. The networking lab (S&T 208) houses 15 Dell Pentium IV computers, 14 Dell Pentium III computers and CISCO networking equipment such as two 2600 series routers, two 2950 switches, one VPN 3000 Concentrator and one IDS-4235 (Intrusion detection system) that came from a CISCO grant. Other labs include: Control System Lab. (S&T 111), Manufacturing Lab (S&T 114), Material Science Lab (S&T 118), Electronics Labs (S&T 220, 217, and 221), Design Lab (S&T 214), GIS Lab (CI 229), and others.

### **Project Description**

The beginning of the pipeline is recruiting. NSF funds are being used to hire A&M-CC computer science undergraduates to recruit students from area high schools and middle schools. Our students go to science and math classrooms in area schools, attend recruiting fairs, host tours at A&M-CC, phone interested prospective students, etc.

Our effort to improve retention of computer science majors has many subcomponents. Because insufficient math skills is one cause for new college students to leave science, NSF and A&M-CC funds are being used to sustain *SCAMP* (summer camp for applied math program). *SCAMP* is a one-week hands-on program for incoming freshman that demonstrates how math is applied to a range of science problems. In an effort to build learning communities among freshman CS majors, A&M-CC begin working with the First Year Learning Communities to offer (Fall 2004) a *freshman seminar* section exclusively targeted at CS majors. Because financial need and pressure to earn money is a leading cause for A&M-CC freshman and sophomores to leave science, and even college, much of the NSF and A&M-CC funds are used to hire undergraduate and graduate researchers.

A major component of this project is the development of four new laboratories for the computer science program, as mentioned above. These laboratories are significantly enhancing our ability to perform advanced research. These labs are also allowing us to expand our undergraduate and graduate curriculum by introducing new courses.

To encourage computer science undergraduates to pursue graduate degrees, this NSF-supported program allows us to expand the number of computer science graduate research assistantships. This is done in initial years through NSF funds, then in later years by other research funds enabled through our enhanced research capabilities.

## Indicators of Success and Major Accomplishments

Thirteen (13) proposals were submitted for research and educational support.  
Twenty four (24) refereed papers were published during this year of the project.  
Twenty five schools (26) and over 2909 high school students were contacted.  
An indicator of success is that Fall 2005 produced a 12.3% increase in freshmen over the previous fall. Spring 2006 produced a 20% increase in freshmen over the previous spring.

### Publications/Patents

- Dulal C. **Kar** and Randy R. Deleon, "On fast estimation of network bandwidth," will be presented in the International Conference on Communications in Computing, Las Vegas, June 2006.
- Dulal C. **Kar** and Randy R. Deleon, "On estimating the downstream bandwidth of an Internet link," In preparation.
- Randy R. Deleon and Dulal C. **Kar**, "Examining TCP/IP timeout and retransmission with ethereal," In preparation.
- Zhuming Lam and Scott A. **King**, "Simulating tree growth based on internal and environmental factors", Proceedings of GRAPHITE 2005, 29 Nov - 2 Dec, Dunedin, New Zealand, 2005.
- Alexis Angelidis, Marie-Paule Cani, Geoff Wyvill and Scott **King**, "Sweepers: Swept Deformation Defined by Gesture", Graphical Models, *to appear*.
- R. **Bachnak**, M. Mendez, J. Esparza, Oliver Fahed, "An Improved Control System for a Remotely Operated Vessel," Proceedings of the Defense & Security Symposium 2006 (DSS), Vol. 6230, April 18-20 2006, Orlando, FL.
- R. **Bachnak**, R. Fox, and R. Chakinarapu, "Teaching Assembly Language with a Taste of Hardware," Journal of Computing Sciences in Colleges, Vol. 21, No. 4, April 2006, pp. 154-160, Huntsville, TX.
- R. **Bachnak**, M. Englert, M. Mendez, C. Ross, W. Wright, "Involving Undergraduate Students in Research Provides them with Ideas for Capstone Projects," Proc. 2006 ASEE/Gulf-Southwest Annual Conference, CD-ROM Proceedings, session T4B2, 13 pages, March 15-17, 2006, Baton Rouge, LA.
- C. **Steidley**, R. **Bachnak**, "Software and Hardware for Web-Based Education," Computers in Education Journal, Vol. XV, No. 4, pp 104-111, October - December 2005.
- Garcia** M. Cisneros R, and Bliss D. "Password Auditing Applications" The Journal of Computing Science in Colleges. 2006 pp. 196 - 201
- Fernandez** J., **Garcia** M., Camacho D., and Evans A. "Software Engineering Industry Experience – The Key to Success" The Journal of Computing Science Colleges, vol. 21, no. 4, 2006, pp. 230 – 236.
- Hytnen R. and **Garcia** M., "An Analysis of Wireless Security." The Journal of Computing Science in Colleges. 2006 pp. 210 - 216
- Phillip Ian Wilson P. I and **Garcia** M. "A Modified Version of the Vigenere Algorithm" IJCSNS International Journal of Computer Science and Network Security, VOL.6 No.2, February 2006
- Mardan M and **Garcia** M. "DDoS Attacks Flooding Networks All Over the World" Accepted for presentation at the The 10th World Multi-Conference on Systemics, Cybernetics and Informatics: WMSCI 2006. Orlando Fl, July 16-19
- Garcia** M, and Sumit A. "What to Teach About Computer Ethics." Accepted for publication and presentation at the 7th International Conference on Information Technology Based Higher Education and Training. Sydney Australia, July 10 – 13.
- Fernandez**, J. & Wilson, P., "Measuring User Satisfaction by Detecting and Measuring Emotion," paper accepted for the 2006 ASEE Annual Conference, June 18-21, 2006.
- Fernandez**, J. & Tedford, P., "Computer Science Recruiting and Retention of Undergraduates to Meet the Needs of the Business Community," paper accepted for the 2006 ASEE Annual Conference, June 18-21, 2006.
- Fernandez**, J., "Computer Science Capstone Courses Supported by Industry-Based Software Tools," paper accepted for the 2006 ASEE Annual Conference, June 18-21, 2006.
- Wilson, P. & **Fernandez**, J., "Facial Feature Detection Using Haar Classifiers," *The Journal of Computing Sciences in Colleges*, vol. 21, no. 4, April 2006, 127-133.
- Fernandez**, J. & Tedford, P., "Evaluating Computer Education Programs Against Real World Needs," *The Journal of Computing Sciences in Colleges*, vol. 21, no. 4, April 2006, 259-265.
- Fernandez**, J., "Human-Computer Interaction Closes the Software Engineering Gap," *Computers in Education Journal*, vol. XV, no. 3, July – September 2005, 96-100.

Li, L., Liu, Y., and **Fernandez**, J., "Content-Rich Web Document Segmentation Based on HTML Tag Structures and Visual Cues", *Proceedings of International Conference on Computing, Communications and Control Technologies*, Austin, Texas, July 24-27, 2005.

**Fernandez**, J., & Young, L., "Comparative Study of Human Computer Interaction Design and Software Engineering," *Proceedings of 2005 ASEE Annual Conference*, Session 3414, June 15-17, Portland, Oregon.

**Fernandez**, J., Fernandez, M.A., and Priesmeyer, R., "Experimenting with an Emotions Measurement Instrument in Usability Testing," *Proceedings of 2005 ASEE Annual Conference*, Session 2658, June 15-17, 2005, Portland Oregon.

### **Outreach Programs and Supported Students**

Students supported:

Graduate Students:

Phillip Wilson, Rita Hodges, Hollie Alexander, Racquel Gilford, Ramya Chakinarapu, Randy DeLeon, Cynthia Rosas

Undergraduate Students:

Delilah Camacho, Reb Spence, Jason Casanova, Jason Shell, Alonzo Balderaz, Oliver Fahad, Jack Esparza, Marc Mendez, Eric Munoz, Stephanie Garza

Outreach:

In general, computer science (CS) recruiting activities included visiting high schools and career fairs to speak with students and encourage them to pursue a degree in CS and to help them understand the advantages of having technical skills and how those skills can propel them to high paying jobs. The CS recruiting team brainstormed and researched different ways to reach underrepresented high school students as well as to provide them information about CS and college life. The brainstorming also included developing ideas on reaching out to undeclared college students. This work resulted in a new presentation with the goal of informing students about our University, college life and the transition from high school, and to help them understand what CS is and inform them about the CS program we offer.

The CS recruiting team, led by Gilford and composed of six undergraduate students, traveled to high school classrooms to present themselves as examples of college students who are not much different from them. This is the most effective recruiting tool, but the team does formally make a presentation during the visit with great emphasis in getting students involved in asking questions and inquiring more about the CS program and TAMUCC. The presentation includes examples of companies that have hired our graduates, internships obtained by students, and estimated salaries that may be earned by CS graduates based on surveys and information from our alumni.

At University sponsored Island Day Fairs, the team has another opportunity to speak with students one on one. Recruiters stand in front of the CS recruiting display developed by the team to attract students and parents who may be walking by during the fair. Many students we visited in high schools come back to see us at Island Days to get more information and see the campus.

The total number of students contacted is 2909 in 26 different school events or classrooms. The total females of 1179 represent a little about 41% of the total students. The total number of Hispanics contacted is 1873, representing about 65% of the total number. One other career event that was not detailed included a walk-through of 251 students for a brief stop by the recruiting table.

	<b>Fall 2005</b>	<b>Spring 2006</b>	<b>Totals</b>	<b>Percent</b>
Total Schools or Other Events	10	16	26	
Total Students Contacted	1080	1829	<b>2909</b>	
Girls	463	716	1179	40.5%
Boys	617	1113	1730	59.5%
Anglos	280	687	967	33.2%
Hispanics	763	1110	1873	64.4%
Blacks	37	28	65	2.2%
Asian	0	4	4	0.1%

The total number of students contacted for the grant period 2003-2006 is 5522. This number includes 2177 females for a percentage of 39.4%. The number of Hispanics contacted through May 2005 is 3778 or 68.4% of the total.

CS recruiters, not only make high school trips to give class presentations, but they also attend career fairs and other events to meet with students. The team made one recruiting trip to a high school football game to visit with students as they came into the stadium. They also give University tours to high school students. We are always looking to expand our recruiting opportunities every semester by extending our reach to high schools outside of our area. We are exploring giving presentations at teen recreation centers during the summer sessions. We are also working on improving our Web site so as to provide more resources for anyone who is interested in A&M CC or our computer science program.

### **Future and Strategic Directions**

#### Networking Education and Research – Kar:

1. Extensive testing on bandwidth estimation algorithm will be performed and the results will be published in a journal.
2. Papers will be published on downstream bandwidth estimation of an asymmetric link, lab experiments on TCP's behavior, and on the security of memory-resident data.
3. New areas of research in security, particularly in image processing such as watermarking and steganography will be sought.
4. Proposals will be written for funding research in information assurance.

#### Visualization Education and Research – King:

1. Finish processing the motion capture data and develop a general model of prosody.
2. Develop a general model that handles coarticulation and prosody, or two tightly coupled models.
3. Complete the experiment on eye gaze and publish paper

#### Computer Systems Engineering Technology Research and Education - Bachnak

1. Building and test a prototype of the Video Conversion and Text Data Overlay system to evaluate its functionality
2. Upgrading the Remotely Operated Vehicle (ROV) by adding a fiberglass cone to the front; incorporating internet-based measurements, control, and data presentation; employing the ROV to map wide coastal areas; and adding autonomous navigation capabilities to the boat

#### Information Assurance – Garcia

Work on Distributed Denial of Service and Wireless Security. Attacks that can compromise computer systems have steadily increased in the past few years. These attacks may range from simply being annoying without being destructive to attacks that can compromise the confidentiality, integrity, and availability of any computer systems in devastating ways. This research will cover attacks designed specifically to compromise the availability of computer systems or resources. Denial of Service and Distributed Denial of Service attacks have not only become more powerful and sophisticated, but they have devastating effects on businesses as well. This research will evaluate several Distributed Denial of Service attacks as well as some of the commonly used tools to launch these attacks. In addition, it is expected to explore various defense tools and techniques to prevent damaging Distributed Denial of Service attacks

#### HCI – Fernandez

This coming year will permit the following activities:

1. Expanding the emotion measurement research to include biofeedback and perform analysis of its utility for usability testing
2. Expanding the facial feature detecting and tracking to work in real-time mode in an attempt to detect user reaction while evaluating a system with the objective of making interactions with computers more intuitive and easier to use.

#### Computer Science Recruiting – Fernandez

Plans for next year include:

1. Contacting students we visited the previous year who have not yet graduated from high school and touch base with them about their future plans and answer any questions they may have.
2. Continuing to reach as many high schools and career fairs as possible and make presentations about college life and computer science at TAMUCC, including making a visit outside the area.
3. Developing a video about CS and the University which may be distributed throughout high schools and community colleges in Texas and other areas.

# **RI: UC Berkeley Wireless Research Infrastructure**

Principal Investigators: Borivoje Nikolić, Robert Brodersen, David Tse

{bora, rb, dtse}@eecs.berkeley.edu

<http://bwrc.eecs.berkeley.edu>

<http://www.eecs.berkeley.edu/wireless>

University of California, Berkeley

Year 2, 9/15/05 to 8/31/05

## **1. Abstract**

A synergistic research in fundamental communications theory and the development of novel communications technologies is covered in this work. We are establishing new and enhancing the existing research infrastructure for research programs associated with the two centers: the Berkeley Wireless Research Center (BWRC) and the Wireless Foundations Center (WFC). The purpose of this new proposed infrastructure is to build a research environment that enables investigation of novel technologies for cognitive radios, high data rate transmission over wireless local area networks (WLANs), and wireless sensor networks.

## **2. Introduction/Background**

The wireless explosion that we are witnessing today can be largely attributed to the availability of unlicensed spectrum bands. New unlicensed bands are being allocated, for example the 5GHz of spectrum available worldwide at 60GHz. At the same time, a large majority of available spectrum is locked in by its pre-allocated use in legacy systems. Methods of overlaying the legacy systems with new wireless systems are being investigated, which allow the use of ultra-wideband (UWB) radios to operate in the 3-10GHz band. Simultaneously, the FCC is considering allowing unlicensed ‘cognitive’ radios to overlay allocated bands, such as TV spectrum overlay.

Our synergistic work takes a broad and thorough inter-disciplinary approach to developing the fundamental understanding of the operation in new bands, such as 60GHz and UWB, spectrum reuse and spectrum recycling by cognitive radios, together with reducing the requirements of these systems to the basic hardware specifications of the underlying technology.

## **3. Goals, Objectives, Targeted Activities**

We are developing a common computational, test and measurement infrastructure that will allow a quantum leap in wireless technology research and its applications. By building a common computational infrastructure, consisting of compute servers, clusters of workstations and FPGA-based emulation, we will foster the propagation of information from theory to prototypes. To fundamentally understand the physical properties of new bands as well as new methods of spectrum utilization and coexistence of various systems, a major investment in test and measurement infrastructure will be made. This common infrastructure supports the research of over a hundred graduate students, tens of undergraduate researchers and more than ten faculty.

## **4. Infrastructure and Facilities**

The infrastructure covered in this proposal will be used by the BWRC and WFC. BWRC was opened in January 1999. The Center is a research unit of the Department of Electrical Engineering and Computer Science and other departments of the University of California, Berkeley, <http://bwrc.eecs.berkeley.edu/> The Center supports more than 60 graduate students and 12 faculty engaged in researching future generations of wireless communications and systems-on-a-chip implementations.

The BWRC is a modern 12,000 square foot state-of-the-art facility containing 41 workstations, a classroom, conference room, library and open forum for informal discussions. The facility also contains a

fully equipped laboratory for testing integrated circuit proof-of-concept prototypes as described below. The Center's infrastructure has extensive computing resources including over 75 desktop and laptop computers and back-end servers (acquired as a part of this research project). The computer network contains a 1.2 TB Network Appliance RAID filer and is connected to the UCB Campus through a DS-3 fiber. A rich set of design tools is available for system, circuit and device modeling and simulation. A complete set of integrated circuit CAD tools are also available to support the design of radio system-on-a-chip ASICs. Chips are fabricated by ST Microelectronics and IBM using CMOS and SiGe BiCMOS processes. Current designs use 90 nm and 130 nm design rules.

The BWRC laboratory occupies an area of 600 square feet and contains 14 test/measurement benches and 1 soldering/rework station. Testing is integrated with the design environment through on-line measurement processing, instrument control, and test reporting using HP VEE.

WFC is a communications research center in the EECS Department at UC Berkeley dedicated to providing the theory and algorithms for tomorrow's wireless systems. The center is comprised of six core faculty members whose interests include information theory, communications, signal processing, and coding theory as well as twenty-five graduate students and two postdocs. WFC is located on the 2nd floor of Cory Hall in The Eugene and Joan C. Wong Center for Communications Research, a newly renovated space that facilitates discussion and interaction between graduate students and faculty.

The EECS Department, U. C. Berkeley, has an extensive computer network, with over 1000 machines that are professionally managed and supported by a computer user support group and an infrastructure support group. The PIs as well as their graduate students have access to high performance workstations and state of the art laptops that are part of this professionally managed network.

## 5. Project Description

Currently the need for new spectrum bands is addressed in several ways, most of them addressed in our research. Most notably, a 5GHz band is opened for unlicensed use worldwide in the 60GHz band. Also, ultra-wideband (UWB) systems are allowed to operate in the 3-10GHz range given that their power spectral density is below the level of noise. Our research projects investigate the use of both 60GHz and UWB systems. Furthermore, we are investigating methods to more efficiently use the spectrum in the new unlicensed bands, and to 'recycle' the spectrum allocated to antiquated systems and applications. The idea of spectrum recycling has gained a wide attention among research and regulatory organizations; the FCC is considering allowing wireless systems to overlay the VHF/UHF bands used for television broadcast, for example.

*Gigabit Wireless Networking.* The increasing demand for high performance wireless broadband access to the Internet, as well as for high-data rate indoor communications, requires innovative approaches, including system and circuit design in the most advanced process technologies.

The research efforts in BWRC are targeting a demonstration of wireless circuit/system architectures that can achieve wirelessly what the best wireline interfaces offer today: multiple Gb/s of bandwidth. The approach involves the investigation of novel communication systems for communication systems in the 60GHz band. To render such systems into a practical form factor requires enormous progress in integrating RF, analog, and DSP into the smallest space at the lowest power. Using high-frequency communications, such as that in the 60-GHz band, allows the use of small form-factor antennas, which could yield adoption in much broader set of applications.

Continued scaling of CMOS technology offers fast transistors with power gain into the mm-wave regime. We have demonstrated working prototype amplifiers and oscillators working at 60 GHz using 130nm bulk CMOS technology, and scaling to 90nm and 65nm promises even better performance, higher gain at lower power consumption and lower noise.

By optimizing the layout of 90nm devices, we have demonstrated a record extrapolated  $f_{\max}$  of 300 GHz (using the Mason's unilateral gain at the most stable measurement points). This is an important achievement since it's the highest reported ratio of  $f_{\max}/f_T$ . But more importantly, we see a significant improvement in the performance metric of the optimized layout in the mm-wave band. In particular, the

minimum achievable noise figure for the device is about 2.5 dB at 60 GHz, low enough to enable high performance mm-wave circuits in CMOS.

*Ultra-Wideband-Based Networks.* FCC has recently allowed the use of 3.1GHz and 10GHz band for UWB communications. The idea is that if the power spectral density of the UWB systems is low enough, below the level of thermal noise, they will not interfere with the existing systems operating in this band. The foci of the joint NSF Grant #ANI-0230963, "Ultra-Wideband Based Next Generation Networking" between the BWRC and WFC were in the precise locationing, network capacity and building the experimental testbed. Our UWB research addressed multiple objectives, including the tracking of the positions of devices within a network of heterogeneous nodes, developing robust and efficient protocols for routing digital communications within such networks and exploring the fundamental capacity limits of such systems. We are developing distributed signal processing algorithms that are network-energy and position aware to take advantage of correlations at the application layer to reduce resource consumption throughout the network hierarchy.

Last year, three students have graduated with Ph.D.'s, finishing up research projects on high speed, power efficient A/D conversion; a low power, baseband, impulse-UWB transceiver front-end; and antenna/circuit co-design for low noise amplifiers (LNA) and pulse transmitters. Work has focused on laboratory testing and verification in addition to writing up and publication. The high speed, power efficient ADC (designed for a low power, sub-sampled 3-5GHz UWB system) achieves very good power efficiency: 600MSample/sec conversion with a >4GHz input bandwidth in 5.3mW with a peak SNDR of 34dB. The low power, baseband, impulse-UWB front-end transceiver integrates a 50-Ohm input matched gain block with 0-42dB of controllable gain, a 1-bit, 1.92GSample/sec ADC, a sub-1PPM trimmable 60MHz third-harmonic oscillator, a delay-locked loop-based sampling clock generator, a pulse transmitter, and subsequent control logic on a single die. Power consumption was measured at 4mW (RX) and 2mW (TX) for 30Mpulse/sec operation which may be duty-cycled down to 0.6mW (RX) and 0.4mW (TX) at 1Mpulse/sec. Antenna circuit co-design was demonstrated for a baseband (100-930MHz) LNA with 13dB of gain, a 4dB noise figure (NF), and -10.2dB IIP3 for 0.4mW and for a baseband H-bridge pulse generator. Additionally a 3-8.6GHz co-designed LNA was implemented for the upper frequency UWB bands with 12dB of gain, a 4-6dB NF, and -11dB IIP3 for 14.4mW of power consumption. All of the projects were fabricated in 130nm CMOS process.

*Cognitive Radios.* Spectrum 'recycling' techniques allow a second set of users to use the same physical spectrum as another system without unduly degrading the performance of the existing system and without requiring any explicit acknowledgment from the existing system that the second set of users even exists. We anticipate that such techniques will be vital when we consider how to improve overall spectral efficiency without having to throw out legacy wireless systems. Spectrum 'reuse' techniques allow multiple wireless systems to coexist over the same physical spectrum by using limited interaction and communication between the systems to help them share the spectrum fairly and efficiently.

One of the most basic approaches to spectrum sharing and the pillar of the cognitive radio paradigm is referred to as "listen then talk." The motivation is that most of the spectrum is empty in most places for most of the time, even when one considers supposedly allocated bands. In order to exploit this whitespace, cognitive radio systems (secondary users) need to be able to determine that a particular band is currently not being used by its primary user. Spectrum sensing is the foundation of the listen and talk protocol. Spectrum sensing presents an enormous challenge because the cognitive radio that is searching the spectrum must detect all radios that it might potentially interfere with. This requirement could mean that the cognitive radio must be at least as sensitive as all the in-band incumbent users. The challenge here is to find a low cost, energy efficient radio architecture that meets these requirements. Progress toward that goal includes implementation of a spectrum sensing algorithm for energy detection using 1024 FFT and programmable number of spectral averages on BEE2 and 2.4 GHz radio, testing its performance in AWGN and wireless channels and compared results with theoretical and simulation predictions. An extensive set of wireless performance experiments were made inside BWRC to characterize cooperation gains of multiple radio spectrum sensing. BEE2 was used for data collection and processing. Future work will focus on analysis and simulation of cyclostationary feature detectors for

spectrum sensing in the presence of analog impairments, implementation of cyclostationary feature detector on BEE2 and performance measurements of cyclostationary feature detectors in AWGN and wireless channels using 2.4 GHz radio board.

Concurrently, in WFC, large scale simulations were performed that verified our analytic estimates of the tradeoff between coordination, cooperation, and sensitivity. This tradeoff occurs because of uncertainty concerning the interference produced by secondary users outside of a given cooperating network. Nearby potentially active users risk confusing us into thinking that the primary user is active. Uncertainty in the secondary transmission level causes a fundamental limit on how weak of a primary signal we can detect. Our approach to dealing with this uncertainty is to follow the wideband example and simply enforce a reduction in the size of this uncertainty by means of a protocol. Secondary users nearby must cooperate by not transmitting while we are sensing. But to be practical, the sense of "nearby" must not be too large or coordination is going to be impossible. This is also fundamentally a fairness issue since otherwise, some secondary networks could deny others access to an empty band. At the same time, it is unreasonable to request nodes to give up too much access to the band simply to enable the sensing of others.

We evaluated such "sensing-MAC" schemes and found that the simplest possible detectors result in impractically large radii of cooperation, even for low powered secondary systems. The detectors are too simple-minded and are thus very non-robust to interference. Thus, the primary signal needs to be modeled more precisely and some form of feature-based detectors are required to get processing gain and to be able to estimate and thereby suppress the uncertainty due to potential interference coming from moderately far away unknown secondary users.

*Configurable Radios.* A practical implementation of wireless systems that efficiently recycle and reuse the spectrum relies on a successful implementation of a radio that implements it. The cognitive radio can be viewed as an ultimate software radio: the complexity of required processing requires its digital implementation. Simultaneously, implementation of the software radio comes with its known challenges: the need for a wideband front-end and a challenging analog-to-digital conversion. We are currently investigating the implementation of the main building blocks of multi-standard wide-band configurable radios.

*A Reconfigurable High-End Computing Platform.* The goal of this effort is to design and construct a scalable high-end reconfigurable computer (HERC) system, based solely on field programmable gate arrays (FPGAs) as the processing elements, and to demonstrate the efficiency of the machine on a set of applications ranging from SOC emulation to high-performance computing.

We have completed the design and fabrication of the Berkeley Emulation Engine (BEE2) FPGA-based compute module. We have also completed detailed design testing and verification of the module, and used these tests to develop a set of manufacture tests for future module production. We also completed the design and fabrication of an ADC board to allow analog signal input to the BEE2.

With a stable hardware platform, our attention then turned to the development of the "gateway" and software needed to make BEE2 a useful system:

- Development of several important variations of FPGA-based DRAM memory controllers and communication MAC-level units for FPGA/FPGA and module/module communication.
- Implementation of the BWRC Simulink/Insecta Emulation and ASIC Design environment on BEE2.
- Porting of the linux micro-kernel to BEE2 and integration of BEE2 into the BWRC networked computing environment.
- Development of the BORPH native-FPGA operating on BEE2.

*New Research Infrastructure.* The Berkeley Wireless Research Center (BWRC) has installed the commercially available Sun Microsystems N1 Grid Engine 6 software. The grid is a collection of computing resources that perform tasks and appears to users as a large system that provides a single point of access to powerful distributed resources.

The grid is allowing our researchers to seamlessly and easily access the group's diverse computational resources, which is efficiently accelerating research results. For example, one of our researchers has found a linear increase in performance by being able to run many MATLAB jobs simultaneously.

The BWRC grid currently consists of forty-one dual-processor Sun Java Workstation W2100z hosts, three four-processor Sun Fire V440 servers, one four-processor Sun Fire V40z server, one dual-processor

Sun Fire V20z server, one eight-processor Dell PowerEdge 8450 server, and two four-processor Dell PowerEdge 2650 servers. Simultaneously, three Sun Fire V40Z servers were deployed in WFC. We have purchased 2 pieces of equipment to establish a high frequency measurement capabilities (40GHz to 110 GHz):

Anritsu 3742A-EW - 65-110 GHz Millimeter Wave Module

Anritsu 3738A Broadband Test Set

These Anritsu instruments uniquely coordinate with our 65GHz Anritsu Vector Network Analyzer, Cascade Probe Station, and 65GHz test setup. The flexible architecture provides accurate S-parameter measurements of components and devices to 110GHz. The broadband test set drives the external mmW module that enables frequency coverage up to 110GHz. The test set also routes the IF outputs from the mmW modules back to the VNA.

## **6. Indicators of Success and Major Accomplishments**

Demonstrated the complete receive chain operating at 60GHz band in 130nm CMOS.

Demonstrated larger gains in 60GHz band using 90nm CMOS.

Demonstrated a very low voltage ADC for UWB.

Demonstrated a fully functional pulse-based UWB transceiver.

Developed a 'listen then talk' strategy for cognitive radios.

Built a prototype of a next-generation FPGA-based reconfigurable computing platform.

## **7. Publications**

Please see:

<http://bwrc.eecs.berkeley.edu/Publications>

<http://www.eecs.berkeley.edu/wireless>

## **8. Outreach Programs and Supported Students**

No students are supported under this program (this is a research infrastructure grant).

To maximize the impact, in addition to the traditional means of publications, the research results are being disseminated through participation in communications standardization processes, participation in government and NSF-sponsored studies on wireless technology and policies, and industry involvement. Research results will be used to form new graduate and undergraduate courses that will be taught in Berkeley and elsewhere.

The new research infrastructure will allow us to promote diversity of both graduate and undergraduate research. The PI's have been actively participating in many initiatives for promoting the diversity in the past. More precisely, we have participated in UC Berkeley's SUPERB. SUPERB-Information Technology, the EECS graduate diversity summer undergraduate research program, continues to prepare future underrepresented graduate students for graduate school.

## **9. Future and Strategic Directions**

Our gigabit networking research effort is now focused on measurement and modeling issues above 65 GHz. De-embedding of test structures above 40 GHz on the Si substrate is particularly challenging and an active research topic. Measurement and modeling of Si structures up to 110 GHz will allow a new generation of wireless communication and imaging devices to be pursued.

We are developing a reconfigurable test bed for cognitive radio exploration, along with the continuing work on the theoretical work in the same field.

Starting the design of next generation of high-end computing platform based on Virtex 5 FPGA's, targeting high-end computing applicatios such as emulation of mixed SoCs, massively parallel processors, and cognitive radio systems.

# **CRI: RAMP: Research Accelerator for Multiple Processors - A Community Vision for a Shared Experimental Parallel HW/SW Platform**

Award Number: CNS-0551739, Project Year 1  
University of California, Berkeley  
PI: John Wawrzynek ([johnw@eecs.berkeley.edu](mailto:johnw@eecs.berkeley.edu))  
Co-PI: David Patterson ([patterson@eecs.berkeley.edu](mailto:patterson@eecs.berkeley.edu))  
<http://ramp.eecs.berkeley.edu>

## 1. Abstract

Processor architectures have crossed a critical threshold. Manufacturers have given up attempting to extract ever more performance from a single core and instead have turned to multi-core designs. While straightforward approaches to the architecture of multi-core processors are sufficient for small designs (2–4 cores), little is really known how to build, program, or manage systems of 64 to 1024 processors. Unfortunately, the computer architecture community lacks the basic infrastructure tools required to carry out this research. We propose to build and support a large-scale, FPGA-based emulation platform to accelerate research in multiprocessors. Advances in FPGA capacity and speed now enable a 1024-processor system to be emulated with just a few dozen FPGAs at very low cost per processor. By creating a research community around a common hardware and software infrastructure, we will pool our resources to build a far more productive environment than we could have achieved individually. Hence the name of the system, Research Accelerator for Multiple Processors (RAMP), as it should ramp up the rate of research in hardware and software for multiple processors. First, we intend to create a set of RTL and software design standards that facilitate easy adaptability and integration of hardware and software components in the multi-core architecture. Second, we intend to use these design standards to create a baseline architecture of 1024+ nodes. Third, we will investigate architectures for fast emulation of large scale multiprocessors. For example, what type of memory controller and caches will speedup emulation of 1024 processors. Fourth, we will create systems to observe the MPP behavior without disturbing the computation. This design will be created, distributed for free on the Internet, and supported through full-time staff.

## 2. Introduction/Background

Most computer architects believe the future of the microprocessor is a large number, perhaps hundreds or even thousands, of processors on a chip. Given such widespread agreement, it is surprising how much research remains to be done in computer architecture, networks, operating systems, file systems compilers, programming languages, applications, and so on to realize this vision. The members of a panel on the topic at ISCA in June 2005 articulated how far we have to go. As Wen Mei Hwu of Illinois pointed out, one of the challenges of making progress is that software people generally become engaged only after the hardware is available. His group used VLIW/EPIC simulators for years, but made far more progress in the first six months after the Itanium hardware arrived. From those results, they were able to quickly propose architectural improvements that were incorporated into next generation Itaniums, which then took years to realize. The “edit-compile-run-debug” loop that takes minutes in software takes years when custom hardware is involved and is the real bottleneck in advancing that hardware and the software around it.

The good news for computer designers is that FPGAs have made it plausible to create sophisticated prototypes of new architectures that can be fully functional, running a full software stack. Rapid iteration is the foundation of FPGAs; you can “tape out” every day. Moreover, FPGAs are large and still growing at Moore’s Law rates. Today’s FPGAs can hold one or two dozen “soft-core” RISC processors and we

expect this to double roughly every 18 months. For example, Altera reported that their FPGA could hold 100 Nios-II 32-bit RISC soft cores. Although a commercial architecture, like PowerPC or SPARC, would take much more space than Nios, we still believe we can place at least one or two dozen on an FPGA.

Hence, we could plausibly build a 1024-processor system from just 40 to 80 FPGAs. Although these processors would be factors of 10 to 20 times slower than custom hardware, 100 to 200 MHz processors are still fast enough to run real applications and operating systems. In addition to their flexibility, FPGA-based systems enable a level of detailed probing and monitoring that has not been available since the days of breadboarding with SSI logic. Depending on the number of boards desired, the relatively low cost—about as expensive as an SMP server—makes such systems affordable to many who would otherwise be excluded from this important research area.

Such a system would not just invigorate multiprocessors research in the architecture community. Since processors cores can run at 100 to 200 MHz, a large scale multiprocessor would be fast enough to run operating systems and large programs at speeds sufficient to support software research. Moreover, there is a new generation of FPGAs every 18 months that is roughly twice as fast and with capacity for twice as many cores, so future multiboard FPGA systems are even more attractive. Hence, we believe such a system would accelerate research across all the fields that touch multiple processors: operating systems, compilers, debuggers, programming languages, scientific libraries, and so on. Thus the acronym RAMP, for Research Accelerator for Multiple Processors.

The RAMP project is a community endeavor. We have formed an initial community of developers and users of the RAMP system. They are: Arvind (MIT), Krste Asanović (MIT), Derek Chiou (UT Austin), James C. Hoe (CMU), Christoforos Kozyrakis (Stanford), Shih-Lien Lu (Intel), Mark Oskin (U Washington), David Patterson (UC Berkeley), and Jan Rabaey (UC Berkeley). The ideas expressed in this report are the result of regular meetings among the group.

We believe the impact of this work is nothing less than the transformation of the parallel computing community in computer science from a simulation-driven to a prototype-driven discipline. The impacts of this will be rapid iteration across interfaces of the many fields of multiple processors, thereby moving much more quickly to a parallel foundation for large-scale computer systems research in the 21st century.

### 3. Goals, Objectives, Targeted Activities

The vision is a scalable, multiboard FPGA-based system that would allow construction of up to a 1024-CPU multiprocessor. This size is an interesting target because many problems that are invisible at 32 processors and awkward at 128 become glaring at 1024. This scale challenge is true across the architecture, network, operating system, and applications disciplines. This shared artifact would consist of hardware and a collection of “gateway” (RTL) and software that members of the community would help create. We would seed the cooperative using the IP from the “open source gateway” found at [www.opencores.org](http://www.opencores.org) and open source software such as Linux, gcc, ORC, and Jikes.

For example, some users would want a system that faithfully emulated each clock cycle of a prototype machine. Clock-cycle accurate performance measurements at high-speed would help software development and benchmark sensitivity studies, as we vary parameters such as clock rate, local memory latency and bandwidth, switch latency and bandwidth, and so on. Creating a system that could have reproducible behavior, where every interrupt happened at exactly the same clock cycle on every run, would be useful for debugging hardware and software. It would be wonderful to have this done well once and then shared across the architecture and related communities.

We imagine that this infrastructure could be valuable to many projects, but here are a few representative examples:

- Testing the robustness of multiprocessor HW/SW under fault insertion workloads
- Developing thread scheduling and data allocation/migration techniques for multiprocessors
- Developing and evaluating instruction set architectures for large scale multiprocessors
- Creating an environment to emulate a geographically distributed computer, with realistic delays, packet loss, and so on

- Evaluating the impact of 128-bit floating point on convergence of parallel programs
- Developing and field testing of hardware and software schemes for improved security.
- Recording traces of complex programs running on a large scale multiprocessor
- Evaluating the design of multiprocessor switches (serial point-to-point, distributed torus, fat trees)
- Developing data flow architectures for conventional programming languages
- Developing parallel file systems
- Testing dedicated enhancements to standard processors (a la Tensilica)
- Compiling DSP software directly into FPGAs

The goal is that our initial distribution of gateware and software would contain a complete configuration of a scalable multiprocessor populated with standard processor cores, switches, operating systems, and compilers. A user should be able to load binaries for the default cores and run them without modification. Hence, the default cores would implement a commercially available ISA that already has an open source implementation available, such as IBM Power or Sun SPARC. Our goal would be that once the base system is assembled and software installed, users could easily run complex system benchmarks and then modify this working system as desired, or start from the ground up using the basic components to build a new system. Users would release back to the community any enhancements and new hardware/software modules. A similar model has led to the proliferation of the SimpleScalar framework which now covers a range of instruction sets and processor designs.

RAMP has the potential of offering scalable multiprocessors that are attractive to both architecture and software researchers. It would provide architects with an economical, easily modified, large-scale platform with tremendous tracing facilities, and software researchers with an economical, large-scale multiprocessor with extensive measurement and debugging support. It would provide software developers with a fast and malleable system to use with real applications, large datasets, and full-scale operating systems.

Our goal is to create an infrastructure to attract experts in architecture, operating systems, compilers, programming languages, and applications at many locations, and creating a community that could share tools and techniques since they share a common platform. This shared infrastructure could lead to greater interaction and research productivity across communities, as well as stretch NSF research funding. We could imagine trying to duplicate interesting results from other projects locally. Imagine the impact on research of FTPing a novel computer overnight and then test driving it.

#### 4. Infrastructure

This section presents two key components of the RAMP infrastructure: the FPGA prototyping substrate and the gateware modular design methodology.

The Berkeley Wireless Research Center (BWRC) at UC Berkeley has developed two successful generations of FPGA-based emulation platforms. The success of our first FPGA platform, BEE, inspired our development of a more powerful and scalable FPGA-based emulation platform. We have developed the next generation emulation platform, BEE2. Each compute module consists of five Xilinx Virtex 2 Pro 70 FPGA chips each directly connected to four DDR2 240-pin DRAM DIMMs, with a maximum capacity of 4GB per FPGA. The four DIMMs are organized into four independent DRAM channels, each running at 200MHz (400DDR) with a 72-bit data interface (for non-ECC 64-bit data width). Therefore, peak aggregate memory bandwidth is 12.8 GBps per FPGA.

Rather than begin the RAMP project by designing yet-another-FPGA-board, we intend to use the BEE2 design for the RAMP 1 system. We will use the BEE2 board to figure out what the wish list of features for the RAMP 2 board will be. We will begin the design of RAMP 2 a year into this project. RAMP 2 will be based on a new design employing the Virtex 5 FPGA architecture (expected to be available beginning in 2006).

A configured RAMP system models a collection of CPUs connected to form a cache-coherent multiprocessor. The emulated machine is called the target, and underlying FPGA hardware (e.g. BEE) is the host. We are developing a design framework to allow flexible, modular, cycle-accurate emulation of

target machines. The framework must support both cycle-accurate emulation of detailed parameterized machine models and rapid functional-only emulations. The framework should hide changes in the underlying RAMP hardware from the module designer as much as possible, to allow groups with different hardware configurations to share designs and to allow RAMP modules to be reused in subsequent hardware revisions. We also require that the framework does not dictate the hardware design language chosen by developers.

The RAMP design framework is based on a few central concepts. A RAMP configuration is a collection of communicating units, where a unit is a relatively large number of gates. Example units include CPUs, coherence engines, DRAM controllers, I/O device interfaces, and network router stages. All communication between units is via messages sent over unidirectional point-to-point inter-unit channels, where each channel is buffered to allow units to execute decoupled from each other. Each unit has a single clock domain. The target clock rate of a unit is the relative rate at which it runs in the target system. In some cases, a unit might use multiple physical clock cycles to emulate one target clock cycle, or even require a varying number of physical clock cycles to emulate one target clock cycle. All channels are unidirectional and strictly point-to-point between two units. The two units at each end of a channel can have different target clock rates and are only synchronized via the point-to-point channels.

The benefit of enforcing a standard channel-based communication strategy between units is that many features can be provided automatically. Users can vary the latency, bandwidth, and buffering on each channel at configuration time. The RAMP configuration tools will also provide the option to have channels run as fast as the underlying physical hardware will allow to support fast functional-only emulation. The configuration tool builds in support to allow inter-unit channels to be tapped and controlled to provide monitoring and debugging facilities. For example, by controlling stall signals from the channels, a unit can be single stepped. Using a separate automatically-inserted debugging network, invisible to target system software, messages can be inserted and read out from the channels entering and leaving any unit. We have encapsulated the concepts of this design framework in a special scripting language, called RAMP Design Language (RDL).

## 5. Project Description

We propose to build and support a large-scale, FPGA-based emulation platform to accelerate research in multiprocessors. Advances in FPGA capacity and speed now enable a 1024-processor system to be emulated with just a few dozen FPGAs at very low cost per processor. The work here will allow us to create a research community around a common hardware and software infrastructure, and allow us to pool our resources to build a far more productive environment than we could have achieved individually. Hence the name of the system, Research Accelerator for Multiple Processors (RAMP), as it should ramp up the rate of research in hardware and software for multiple processors.

We expect RAMP to provide multiple orders of magnitude speedup over software-based simulation. This speedup will provide a qualitative leap in the quality and range of future computer architecture research. By running a full RTL model at high-speed, researchers will have much higher confidence in the accuracy of their results and the feasibility of their ideas.

In fact, the resulting prototype systems will have sufficient speed to interest a much wider research community, including compiler researchers, operating system developers, and distributed system designers. Previous experience has shown that software researchers are only inspired to work with new computer architectures when a hardware prototype is available. Furthermore, RAMP allows rapid changes in the hardware based on feedback from software developers; unlike conventional chip fabrication, we can tape out every day.

Because RAMP is a research platform, we can offer capabilities not found in commercial multiprocessors. For example, RAMP could have reproducible behavior, where every processor's memory references and interrupts would happen at exactly the same clock cycle on every run. RAMP would be valuable for many forms of software experiments, as well as aiding in software debugging.

Rather than having many institutions build their own RAMP-like boards and accompanying software, we intend to develop a common platform aided by industrial donations. The goal would be that any site could obtain a copy of the hardware at modest cost, and download a fully-functioning large-scale multiprocessor including hardware design and software environment. This will lower the barrier to entry for access to a multiprocessor research capability, and will dramatically expand the set of individuals and departments who can participate in this new wave of architecture and software research.

Both parallel software and multiprocessor architecture are critical research areas, and RAMP will provide a platform to support much richer interaction between these communities. We believe RAMP will offer an unprecedented combination of speed, accuracy, repeatability, and adaptability in a large-scale multiprocessor platform with reduced cost and risk to both NSF and individual researchers.

## 6. Indicators of Success and Major Accomplishments

While the project just begun several months ago, we have made significant progress toward our goals. Our activities have centered on 1) building industrial support and collaboration, 2) outreach to developers, and 3) internal development of the RAMP infrastructure and prototypes. In January 2006 (before the start of the NSF funding) we held a three-day workshop at UC Berkeley where we gave RAMP developers hands-on experience with the BEE2 FPGA computing modules and its associated design tools. We also used this as an activity to engage with industrial affiliates. The result of this workshop and other negotiations has resulted in the donation of three industrial processor cores for inclusion in RAMP: the Xilinx Microblaze, OpenSparc from Sun, and PowerPC from IBM (pending).

We have released several versions of the RAMP Description Language (RDL) and it is beginning to gain acceptance within our development community. We are currently developing two RAMP prototypes, one at Stanford and one at UC Berkeley, named "RAMP red" and "RAMP blue", respectively. Both are expected to be completed within several months. The senior RAMP personnel continue to meet via teleconference bi-weekly and will all meet, along with graduate students, and industrial visitors at a RAMP workshop in Boston, June 21-23, 2006. More information is available on the RAMP website: <http://ramp.eecs.berkeley.edu>.

## 7. Publications/Patents

None.

## 8. Outreach Programs and Supported Students

Although the universities involved in the development in RAMP and the initial group of external supporters come from the usual leading research institutions, we believe the low cost and large scale will allow a much wider range of departments to have their own massively parallel processor for use in courses and in research. If the RAMP 2.0 can be sold for \$10,000 per board, including FPGAs and DRAMs, then departments can have a 128 to 256 parallel processor for the cost of a few fully loaded PCs. We believe that parallel processing can rapidly advance when most CS departments can afford to have their own MPP, where you don't have to ask permission from anyone to try crazy ideas and you can run programs as long as you want. Besides making boards and software directly available for purchase, remote users can sign up to run large experiments on an 1024 node or larger system at Berkeley.

In addition to broadening the types of departments to have access to the large scale parallel processing, the RAMP project will also engage underrepresented groups. For example, some women and under-represented minorities wrote letters of support, so these supporters will surely be involved. Also, many grad students on the developer's projects are from underrepresented groups. For example, most of the developers of Infini-T at MIT are female, while the main developers for the FPGA frameworks at Stanford are African-American.

## 9. Future and Strategic Directions

See section 3 above.

# **CRI: EMSTAR: A COMMUNITY RESOURCE FOR HETEROGENEOUS EMBEDDED SENSOR NETWORK DEVELOPMENT**

**CNS-0453809; 2005**

**University of California, Los Angeles**

**PIs: Deborah L. Estrin (destrin@cs.ucla.edu)**

**Richard G. Guy (rguy@cs.ucla.edu)**

**William J. Kaiser (kaiser@ee.ucla.edu)**

**Edward Kohler (kohler@cs.ucla.edu)**

**Mani B. Srivastava (mbs@ee.ucla.edu)**

**Senior Personnel: Lewis Girod (girod@nms.lcs.mit.edu)**

## **1. Abstract**

Recent work in wireless embedded networked systems has followed heterogeneous designs, incorporating a mixture of elements from extremely constrained 8- or 16-bit “Motes” to less resource-constrained 32-bit embedded “Microservers”. Emstar is a software environment for developing and deploying complex applications on networks of 32-bit embedded Microserver platforms, that can also integrate with networks of Motes. Emstar consists of tools that support simulation, emulation, and visualization of live systems, (both real and simulated); services that support networking, sensing, and time synchronization; and libraries that implement message-passing IPC primitives. This project seeks to continue development of Emstar and promote it as a standard software framework for implementation of heterogeneous wireless embedded systems.

## **2. Introduction/Background**

Wireless embedded systems are invigorating CISE research areas from operating systems, distributed embedded computing, architecture, and networking to signal processing, algorithms, and data management, and opening up new broad-impact applications ranging from wide area environmental management to biomedical monitoring. These systems are increasingly focused on a particularly powerful and exciting class of deployment: the *heterogeneous* or *tiered sensor network*. A heterogeneous sensor network contains nodes with different capabilities, such as both tiny, low-power, 8-bit “motes” and higher-powered, 32-bit “microservers”. Motes are inexpensive and require no infrastructure for long-term deployments, but are also extremely constrained in memory, CPU power, and communication. Microservers, in contrast, are more efficient than motes at many computation- and memory-intensive tasks, and more readily interfaced to high-bandwidth peripherals, such as high-rate ADCs and network interfaces; but their higher energy consumption requires power infrastructure, such as solar panels, in long term deployments. A heterogeneous system containing both motes and microservers can combine the advantages of both devices, using motes to achieve the desired spatial sensing density and microservers to achieve the desired processing power. The whole is greater than the sum of its parts.

Wireless embedded sensor systems present the CISE community with an array of intertwined research challenges: real time sensing of complex and diverse phenomena, embedded computing constrained in bandwidth, energy, memory, and storage, controlled mobility, and the autonomous coordination of vast numbers of network nodes. Even despite these challenges, implementing and testing sensor network applications is daunting: many of the constraints that yield low-power, long-lifetime systems also undermine traditional instrumentation methods for understanding program behavior in deployed systems. No research group could address all of these challenges simultaneously. Coordinated community infrastructure can thus act as a tremendous research accelerator, and without shared infrastructure, research in the field will be significantly hampered. The well-supported TinyOS repository has so far filled this need, representing a free and open coordination point for both researchers and industry. Unfortunately, TinyOS is limited to motes, and since it rightly focuses on responding to resource constraints, it doesn’t generalize to the greater functionality available on other types of node—missing out on opportunities for increased robustness, flexibility, and sophistication. These features are critical to maximizing the

potential of deeply embedded systems. The need is pressing: wireless sensor networking research demands a concentrated effort to develop a community resource for heterogeneous sensor systems.

### **3. Goals, Objectives, Targeted Activities**

The primary goal of this project is to broaden our user base and provide Emstar as a community resource. At the beginning of this project we were using Emstar internally for several projects and it was being used on a trial basis at a few other sites. Through this funding, we will improve the usability and utility of Emstar and broaden our user base, both internal and external. We will achieve this by:

- Making Emstar easier to use through selective restructuring and redesign based on lessons learned
- Greatly improving Emstar documentation and reference implementation examples.
- Supporting Emstar on the most common platforms
- Outreach to the user community through conferences and by providing easy startup kits

In addition, we intend to continue to support the core goals of Emstar, including real-code simulation, natural integration with testbeds, and the capability for debugging and diagnosis when used in deployment.

Our strategy in this is to develop applications internally and use the requirements of those driver applications along with user feedback to prioritize future development effort.

### **4. Infrastructure and Facilities**

Our project is primarily a collection of software, and thus the infrastructure and facilities are virtual. We provide a CVS and HTTP server from which we permit free download of all of our software and tools packages. Our HTTP server hosts a documentation web site as well as a wiki that the larger user community may join upon request. We also provide on-line support via several email mailing lists, for which the archives are public and searchable sources of information.

We also maintain some physical infrastructure. We maintain a testbed of 50 Mica2 motes and 13 Stargate PXA255 Linux-based systems, deployed on the ceiling throughout our facility (and paid for through other resources). Using Emstar, these devices can be used to perform wireless experiments and test application software. This testbed has been used for numerous protocol design projects as well as to test new versions of system software prior to deployment.

In addition, we maintain a portable system of 10 packaged Acoustic ENSBox nodes that we have developed in-house with partial support from this grant. These Emstar-based nodes are powerful platforms for distributed acoustic processing, and have been used in several test deployments. Based on these successful tests, UCLA researchers plan two remote deployments later this year.

Although we do not have sufficient resources to make these facilities generally available, we do make them available to groups in the CENS and UCLA community, as well as other groups based on availability and convenience, on a case-by-case basis. In addition, we provide information on our wiki and websites detailing how these systems can be replicated, as well as the Emstar software which is available for free download.

### **5. Project Description**

Emstar provides several common tools needed for development, debugging, and deployment. This includes several simulation and emulation modes capable of simulating and emulating heterogeneous systems. Emstar's simulator implements a "real code" simulation environment in which no changes to the source code or config files are required to move between a real deployment and a simulation, enabling rapid iteration while debugging. It also supports "emulation modes" in which some devices, for example radio transceivers, are implemented by real hardware

deployed in an environment, rather than modeled. By using real radio transceivers, real channel characteristics are observed, as well as the real characteristics of the MAC layer, including bugs and undocumented features. A visualization tool, EmView, can be used to visualize or trace a running simulation, emulation, or active deployment.

## 6. Indicators of Success and Major Accomplishments

During the covered period, we have numerous major accomplishments.

**Packaging and Documentation.** We have continued to improve the Emstar documentation website and wiki. As part of our preparation for an Emstar tutorial this year at SECON, we developed a packaged version of the Emstar development environment in the form of a bootable “Live CD”. This Live CD allows any user with an Intel-based PC to boot into a full-fledged Linux development environment that is pre-configured to support Emstar development. User feedback from a variety of sources has indicated that one of the primary stumbling blocks in efforts to try out Emstar has been the installation process, as well as the dependency on Linux and on a custom kernel module. This Live CD eliminates these issues by enabling a user to try out the complete system without any modification to their existing system, even if they are currently using Windows or some other operating system.

**Classroom and Educational Use.** Emstar has also been used in a classroom setting for this year’s CS 113 undergraduate class in embedded networks. Acoustic ENSBox hardware and its Emstar-based software stack was used by two groups of students in a graduate class in embedded networked sensing at USC. As part of Network Security class at the Ohio State, two groups of students are doing class projects to add secrecy and authentication to the network services, viz. flood and ping, available in the EmStar package. The groups consist of both undergraduate and graduate students in CSE. Emstar was also used and modified by summer interns as part of the CENS undergraduate summer internship program. This internship program emphasizes the inclusion of women and minorities, leveraging additional NFS funding specific to that purpose.

**Key Functionality and Platform Support.** Important new functionality has been added to the Emstar system, including a highly accurate automated acoustic localization and calibration system. Several important management functions have been added, including the Remote Broadcast SHell (RBSH) and the Delay Tolerant Shell (DTS) facilities for remote simultaneous administration of collections of intermittently connected systems. Support for a number of platforms has been improved. The Acoustic ENSBox platform is completely supported and includes a base software stack as part of its standard image. Support for the MSP430/Chipcon2420 based Telos mote has been implemented, enabling new high rate sensing applications with Mote class hardware. On the standard Intel x86 platform, the build process is now auto-configuring and supports GCC 4.0 as well as all recent distributions of Linux. We have produced two major releases of the Emstar software build and support TinyOS 1.1.13.

**Research Projects Using Emstar.** A number of research efforts in the CENS program have used Emstar in the development and deployment of embedded sensing systems. These research programs include:

- e. A programmable and robust distributed observation system that is supporting investigation of:
  - ecological phenomena such as cold air drainage, and energy balance, being performed at the James Reserve by UC Riverside and UCLA researchers
  - contaminant transport in ground water performed by researchers at UC Merced and UCLA.
  - arsenic contamination of wells in rural Bangladesh using a rapidly deployable ground water system.
- f. An Emstar-based Acoustic ENSBox, that is being used to study the behavior of:
  - Acorn woodpeckers at the James Reserve
  - Marmots in Colorado and Dusky Antbirds in Chiappas, Mexico.
  - Wolves in San Diego
- g. The Emstar system has also being integrated into an experimental portion of the Meso-American Subduction Experiment (MASE) networked seismic sensor facility.

## 7. Publications/Patents

Nithya Ramanathan, Laura Balzano, Deborah Estrin, Mark Hansen, Thomas Harmon, Jenny Jay, William Kaiser, and Gaurav Sukhatme, **Designing Wireless Sensor Networks as a Shared Resource for Sustainable Development**, *First International Conference on Information and Communication Technologies and Development*, 2006.

N. Ramanathan, K. Chang, R. Kapur, L. Girod, E. Kohler, and D. Estrin, **Sympathy for the Sensor Network Debugger**, *Third ACM Conference on Embedded Networked Sensor System (SenSys)*, November 2-4 2005.

Martin Lukac, Lewis Girod, and Deborah Estrin, **Delay Tolerant Shell**, *In ACM SIGCOMM workshop on Challenged Networks (CHANTS 2006)*.

Lewis Girod, **A Self-Calibrating System of Distributed Acoustic Arrays**, Ph.D. Thesis, UCLA, December, 2005.

Vladimir Trifa, **Bird Call Recognition based on Acoustic Data collected using Sensor Networks**. Masters Thesis, UCLA, March 2006.

Mani Srivastava, Mark Hansen, Jeff Burke, Andrew Parker, Sasank Reddy, Ganeriwal Saurabh, Mark Allman, Vern Paxson, Deborah Estrin, **Wireless Urban Sensing Systems**, in *CENS Technical Report #65*, April 2006.

Richard Guy, Ben Greenstein, John Hicks, Rahul Kapur, Nithya Ramanathan, Tom Schoellhammer, Thanos Stathopoulos, Karen Weeks, Kevin Chang, Lewis Girod, Deborah Estrin, **Experiences with the Extensible Sensing System ESS**, in *CENS Technical Report #61*, March 29 2006.

N. Ramanathan, L. Balzano, M. Burt, D. Estrin, T. Harmon, C. Harvey, J. Jay, E. Kohler, S. Rothenberg, and M. Srivastava, **Monitoring a Toxin in a Rural Rice Field with a Wireless Sensor Network** in *CENS Technical Report #62*, 2006.

Thanos Stathopoulos, Lewis Girod, John Heidemann, and Deborah Estrin, **Mote Herding for Tiered Wireless Sensor Networks** in *CENS Technical Report #58*, December 7 2005.

Ben Greenstein, Alex Pesterev, Christopher Mar, Eddie Kohler, Jack Judy, Shahin Farshchi, and Deborah Estrin, **Collecting High-Rate Data Over Low-Rate Sensor Network Radios** in *CENS Technical Report #55*, 2006.

Lewis Girod, Jeremy Elson, Thanos Stathopoulos, Nithya Ramanathan, Martin Lukac, Andrew Parker, and Deborah Estrin, **EmStar: a Software Environment for Developing and Deploying Heterogeneous Sensor-Actuator Networks**, *In submission to ACM Transactions on Sensor Networks*.

Lewis Girod, Martin Lukac, Vladimir Trifa, and Deborah Estrin, **The Design and Implementation of a Self-Calibrating Distributed Acoustic Sensing Platform**, *In submission to ACM Conference on Embedded Networked Sensor Systems (SenSys 2006)*.

## 8. Outreach Programs and Supported Students

During the covered period, Emstar has also been used in a classroom setting for this year's CS 113 undergraduate class in embedded networks. Acoustic ENSBox hardware and its Emstar-based software stack was used by two groups of students in a graduate class in embedded networked sensing at USC. Emstar was also used and modified by summer interns as part of the CENS summer internship program. This internship program emphasizes the inclusion of women and minorities, leveraging additional NFS funding specific to that purpose.

A number of independent research efforts in the CENS program are actively using Emstar-based systems: ecological studies of above and below ground microclimate, environmental engineering studies of ground water in both long term and rapid-deployments, acoustic observations of various bird and mammal species, and a wide area seismic array (both long-term deployed and shorter-term smaller scale deployments)

During this period we directly supported three PhD students: Lewis Girod, Martin Lukac, and Thanos Stathopoulos.

## **9. Future and Strategic Directions**

**New postdoc.** To support continued effort on this project, we have arranged to hire Vinayak Naik, a new postdoctoral staff member who will start in his post in August 2006. Vinayak comes into the position with considerable Emstar experience, having used it in his doctoral work at the Ohio State University. In his previous work, he implemented an adaptive routing protocol in Emstar, in addition to designing Kansei, a network testbed containing 210 XSM Motes and 210 Stargate systems. This testbed uses Emstar as a heterogeneous simulation support tool.

**Short term strategy.** In the short term, we plan to take a new approach to packaging Emstar, by packaging our most popular applications as Emstar Software Development Kits (SDKs). These kits will extract a very well-tested tag of precompiled libraries and utility programs, along with certain portions of the Emstar source tree, and will include documentation and sample code demonstrating how to use them. We anticipate that these kits will make Emstar much easier for applications developers to use, because of the more tractable, small amount of code they would need to understand to get going. We have already tried this approach with the Acoustic ENSBox, which comes pre-loaded with the Emstar stack and requires only limited understanding of that system to develop applications.

**Long term strategy.** In the longer term, we are beginning a lessons-learned and future directions process. We are holding an all-hands design meeting on June 15, 2006 to discuss in more detail a number of ideas for Emstar V2, including revisiting the basic design decisions underlying Emstar V1, in the context of our current understanding of the requirements of applications. These directions may include a movement towards simpler programming models, towards type safe languages, or towards the adoption of existing programming models and the elimination of existing kernel dependencies.

# **CRI: FWGrid: A Research Infrastructure with Fast Wireless, Wired, Compute, and Data Infrastructure for Next Generation Systems and Applications**

Andrew A. Chien, David Kriegman, Joseph Pasquale, Stefan Savage  
Department of Computer Science and Engineering University of California, San Diego  
{achien, kriegman, pasquale, savage}@cs.ucsd.edu  
<http://fwgrid.ucsd.edu/>

Award # NSF EIA-0303622, Year 3 Annual Report, June 2006

## **I. Abstract**

FWGrid enables exploration with real systems and users, radical new applications, novel application structures, system architecture, resource management policies, innovative algorithms, as well as system management. It will fundamentally enhance undergraduate and graduate education and all activities of the department. Because FWGrid will be a “living laboratory”, it enables access to real users and actual workloads. FWGrid models the world of 2010 where widespread high bandwidth wireless, extreme wired (fiber) bandwidth, and plentiful wired computing and data resource are the norm. FWGrid will be used to support a wide variety of research, ranging from low-level network measurement and analysis, to grid middleware and modeling, to application-oriented middleware and new distributed application architecture and finally to higher level applications using rich image and video – both off and on line. Driving applications examples for FWGrid include Smart Vivarium (intelligent behavioral capture for medical mouse experiments), digital diary (capture of everything seen and experienced by a person every day of their life), personal TiVO (the ability to retrieve and integrate this data on demand), photorealistic augmented reality, and omniscient digital observer (capture of everything happening digitally in a particular physical or logical space enabling study of detailed behavior).

## **II. Introduction/Background, Project Description**

In the future digital fabric, computing elements will be distinguished by their networked bandwidth, compute and data capabilities, and their particular input/output capabilities (e.g. graphical displays, cameras, GPS, laser ranging, etc.). These elements will be knit into large scale distributed applications and resource pools – together an entity increasingly known as the “computational grid”. We propose a research infrastructure (FWGrid) which enables research in innovative and radically new applications, systems and system architectures, and the emerging technical and even social use of systems and services built for this future digital fabric.

Key aspects of this infrastructure include:

- highly capable mobile image/video capture and display devices (interact with the world)
- high bandwidth wireless 100-500 Mbps (to couple mobile elements to wired resources),
- rich wired networks of 10-100Gbps (move and aggregate massive data and computation), and
- distributed clusters with large compute (teraflops) and data (10’s of terabytes) capabilities (to power the infostructure).

FWGrid will have a transformative impact on the department, accelerating our growth into experimental computer science, continuing the transform and broaden both undergraduate and graduate education, providing a modern Grid infrastructure on which to perform research experiments, but one that also extends to wireless and peer-to-peer computing/storage. Because of the fortuitous timing with respect to our new CSE building, the infrastructure will have a deep impact on the department’s educational, research, and social environment. It will support experimental research, multi-faculty and multi-disciplinary collaborative research and education. FWGrid will form a nexus for systems research within the department, and couple all of us to researchers on the campus, to SDSC and CalIT2, the metropolitan area, and wide area at high bandwidth. This will enable innumerable shared experiments and collaboration. In short, it will couple us to the future at high speed.

## **III. Goals, Objectives, Targeted Activities**

The major goals in the third year of the project are: 1) to design, plan, and execute the third increment of FWGrid equipment infrastructure, 2) operate this infrastructure in support of a wide range of computational, systems, and networking research, 3) couple the FWGrid infrastructure to the emerging OptIPuter UCSD Campus and national testbeds, and continue to 4) engage in vision and image analysis algorithms research which supports the driving applications, 5) research into remote agent middleware, 6) research for novel high bandwidth and manageable

wireless networking, and 7) to work with key new driving applications, the Smart Vivarium project. We describe each of these efforts and progress below.

**1. Core Computing and Storage Infrastructure** Major goals of the 3rd year of the project is the design, planning, and execution of the 3rd increment of FWGrid equipment infrastructure. Our first and second FWGrid infrastructure consists of a multi-level high bandwidth Ethernet network infrastructure (gigabit aggregated to 10Gigabit) which provides a full 1Gbps bisection for each node. The system is a number of large clusters of dual Opteron & Xeon systems, providing scalable compute and storage capabilities.

Cluster Nodes Detailed cluster configurations now include 96 HP DL145 dual-Opteron 1.6GHz servers with 2GB RAM, two HP DL145 dual-Opteron 2.2GHz servers with 16GB RAM, 64 Dell SC1425 dual-Xeon 2.8GHz servers with 4GB RAM, and 160 Dell SC1425 dual-Xeon 3.2GHz servers with 4GB RAM. Each of the servers have two 250GB hard disks, providing an aggregate cluster storage capacity of nearly 160 terabytes. We dedicate one of the DL145 1.6GHz servers as a computing frontend, an access portal. It monitors and maintains the cluster resources with SGE job scheduler and the OS images for cluster nodes. Another DL145 1.6GHz server is responsible for user administration; managing user credentials as well as hosting all user home directories. The remaining 318 servers make up the computing cluster which provides, aggregately, at least 1.5 Tflops CPU power, ~1TB RAM and ~160TB of disk storage.

Cluster Networking All of the compute servers have two build-in 1Gb network interfaces. The 320-node cluster is divided into 10 rack groups (32 nodes/rack). The network connectivity within a rack is provided by either an Extreme Networks Summit 400 switch or an HP 3400cl (48 Gigabit Ethernet ports and 2 10Gb uplinks). The network connection amongst the first 5 racks is provided by Fujitsu XG800 (12-port 10Gb L2 Ethernet switch). The last 5 racks are connected by HP 5406zl (12 port 10Gb L2 Ethernet switch). This configuration provides a research network with 26Gb bisection bandwidth. The network configuration is shown in Figure 1. All the computing resources (cluster front end, cluster computing nodes and the 64-bit computing servers) have direct network connectivity to the public Internet (1Gb shared bandwidth) and the 10Gb OptIPuter network to several sites, including SDSC, SIO, UIC (Chicago) and Amsterdam, Netherlands.




Figure 1. FWGrid Network Configuration

Wireless Networking The vision of FWGrid posits a high speed wireless front-end network which enables a wide range of mobile and fixed devices to couple tightly with the high performance cluster and networking infrastructure. By using open spectrum and collections of commodity devices (802.11x, WiMax, and wireless USB), FWGrid is exploring the construction of gigabit/second wireless networks in spaces as small as a 300 square foot room. Such bandwidth densities are possible through aggressive use of frequency space, dense deployment of inexpensive wireless access points, and intelligent transport and management algorithms which are the subject of our FWGrid project research.

2. **Large Scale Use of the FWGrid Infrastructure** The FWGrid cluster forms a computational and storage infrastructure which has been used for a wide variety of research, including Bioinformatics, VLSI/CAD (Computer-Aided Design), Large-Scale computing, and Machine Learning. These are in addition to the core areas of Vision/Image processing, Wireless, Middleware, and so on. To date, the FWGrid system has been used by 65 users in 11 research groups, and run over 77,000 jobs.

3. **Coupling of FWGrid to UCSD Campus and National OptIPuter Testbed** The OptIPuter, so named for its use of Optical networking, Internet Protocol, computer storage, processing and visualization technologies, is an IP-based infrastructure that tightly couples computational resources over parallel optical networks using the IP communication mechanism. Through Chien's involvement as the OptIPuter system software architect, the FWGrid cluster has been connected directly to the 100Gbps+ core optical switch and Chiaro router at the core of the UCSD campus testbed. Through this connection, researchers are able to couple FWGrid at 10Gbps to resources across campus, and through CaveWave and National Lambda Rail to resources in Chicago and via Starlight to Amsterdam and many other nations at 10Gbps+. In the near future, we expect these rates to increase to ~50Gbps and soon thereafter to 100Gbps.

4. **Vision And Image Processing Algorithms** which support applications rich in image and video data such as Smart Vivarium, digital diary, personal TiVO, or omniscient digital observer. For our research in computer vision and computer graphics, our research activities have been utilizing the FWGrid resources in three ways: 1) Toward annotating video streams, we have installed a collection of IEEE 1394 cameras, 2) Looking into an office space for the purposes of providing multiple video feeds to the FWGrid. To date our focus has been on real-time tracking and recognizing people in and across multiple video streams, and we have a multi camera, multi processor implementation running in our laboratory, 3) Toward photorealistic augmented reality, we have purchased a Trevesio head mounted display (goggles) which has a pair of 800 by 600 displays to provide stereoscopic depth cues along with a stereo pair of IEEE1394 cameras mounted on the front of the goggles. To date, we have been working on developing the software infrastructure for using this device along with the ARtoolkit. One of our key research directions for using the HMD in photorealistic augmented reality is to estimate the lighting distribution in a real scene using the cameras and in turn to render an artificial object inserted into the scene such that the lighting on the synthetic object matches that of the actual scene.

5. **Middleware Architectures** support execution of novel FWGrid applications. John Ehrhardt and Joseph Pasquale have been working on "Extended Remote Agents: Middleware for FWGrid Applications". This remote execution middleware infrastructure, called Extended Remote Agents (ERA) supports a simple remote execution programming model, and allows for flexibility in the amount of local resources used by participating devices. We take a more structured approach to simplify programming while benefiting from remote execution.

6. **Dense Wireless Spaces** involves the exploration of novel wireless aggregation and management techniques. The effort involves Apurva Sharma, Jason Krasnow, and Andrew Chien and explores the aggregation of multiple wireless channels and commodity technology to achieve gigabit wireless into small spaces in support of distributed sensor data collection such as is needed in the Smart Vivarium project.

To allow application development to keep pace, we have developed the FWComm system. FWComm emulates the technology environment of the future, enabling experimentation with new applications now. Currently, 802.11a/g WiFi technologies can deliver approximately 25Mbps of TCP bandwidth. The FWComm system presented almost triples that bandwidth, delivering close to 70Mbps by aggregating multiple 802.11 channels within a dense wireless cell. Furthermore, FWComm is hardware independent, allowing it to scale with new wireless technologies, always keeping it ahead of the performance curve.

6. Smart Vivarium is an emerging research application uniquely suited as a driver for FWGrid. Modern medical research facilities to house thousands of caged mice, rats, rabbits, and other mammals in rooms known as vivaria. At large-scale research facilities such as UCSD, there are often 100,000 or more mice in ~30,000 cages, supported by a staff of 100+ for basic maintenance and care, and a much larger cadre of researchers. Experiments involve environmental and physiological monitoring to determine the effects of the procedure and the health of the animals involved. Such monitoring is currently performed by human observers, and for practical reasons, only a small subset of cages can be inspected for limited amounts of time. In response to this bottleneck, we have initiated the Smart Vivarium project, a multidisciplinary effort to develop a system for automated, continuous animal behavior monitoring.

The idea behind Smart Vivarium is to use an array of video cameras, advanced image and vision algorithms, and large scale deployment to manage mouse vivaria for scientific research. The initial research efforts behind the Smart Vivarium focused on the computer vision and machine learning aspects of the problem for a single cage. Faced with the challenge of growing the Smart Vivarium to encompass a desired number of 1000s of cages, one is presented with formidable computational, data rate, and data storage challenges -- challenges for which the FWGrid and NSF OptiPuter ([www.optiputer.net](http://www.optiputer.net)) projects are ideally suited. Using FWGrid and OptiPuter infrastructure at UCSD, there is the opportunity to couple a large number of cameras (one per cage) thru a high speed optical network to the FWGrid infrastructure, using the back-end for video storage, online-analysis, and offline analysis.

#### IV. Infrastructure

The FWGrid infrastructure includes a wealth of computing elements, ranging from small video cameras to rich fixed and head-mounted displays to clustered 64-bit fast compute engines to 10's of terabytes of storage (100x current infrastructure), from wireless clients with 100-500Mbps (100x current practice) and wired network infrastructures with 10-100Gbps switched connections (100x current practice). The wired networks' massive bandwidth allow the fixed resources (computing and data bricks) to be dedicated to local activities, clustered or even aggregated in units of any size up to the entire infrastructure (and with resources outside the building). The fast wireless networks will be achieved by using commercial products configured to use multiple channels of the allocated frequency band simultaneously, leveraging commercial products and achieving multiples of the current nominal bandwidth. Such high wireless bandwidth will enable mobile clients to stream/receive multiple high resolution video or data streams, enabling rich interaction with the environment, based on computing and data in the core.

#### V. Indicators of Success and Major Accomplishments

FWGrid supports a wide range of research activities such as Computational and data grids, Vision and Graphics, Distributed computing, and Network measurement and analysis. Listed below are some FWGrid user accomplishments in a few of the wide variety of areas that FWGrid adheres to:

**Bioinformatics** – (Stephen Tanner, Pavel Pevzner, Nitin Gupta) During the last year, we used the FWGrid computational research to annotate over 20 million peptide mass spectra. This work included the search for post-translational modifications (PTMs), particularly in lens proteins. Such modifications are important for protein regulation and disease progression. We constructed a new full-proteome database, and searched it to discover novel exons and improve annotations of human genes.

(Eugene Ie, Yoav Freund) Our work on the FWGrid resource involves indexing a large protein sequence database with patterns represented by position specific substitution matrices (PSSMs). Our initial set of patterns consists of ~28,000 BLOCKS multiple sequence alignments and we process them over 2.6 million sequences from the UniProtKB sequence database (version 6.5). Using FWGrid and our pattern clustering and refinement procedures, we reduce the number of patterns to ~15,000 PSSM patterns. With these refined set of patterns, we are able to achieve 55% coverage of the UniProtKB sequence database.

(Vikas Bansal, Vineet Bafna) We have a computational method to identify inversion polymorphisms using population data for 1 million SNPs in the human genome. FWGrid has tremendously helped in developing the method by allowing us to run our program on the 22 human chromosomes in parallel and improving the method using the results. On a single computer, it would take months to run some of our programs.

**VLSI/CAD** – (Minjian Liu, Yuan Taur) The major achievement in last year is the study of the power/performance trade-offs of 10 nm CMOS. Using FWGrid, our research progress is much faster than before since many tasks can be executed simultaneously. We finished the minimum power supply voltage study from noise margin considerations, gate to source/drain overlap and S/D gradient studies.

**Large Scale Computing** – (Eric Weigle, Andrew Chien) Over the last year we have used the FWGrid hardware as part of a complete regression testing framework for my research code. This has helped to stabilize the code and validate my results. In the coming year I plan to use the resources for large scale network simulations that can not be accomplished in other environments. This resource is critical to the success of our future research. (Pietro Cicotti, Scott Baden) We demonstrated excellent scaling in synthetic benchmark use to study communication overlap in a discrete event simulation.

**Machine Learning** – (Timothy Sohn, Bill Griswold) FWGrid was used to test different machine learning algorithms for detecting mobility using GSM signals. With the FWGrid cluster, I was able to run more than 30 different combinations of features and algorithms in order to find the best results. The results was an algorithm that can successfully distinguish if a person is walking, driving, or staying at one place based solely on that person's phone and observed cell towers.

#### **VI. Outreach Programs and Supported Students**

There is no student funding supported as part of the NSF-support for this project. In prior years, the following students were supported using matching funds from UCSD. Below is a list of the supported students and their accomplishments: 1) Jason Krasnow - Developed SLSP Striped Wireless Research Prototypes, experiments with wireless infrastructure, 2), Apurva Sharma - Tested SLSP Striped Wireless Prototypes, evaluation, construction of demonstrations. exploring new array-based management for high bandwidth, 3) John Bellardo - Developed infrastructure for programmable 802.11a access points, and novel management/configuration algorithms which use measurement to optimize configuration, 4) John Ehrhardt - Designed and developed prototype FWGrid middleware based on Java RMI and an agent architecture, 5) Jongwoo Lim - Develop motion tracker algorithms and software.

#### **VII. Publications/Patents in Current Project Year**

Dekel Tsur, Stephen Tanner, Ebrahim Zandi, V. Bafna, P. A. Pevzner, Identification of Post-translational Modifications via Blind Search of Mass-Spectra, Nature Biotechnology 23, 1562-2567 (01 Dec 2005).

E. Ie, Y. Freund, BioSpike: Efficient search for homologous proteins by indexing patterns, UCSD CSE Technical Report, CS2006-0858, April 26, 2006.

M. Liu, M. Cai, B. Yu and Y. Taur, "Effect of Gate Overlap and Source/Drain Gradient on 10 nm CMOS Performance", submitted to IEEE trans. Electronic devices.

M. Liu, M. Cai and Y. Taur, "The Scaling Limit of Power Supply Voltage from Noise Margin Considerations", SISPAD, 2006

M. Liu, W. Lu, W. Wang and Y. Taur, "Scaling to 10 nm: Bulk, SOI, or Double- Gate MOSFETs?", ICSICT, Oct, 2006

N. Jones, P. Pevzner, Comparative genomics reveals unusually long motifs in mammalian genomes, to appear at ISMB 2006

Pietro Cicotti, Scott Baden, Asynchronous Cellular Microphysiology Simulation (Abstract and talk), 12th SIAM Conference on Parallel Processing for Scientific Computing, San Francisco, CA, February 2006

Pietro Cicotti and Scott B. Baden, Asynchronous programming with Tarragon, The 15th IEEE International Symposium on High Performance Distributed Computing, Paris, June 19-23 2006 (To Appear).

S. Tanner, H. Shu, A. Frank, L. Wang, E. Zandi, M. Mumby, P.A. Pevzner, and V. Bafna. Inspect: Fast and accurate identification of post-translationally modified peptides from tandem mass spectra. Anal. Chem., 77(14):4626-4639, 2005.

#### **VIII. Future and Strategic Directions**

We have made significant additions to FWGrid's core infrastructure that will be brought online in June of 2006. Doubling the compute power and storage by adding 160 Dell SC1425 3.2 Xeon nodes (4GB Mem and 500GB's storage). To connect these additional nodes to the existing high speed 10Gb wired network, we've added a HP5406zl 10Gb switch which will connect the additional racks with a HP3400cl switch via 10Gb copper connections. FWGrid's total capacity is now ~1.5 Tflops CPU power, ~1TB RAM and ~160TB of disk storage, all connected over the fast 10Gb network. Our plan include the further extension of the system with additional cameras, wireless access points, and other novel input devices as part of the overall FWGrid open testbed system. These elements will enable a broader range of interactive, visual, mobile, and simulation experiments.

# CRI: Standard ML Software Infrastructure

Proposal ID 0454136; 2005

University of Chicago and Toyota Technological Institute at Chicago

David MacQueen, John Reppy, and Matthias Blume

[macqueen@cs.uchicago.edu](mailto:macqueen@cs.uchicago.edu), [jhr@cs.uchicago.edu](mailto:jhr@cs.uchicago.edu), [blume@tti-c.org](mailto:blume@tti-c.org)

<http://people.cs.uchicago.edu/~dbm/>

<http://people.cs.uchicago.edu/~jhr/>

<http://ttic.uchicago.edu/~blume/>

## 1. Abstract

Standard ML of New Jersey is a software system consisting of a compiler for the Standard ML programming language and related tools and libraries. It is widely used for implementing other systems in programming, formal methods of software development, applied logic and other areas. It is also a system whose component pieces have been used in several other language implementations and for the construction of program analysis tools. The purpose of the Standard ML Software Infrastructure Project is to develop and enhance SML/NJ so that it will continue to play a valuable role as a programming tool and as an infrastructure for programming language research.

## 2. Introduction/Background

Standard ML of New Jersey is a software system consisting of a compiler for the Standard ML programming language and related tools and libraries. It is used to develop research software and tools by numerous research projects in areas such as programming languages, compilers, software tools, applied logic and theorem proving, formal methods of software development, concurrent programming, and domain specific languages. The development of Standard ML of New Jersey began at Bell Laboratories and Princeton University in 1986, and has been continued by a group of collaborators as an open source project.

Standard ML of New Jersey is the most widely used implementation of Standard ML, and it supports the work of a wide community of researchers. Many of the people actively working on SML/NJ were employed by Bell Laboratories during the period from 1986 through 2002, but all currently active collaborators are now at other institutions, with the core development community now located in Chicago at the University of Chicago and the Toyota Technological Institute. The purpose of this project is to strengthen the resources available for developing SML/NJ and to replace the support lost when SML/NJ related activity ceased at Bell Laboratories.

### Standard ML

Standard ML is a safe, modular, strict, functional, polymorphic programming language with compile-time type checking and type inference, garbage collection, exception handling, immutable data types and updatable references, abstract data types, and parametric modules. It has efficient implementations and a formal definition with a proof of soundness.

The original ML language was an integral part of the Edinburgh LCF theorem prover, where it served as the meta-language (hence "ML") for manipulating logical formulas and programming proof tactics. LCF and the original version of ML were developed by Robin Milner and his research group at Edinburgh University in the 1970s. It was soon realized that ML would serve as a very good general purpose programming language in its own right, and in the late 70s and early 80s free-standing implementations began to appear. In 1983 Robin Milner organized an effort to design a new dialect of ML incorporating ideas from some of the earlier ML implementations and related work from the functional programming language Hope which was also designed and implemented at Edinburgh in the late 70s by Rod Burstall and David MacQueen. This design effort lasted from 1983 to 1989 and produced the Standard ML language design and its formal definition (<http://www.smlnj.org/sml.html>).

Based on additional experience gained from implementing Standard ML (e.g. Standard ML of New Jersey) and from some substantial programming projects using Standard ML, a revised version of the language, known as SML '97, was developed between 1995 and 1997, and a revised formal definition was published (<http://www.smlnj.org/sml97.html>). Work continued on a specification of the basic programming libraries for

Standard ML, resulting in the publication of the Standard ML Basis Library specification (<http://www.standardml.org/Basis>).

As a programming language, Standard ML is notable for a number of reasons. It has a strong type system, with static type checking and automatic type inference, with the added flexibility of parametric polymorphism (sometimes known as “generics”). It has pattern matching over “algebraic” data types, a very powerful and flexible module system, and an elegant exception handling mechanism. Its feature set is rich and powerful enough to make it an expressively efficient general purpose language, yet it remains small and simple enough that programmers typically master and use the full language.

Standard ML has become a paradigmatic language for the programming language research community, and it has been the source of a large literature of research papers and Ph.D. theses (citations?, general SML bibliography? select samples?).

### **Standard ML of New Jersey**

Standard ML of New Jersey (SML/NJ) is a compiler and programming environment for the Standard ML programming language. It was originally developed jointly at Bell Laboratories and Princeton University, starting in 1986, and is now a joint project between researchers at the University of Chicago, Toyota Technological Institute, with significant participation from Yale University, AT&T Research, and individuals at a number of other institutions.

Standard ML of New Jersey is written in Standard ML, except for the runtime system, which is written in C. SML/NJ has been used to implement a significant number of large systems, mainly in the fields of applied logic and verification (HOL 90, Isabelle, Step), and program analysis and advanced compilers. See the more detailed list of projects built using SML/NJ below.

SML/NJ is an open source project currently hosted at SourceForge. The open source license resembles the MIT X Windows and BSD licenses.

Here is a summary of some of the major features of the SML/NJ system:

- The core of the SML/NJ system is an aggressively optimizing compiler that produces native machine code for most commonly used architectures: x86 (IA32), Sparc, MIPS, IBM Power 1 (PowerPC), HPPA, and Alpha.
- SML/NJ runs under Windows 95 and NT as well as many flavors of Unix, including Mac OS X.
- SML/NJ provides an interactive top level based on incremental compilation, but it can produce stand-alone executables.
- SML/NJ uses Matthias Blume's Compilation Manager, to greatly simplify the development of large software projects.
- A variety of general-purpose data structures, algorithms and utilities (such as finite sets and maps, regular expressions, pretty-printing) are provided by the SML/NJ library.
- Concurrent programming in SML is supported by the Concurrent ML library (<http://cm.bell-labs.com/cm/cs/who/jhr/sml/cml/index.html>).
- eXene is a graphical interface toolkit for X-windows based on Concurrent ML.
- SML/NJ extends the SML '97 language with higher-order functors, OR-patterns, first-class continuations, vector expressions and patterns, and a number of other useful features.
- Interoperability with C code is provided by the nlfci subsystem, which can automatically generate stub(?) code from C header files that allows SML programs to access libraries written in C.
- Continuing the tradition of the original LCF/ML dialect, support for manipulating "object languages" (e.g. logics, calculi, simple programming languages, specification languages) is provided by a simple quote/unquote mechanism.

Just as the Standard ML language has been a fruitful source of language research topics, SML/NJ has been a research test-bed for compiler technology for ML and related languages.

### **Components of the SML/NJ System**

Many components of the system are individually valuable as resources to the community even outside the context of their use in the SML/NJ implementation.

Major components of the SML/NJ compiler itself include the Front End, which includes the lexer, the parser, and the type checker for the core and module languages, the FLINT middle end, which uses a typed intermediate language to support a wide range of architecture independent optimizations, MLRISC, a customizable optimizing compiler back-end, and CM, the compilation manager, which provides source code dependency analysis and supports a smart recompilation system.

The compiler is augmented with a suite of tools and libraries. The core library is the Standard ML Basis Library, whose code is currently integrated with the compiler. The SML/NJ Library provides a wide range of additional libraries, including a collection of standard data structures for various forms of collections and maps. Other libraries implement the Concurrent ML language extension and the eXene X-Windows GUI toolkit, and the C-Kit front end for the C language.

### **3. Goals, Objectives, Targeted Activities**

Standard ML and SML/NJ have already enormous contributions to research in programming and related fields. But to remain viable and competitive, a programming system needs to be continually maintained and improved, and the work done under this project is helping to do this. In order to maintain and increase the value of SML/NJ to its user community, we are addressing the following needs: (1) to keep SML/NJ in step with advancing technologies (e.g. support for 64-bit architectures, Unicode character and string types, improved interoperability with other languages, support for new operating system releases and new platforms), (2) to improve the system's performance, reliability, capabilities, and usability, both as a programming tool and as a platform for research and education, and (3) to manage the open source development process (e.g. managing bug reports and fixes, integrating contributed code, testing, building releases, writing and revising documentation, etc.). In addition, the hardware support infrastructure, that we are building, consisting of instances of supported architectures and operating systems, helps maintain the valuable diversity of platforms that SML/NJ is known for.

### **4. Infrastructure and Facilities**

#### **Hardware**

We acquired a new single processor x86 server to support Gforge project software (web site, software distribution, mailing lists, bug management, *etc.*). We will move the SML/NJ project from sourceforge.net to this local server. We also acquired a new compute server for the project software development. This is a quad, dual core AMD machine running Debian Linux. Finally, we acquired an Apple Intel iMac, which was used to perform Intel Macintosh port and support general project development.

#### **Software**

The Gforge software was installed on the first machine mentioned above by David Press. This is expected to become operational as the main project host machine this summer.

Networking services and system administration have been provided by sourceforge.net and University of Chicago Computer Science department.

### **5. Project Description**

The Standard ML Software Infrastructure Project is concerned with enhancing the value of the SML/NJ system as a programming tool and as a platform for continuing programming language research. This effort has a number of major components:

- Compiler modularization. We are working to improve the coarse-grain modularity of the system. This will, of course, improve the maintainability of the system and make it easier to evolve in the face of changing needs, but it will also make it easier to break out major components of the system (such as the front end, FLINT, MLRISC) and use them as components in other systems and tools.
- We are revising or replacing old tools and creating new tools for system building. Examples are the new lexgen and ml-antlr tools for generating lexical analyzers and parsers, and the continuing development of the NLFFI system for interoperability. We are also working on connections with programming environments such as Eclipse.
- We are porting to new platforms and architectures (e.g. Apple's Intel Macintoshes), and maintaining compatibility with new operating system and compiler releases.
- We are working to improve the usability of the system, for instance through improved type error reporting.

## 6. Indicators of Success and Major Accomplishments

Over the past year we issued 5 mini-releases, so-called “working versions” of SML/NJ (110.55-110.59). These working versions are used to make bug fixes and feature enhancements available to our users on a regular, continuous basis. Changes and additions to the implementation can be categorized by the sub-systems they affect as follows:

### Tools:

The area of tool support for program development has seen the most significant developments. In addition to various improvements to existing tools, a new lexer generator (lexgen) has been developed. Lexgen is designed to replace the aging lexer generator ml-lex. Development of a new parser generator (ml-antlr) is in progress. In contrast to ml-yacc, which is an LALR parser generator modeled after the Unix yacc tool, ml-antlr is an LL parser generator modeled after ANTLR. It will, therefore, not replace ml-yacc but complement it. Summary of lexgen:

The new lexgen tool uses "RE derivatives" to directly translate from REs to the DFA implementation. The generated lexers have fewer states than the ones produced by ml-lex. Lexgen was included in the 110.58 release and is currently being used as the default scanner generator for the SML/NJ system. It accepts the old ml-lex syntax, but there is a new syntax in the works that also includes support for Unicode. We also plan a functional stream-based interface to the lexer that will be used by the ml-antlr tool. The lexgen tool was integrated into the compiler and used to build the compiler's lexical analyzer in releases 110.58 and 110.59.

The ml-antlr tool is an alternative to ml-yacc that is based on LL(k) technology (similar to the Antlr tool for Java). In addition to k-token lookahead, it also supports regular-right-part grammars, syntactic and semantic predicates, and backtracking. The generated parsers use recursive descent as their implementation strategy, which is a good fit for languages like SML. We are currently extending this tool with several additional capabilities.

### Platforms

SML/NJ has been ported to three new platforms: Mac OS X on Intel processors (new OS/architecture combination), FreeBSD 3.x series (upgrade of existing port), and Linux on AMD64 systems (upgrade of existing port). In addition, we dropped support for Linux systems using kernels that are older than version 2.2. As a result, it was possible to clean up the signal-handling code in the runtime system (see below). All changes related to platform support were released as part of version 110.58.

### Programming Environment

Work on an SML/NJ plugin for the Eclipse system, an interactive development environment (IDE), has begun. In addition, we have improved some of the scripts that are part of SML/NJ's command-line interface (CLI).

### Summary of Eclipse plugin:

Our approach is to use the existing SML/NJ compiler infrastructure for the analysis of SML code and treat Eclipse as a graphical interface to that infrastructure. We have designed an extensible protocol for our Eclipse plugin to communicate with SML/NJ and have built a prototype of the plugin.

### Libraries

SML/NJ ships with a number of code libraries for use by other SML programmers. Some of these libraries are also used by the implementation of the SML/NJ compiler, its compilation manager CM, and its tools. Extensive fixes and updates to these libraries were incorporated in this year's releases. One of the major libraries, the eXene X-11 GUI library is now maintained by a group at Kansas State University: <http://www.cis.ksu.edu/~stough/eXene/index.html>

### Runtime system

The runtime system provides the interface between the operating system and the part of SML/NJ that is implemented in Standard ML. This includes the generational garbage collector and a large collection of interface stubs for accessing low-level OS functionality. This year support for HP-UX 11 was added, and improvements were made in the support for Cygwin (on Microsoft Windows), x86/Linux, and NetBSD 3.x.

## **Compiler**

In version 110.56 we provided a bug fix for a problem with handling polymorphic equality in the compiler and improved the code generated for equality on datatypes.

## **MLRISC**

MLRISC is the retargetable backend of SML/NJ. It implements instruction selection, register allocation, a number of low-level optimizations, and machine-code generation. Improvements in the MLRISC subsystem related to the x86 register allocator, interaction with the compilation manager CM, and support for Intel Macintoshes.

## **Foreign-function interface**

SML/NJ's new foreign function interface (NLFFI) provides low-overhead calls to C functions and "data-level interoperability", i.e., facilities for directly manipulating C data structures using ML code. We have provided some modifications in order to support new platforms. In particular, beginning with version 110.58 we support the Darwin-specific Intel API. The current solution is not fully general and plan to replace it in the future to eliminate current limitations.

## **Compilation Manager**

In version 110.58 we re-implemented the internal error handling in CM, the compilation manager. As a result, there are fewer "spurious" and "follow-up" errors being reported. The resulting reduction in unnecessary and misleading error diagnostics improves the subjective usability of the system.

## **Installation:**

In 110.58 we provided an improved installation procedure which now handles "pre-loading" of libraries such as `$/html-lib.cm` even if they are not used by SML/NJ itself. (Only libraries used by SML/NJ itself are shipped in pre-compiled form; the old installer could not pre-load libraries that were not shipped that way.)

## **CML**

Concurrent ML is a dialect of Standard ML that provides concurrency primitives. The design is based on an abstraction of first-class synchronous events. The implementation of CML takes the form of a CM-managed library which is shipped (in source-form) with SML/NJ. Although CML is fairly stable and, therefore, undergoes very little change, we provided a number of important bug fixes which were shipped with SML/NJ working versions 110.58 and 110.59.

## **7. Publications, Patents**

The principle product of this project is improved software for programming in Standard ML, and the main publications will consist of documentation and manuals. Plans are also being made for a series of papers describing several of the major system components.

Since SML/NJ is free, open-source software, we do not anticipate filing patents on the new developments.

## **8. Outreach Programs and Supported Students**

During the summer of 2005 the project supported two graduate students. David Press worked on setting up a Gforge server, and Adam Shaw did preliminary work on an Eclipse-based environment for Standard ML. Undergraduate Aaron Turon has been supported for work on `lexgen` and `ml-antlr`. During the summer of 2006, graduate student George Kuan is being supported for work on streamlining the front end.

## **9. Future and Statagic Directions**

A major task for the future is to develop true 64 bit support for 64-bit architectures like x86-64. We also plan to improve support for interoperability with C through further development of NLFFI. Improved support for profiling and debugging is another high priority need. Garbage-collection performance is another area that needs attention, and considerable design work has already been completed on this problem.

The current restructuring work on the Front End is partly motivated by the need to provide a cleaner interface between the Front End and the FLINT Middle End. In the longer term, we need to streamline or possibly replace the FLINT middle end, which has become more complex than necessary.

A long term goal is to use SML/NJ as the basis for experiments in language evolution and experimentation. Many new or modified features of interest can be implemented almost entirely in the Front End, if the interface between the Front End and FLINT is made sufficiently clean.

# CRI: Computing Infrastructure for the UConn Health-Grid Initiatives

CNS 0551549, 2006

University of Connecticut

Chun-Hsi Huang

huang@cse.uconn.edu

<http://www.cse.uconn.edu/~huang>

## 1. Abstract

The UConn *Health-Grid Initiatives* aim at advancing the application of modern information technology to various disciplines of life science research and practice by promoting and reinforcing awareness of the advantages linked to the development and deployment of modern Grid technologies towards an infrastructure for automated information integration and analysis. The computing and information infrastructure is based on a campus-wide computational and data Grid, an ongoing effort initiated in 2004. Also an ongoing effort is the development of general-purpose middleware support for secure transfer of sensitive data over the computing infrastructure. The educational programs associated with the Health-Grid Initiatives include the development of new and re-development of current courses to incorporate inventions coming out of the research, a new cross-disciplinary e-Health minor, as well as new degree tracks emphasizing the application of information technology to different disciplines of life science. In addition, an annual scientific meeting, the international *Bio-Grid* Workshop, will be held in conjunction with the research and education enterprise.

An infrastructure proposal based on the Health-Grid initiatives is under preparation. This planning proposal requests equipment support that will allow the purchase of ten to fifteen computing nodes with storage devices to join the machines in the PI's lab. These machines will be configured to simulate the job activities on a large-scale campus grid. Starting with a few seed projects, the machines will be used to study the computation load distribution, data migration frequency, bandwidth consumption rate, user statistics, ease of use of the web portal, as well as the completeness of user APIs. This study will provide useful information while determining/implementing the following for the infrastructure proposal: (1) the preliminary Health-Grid configuration for computation and storage use by projects throughout the campus; (2) a set of robust APIs in support of various applications; and (3) a modularized data manager for general-purpose middleware support for secure data migration over the campus grid.

## 2. Introduction/Background

Research in life-sciences increasingly relies on globally distributed information and knowledge repositories. The quality and performance of future computing and storage infrastructure in support of such research depends heavily on the ability to exploit these repositories, to integrate these resources with local information processing environments in a flexible and intuitive way, and to support information extraction and analysis in a timely and on-demand manner.

For example, the epidemiological studies and medical imaging produce tremendous amount of data that are usually geographically distributed among hospitals, clinics, research labs, radiology centers, etc. For research, training or clinical purposes, physicians and medical researchers often need to consult and analyze medical data from distributed sites. Thus, an infrastructure supporting on-demand secure retrieval and efficient integration of sensitive clinical and health information will provide significant convenience by, for example, allowing physicians to get secure access to their patients' images and send hybrid requests over distributed data bases; and allowing the radiologists from geographically dispersed hospitals to share standardized mammograms, to compare diagnoses, and to perform sophisticated epidemiological studies across national boundaries, in a real-time manner.

Also, the Swiss-Prot database, release 50.0 as of May 30, 2006, of protein sequences contains 222,289 sequence entries, comprising 81,585,146 amino acids abstracted from 142,438 references. 2,936 sequences have been added since release 49.7, the sequence data of 189 existing entries has been updated and the annotations of 82,730 entries have been revised. This represents an increase of 1%. Swiss-Prot is only one of dozens of publicly accessible protein databases available worldwide, each focusing on different types of proteins, organisms, diseases, tissue, protein-protein interaction, protein expression, or organization of its content. In addition to proteins, there are hundreds of molecular-biology databases storing an ever-increasing amount of information on DNA and RNA sequences, comparative genomics, genetic and physical maps, transgenics, metabolic pathways and cellular regulation, gene expression, and scientific article repositories like MedLine.

Modern Grid technology represents an emerging and expanding instrumentation, computing, information and storage platform that allows geographically distributed resources, which are under distinct control, to be linked together in a transparent fashion. The power of the Grids lays not only in the aggregate computing ability, data storage, and network bandwidth that can readily be brought to bear on a particular problem, but also on its ease of use. After a decade's research effort, Grids are moving out of research laboratories into early-adopter production systems, such as the Computational Grid for certain computation-intensive applications, the Data Grid for distributed and optimized storage of large amounts of accessible data, as well as the Knowledge Grid for intelligent use of the Data Grid for knowledge creation and tools to all users.

The Grid-based infrastructure supporting on-demand secure retrieval and efficient integration of sensitive clinical and health information will offer significant convenience to life-science researchers and health-care providers by, for example, allowing physicians to get secure access to their patients' images and send hybrid requests over distributed data bases; and allowing the radiologists from geographically dispersed hospitals to share standardized mammograms, to compare diagnoses, and to perform sophisticated epidemiological studies across national boundaries, in a real-time manner.

### **3. Goals, Objectives, Targeted Activities**

While promoting and reinforcing awareness of the possibilities and advantages linked to the development, deployment and evaluation of Grid technologies towards the automated integration and analysis of health information in a broad sense, part of the research efforts will be directed towards further exploring research issues arising from the Health-Grid, as a potential solution, via on-campus interdisciplinary collaboration on grid-enabling research projects, as well as the integration tests with two international research teams from Italy and UK.

Architectures proposed in our Health-Grid initiatives will allow to build a general-purpose data manager, taking advantage of classical theory (transactions concept) and proposing solutions to implement a secure and high-performance data manager by decomposing transactions into concurrent tasks and requests. Further equipped with customized web portals and application-specific user interface, this infrastructure may potentially provide a solution to secure and on-demand integration/analysis of life-science data in a broader sense, including information from the population level (social health-care) to the individual (clinical practice) and molecular level (genetic and proteomic information).

The educational programs will produce software engineers who are prepared to formalize and solve problems emerging from life science disciplines, as well as clinical scientists and professionals with strong information processing techniques.

#### 4. Infrastructure and Facilities

Major research facilities for this project are maintained by the Taylor L. Booth Engineering Center for Advanced Technologies (BECAT) at the University of Connecticut. The BECAT maintains a modern set of networked laboratory facilities available to Computer Science and Engineering, including (1) a 64-processor SGI Altix 3700 Bx2 supercomputer, (2) several high performance computing systems and servers including the BECAT GRID which consists of over 24 nodes supporting grid and cluster middleware architectures, and (3) numerous computing workstations which are available for small-scale and prototype research projects ranging in platforms from Solaris to Wintel to Linux.

The Information Technology Engineering Building (ITEB) is the home of the Computer Science and Engineering Department. It houses newly renovated laboratory space for graduate students in close proximity to many faculty offices. Computing equipment for this space includes 40 Pentium II Windows NT systems, a Windows NT server, and 20 Sun Ultra workstations. Graduate students on funded projects have first priority for these facilities, as it affords the serendipitous communication that is so crucial for effective research progress. ITEB houses many of our computer science undergraduate laboratories, including several high-tech classrooms, which have PCs connected to the Internet, Barco projectors for both PC and VCR, satellite down-links, and area cable television.

#### 5. Project Description

The current task contains two phases. The first phase is to complete the configuration of a campus-wide computational and data grid, the Health-Grid. The second phase is, on top of the Health-Grid infrastructure, to develop middleware support for secure and automated integration of sensitive medical data. For global compatibility, development and integration test in the second phase involve two international research teams in Italy and UK. A few projects conducted by faculty members from the Schools of Engineering, Medicine, and Public Health are currently being grid-enabled. One representative project, Cooperative Genomic Knowledge Inference, led by the PI, is described next as an example.

It is crucial that the massive genomic data produced are well represented so that useful biological information may be efficiently extracted or inferred. A useful tool for effective knowledge representation is the semantic network system. A *semantic network* is a conceptual model for knowledge representation, in which the knowledge entities are represented by nodes (or vertices), while the edges (or arcs) are the relations between entities. A semantic network is an effective tool, serving as the backbone knowledge representation system for genomic, clinical and medical data. Usually these knowledge bases are stored at locations geographically distributed. This highlights the importance of an efficient distributed semantic network system enabling distributed knowledge integration and inferences. The semantic network is a key component of the Unified Medical Language System (UMLS) project initiated in 1986 by the U.S. National Library of Medicine (NLM). The goal of the UMLS is to facilitate associative retrieval and integration of biomedical information so researchers and health professionals can use such information from different (readable) sources. Our research team has developed a distributed semantic network system, based on a task-based and message-driven model to exploit both task and data parallelism while processing queries. The system also features multi-threading and dynamic task migration to support communication latency hiding and load balancing, respectively. In the task-based message-driven model, queries are decomposed into tasks and distributed among processors for execution. Other system support activities are also decomposed into system tasks and distributed as well. When a task is completed, a message is generated to either spawn new tasks or trigger further processing, depending on the property and current status of the task. This process is carried out by two collaborating components: the host system and the slave system. The host system interacts with the user and processes the information for the slave system, while the slave system performs task execution. This task-based and message-driven system is particularly suitable for grid environments. The next phase of this project is to test the efficiency of the grid-enabled distributed semantic network system on the Health-Grid. The design experience of the distributed UMLS, particularly the task model for cooperative inference and the layered architecture for

the host and slave systems, complies with the health data management (HDM) unit in our middleware development for the Health-Grid, in which the metadata management capabilities can easily be designed. The HDM design for secure retrieval of sensitive bio/medical/health data was in fact motivated by our research on a distributed UMLS.

The educational development program will produce course materials for a redeveloped undergraduate Introduction to e-Health course. This is an initial step towards implementation of an undergraduate e-Health minor to supplement the UCONN Computer Science and Medical degrees. The course will involve fundamental concepts of health informatics, as well as a reasonably complete introduction to the formal aspects of information processing. The e-Health minor will also include a course on grid computing developed by the PI, an (existing) course introducing medical and health informatics, and a capstone senior design course focusing on portable implementation of parallel algorithms for biomedical and healthcare problems. The minor will produce software engineers who are prepared to formalize and solve emerging medical and health applications, as well as clinical scientists and professionals with strong information processing techniques. The PI has recently served on the curriculum committee of the department of Computer Science and Engineering and has been working on proposing two new graduate courses, Advanced Computational Biology and Medical Informatics, jointly with the department of Neuroscience at the UConn Health Center and the School of Allied Health of UConn. The proposed computing infrastructure will serve as the platform for class projects.

## 6. Indicators of Success and Major Accomplishments

Research programs incorporating cross-disciplinary expertise are among the strategic research directions of the School of Engineering at the University of Connecticut. The Health-Grid Initiatives have been enthusiastically endorsed by the Schools of Medicine and Allied Health and have attracted a large number of faculty participants. Research outputs related to the initiatives have been published at major journals and conferences.

## 7. Publications/Patents

Selected publications related to this project are listed below.

1. V. Lanza and C.-H. Huang, "Advanced Course for Doctor as Department Network Administrator in Anesthesia and Intensive Care Units", to appear in *Journal of Clinical Monitoring and Computing*, 2006.
2. C.-H. Huang, V. Lanza, S. Rajasekaran and W. Dubitzky, "HealthGrid: Bridging Life Science and Information Technology", *Journal of Clinical Monitoring and Computing*, Springer, 19(4-5):259-262, 2005.
3. C.-H. Huang, "BioGrid: A Collaborative Environment for Life Science Research", in *Proceedings of the 20<sup>th</sup> International Symposium on Critical Care Medicine*, 123-132, Nov. 18-21, 2005, Trieste, Italy.
4. C.-H. Huang, "Latency Reduction in Translational Research", to appear in *Proceedings of the 21<sup>st</sup> International Symposium on Critical Care Medicine*, Nov. 10-13, 2006, Venice-Mestre, Italy.
5. S. Rajasekaran, S. Balla and C.-H. Huang, "Exact Algorithms for Planted Motif Problems", *Journal of Computational Biology*, 12(8):1115-1126, 2005.
6. C.-W. Lee and C.-H. Huang, "Towards Cooperative Genomic Knowledge Inference", *Parallel Computing Journal*, 30(9-10):1127-1135, 2004.
7. C.-H. Huang and L. Yin, "Distributed Clinical Knowledge Inference", in *Poster Book of the 8-th International Conference on Research in Computational Molecular Biology (RECOMB)*, 556-557, Mar. 27-31, 2004, San Diego, CA.
8. C.-H. Huang and S. Rajasekaran, "Parallel Pattern Identification in Biological Sequences on Clusters", *IEEE Transactions on NanoBioscience*, 2(1):29-34, 2003. Preliminary version in *Proc. of the 4<sup>th</sup> IEEE International Conference on Cluster Computing*, 127-134, 2002.

9. S. Balla, V. Thapar, T. Luong, T. Faghri, C.-H. Huang, S. Rajasekaran, J. Cambo, J. Shinn, W. Mohler, M. Maciejewski, M. Gryk, B. Piccirillo, S. Schiller and M. Schiller, "Minimotif Miner, A New Tool for Investigating Protein Functions", *Nature Methods*, 3(3):175-177, 2006.
10. S. Rajasekaran, V. Thapar, H. Dave and C.-H. Huang, "Randomized and Parallel Algorithms for Distance Matrix Calculations in Multiple Sequence Alignment", *Journal of Clinical Monitoring and Computing*, Springer, 19(4-5):351-359, 2005.

### **8. Outreach Programs and Supported Students**

The PI has initiated in 2003 the International Workshop on Biomedical Computations on the Grid (BioGrid), currently funded by the NIH. The aims of the BioGrid workshop are to promote and reinforce awareness of the possibilities and advantages linked to the development, deployment and evaluation of Grid technologies in broadly biology-related research and practice. The first event, BioGrid 03-06 have been held in Tokyo (Japan), Chicago (US), Cardiff (UK) and Singapore, respectively, all featuring invited keynote speeches, technical program of peer-reviewed papers as well as plenary sessions. Specifically, the workshop concentrates on all aspects of grid-enabled infrastructures, testbeds, management and security in support of such research areas as

- Computational Genomics
- Computational Proteomics
- Biomedical Information Retrieval
- Biomedical Modeling and Simulation
- Biomedical Image Processing and Simulation
- Distributed Medical Database Management and Integration
- Integration of Biological Information
- Mining and Visualization of Biomedical Data
- Grid-Based Approaches to Systems Biology
- Tele-systems for Diagnostic, Prognostic, and Therapeutic Applications
- Health Data Storage and Retrieval
- Social Health-care
- Pharmaceuticals and Clinical Trials
- Computerized Epidemiology
- Collaborative and Proprietary Health-Grids

In addition to the focus areas, research articles reporting on original results of developing, deploying and evaluating grid techniques in novel topics in bioinformatics, clinical informatics, bioimaging and public health informatics are also solicited.

The PI will continue hosting this workshop as part of the UConn Health-Grid initiatives throughout the project period.

### **9. Future and Strategic Directions**

This project investigates the infrastructure for high-performance, automated integration and analysis of information from the population level (social health care) to the individual (clinical practice) and molecular level (genetic and proteomic information). The project will establish a grid-enabled research network throughout the UConn campus in support of on-campus interdisciplinary life-science related research projects. In the long term, the infrastructure will aim at supporting regional collaborative life-science projects by research institutes and health care providers in Connecticut, as well as becoming a "virtual" life-science data repository and computation center for use nationally and internationally.

# CRI: A Configurable, Application-Aware, High-Performance Platform for Trustworthy Computing

*NSF Award Number: 0551665; Year One*  
University of Illinois, Urbana-Champaign  
Wen-mei Hwu (Principal Investigator)  
William Sanders (Co-Principal Investigator)  
Ravishankar Iyer (Co-Principal Investigator)  
Klara Nahrstedt (Co-Principal Investigator)  
{w-hwu,whs,rkiyer,klara}@uiuc.edu

## 1. Abstract

The goal of this project is to enable ground-breaking experimental research in creating large scale, demonstrably-trusted, cluster computing platforms for what is variously referred to as on-demand/utility computing or adaptive enterprise computing. The fundamental challenge is to provide application-specific level of reliability and security in a manner that preserves standard software interfaces while delivering optimal performance. The intellectual merit lies in investigating new set of application-aware methods to provide customized levels of trust (specified by the application) via an integrated approach involving re-programmable hardware, novel compiler methods to extract security and reliability properties and supported by a configurable OS and middleware.

We plan to transform a large Linux based cluster, using augmented hardware (demonstrated via FPGA implementations), smart compilers capable of extracting and programming into hardware, application-specific reliability and security guarantees and supported by an Operating System and Middleware configured to support the application execution. The proposed infrastructure, which we call the Trusted ILLIAC with massive reconfigurable hardware in each node, enables rapid design and deployment of low-cost application aware hardware engines and supporting OS and middleware configuration that provide an unparalleled opportunity to devise systems and applications with provable security and reliability guarantees. Tools and well-defined procedures will support configurability of system services exposed to end users, including error and intrusion detection, security vulnerability masking, and recovery.

The Trusted ILLIAC infrastructure will support a rich set of research projects that span online hardware-software assessment, efficient programming environment for heterogeneous multiprocessor systems, software bug detection, hardware validation, configurable trust-providing mechanisms, automated fault management, on-line model-based adaptation strategies, middleware support for trust, application-based placement of detectors, and smart card utilization. These research projects will benefit from a common infrastructure and will together advance the knowledge base of building large scale trustworthy systems.

A distinctive, integral part of the Trusted ILLIAC is a validation framework, which constitutes a corner stone for quantitative assessment of alternative designs and solutions. Such evaluation is crucial in making design decisions, which require understanding tradeoffs such as cost (in terms of complexity and overhead) versus efficiency of proposed mechanisms. The framework leverages years of experience we have in experimental evaluation of highly reliable and secure systems and is based on comprehensive fault and attack injection technology and tools developed in Illinois.

In a broader context, the Trusted ILLIAC approach represents a fundamental change in how computing is accomplished, that of direct representation of tasks in silicon, and enables that paradigm by merging the new architecture with existing cluster and operating system functionality. In the field of trust, it will provide customizable computing technology to the broader community of students, researchers, and institutions to enable create their own integrated trusted computing testbeds. The proposed infrastructure will benefit technology transfer efforts from research to the *real world environments* and will enable researchers to collaborate with developers from the government and industry to determine how trustworthy hardware assists and software stacks can be integrated into products.

## 2. Introduction/Background

Historically, security in large-scale systems has been achieved by placing defense mechanisms at the perimeter of these systems. By concentrating defenses on the perimeter via firewalls for an enterprise network or access control for an operating system, the system designer has a clean and simple model for reasoning about the security of the system that corresponds to existing models of physical security. However, the perimeter metaphor breaks down in today's computational environment. Perimeter security does not address insider attacks (either intentional insiders or accidental inside attackers infected by malware). In addition, the definition of "inside" vs. "outside" is not nearly as clear today as it has been in the past. In today's fast-paced business environment, trust relationships are constantly shifting via partnering and outsourcing. Because of these changing business relationships, boundaries are not static and trust is no longer absolute. For example, a manager might trust a contract employee to have access to company confidential information, but she would want additional monitoring of the requested accesses. To provide a secure environment in spite of insecure and changing boundaries, systems must provide multiple layers of trust-aware mechanisms beyond basic perimeter access control.

Augmenting a traditional enterprise cluster with an adaptable, configurable, and integrated HW/SW approach to security and reliability can greatly increase the trustworthiness of enterprise systems with minimal performance impact. This adaptive, configurable, and integrated HW/SW approach to security builds on much of the work performed in reliable systems. When designing reliable systems, the developer attempts to catch errors as close to the source as possible, e.g. before an error leaves the microprocessor or the application process boundaries. Providing efficient low-overhead error containment is a challenging problem in particular in the face of: (i) the technology scaling and power reduction, which contribute to higher error rates and (ii) pervasiveness of network environments and increasing system complexity which increases the chance of error propagation. To address these challenges, both industry and academia have considered employing a hierarchical system of detection and recovery schemes/mechanisms some of which can be embedded into the hardware (e.g., processor or dedicated FPGA-based modules) while others can be integrated with the operating system or application (e.g., via a robust middleware).

## 3. Goals, Objectives, Targeted Activities

The proposed infrastructure will support research in referred to as on-demand computing, utility computing, or adaptive enterprise computing, where a large number of applications from diverse organizations co-exist and share hardware/software resources. These applications can have very diverse trust requirements. A major intellectual challenge is to provide application specific levels of reliability and security while maximizing performance. The current hardware and software approaches to providing trust, which can only be described as "one-size-fits-all," is both insufficient and inefficient. We argue that the answer lies in developing application-aware methods to provide customized levels of trust enforced via an integrated approach involving the compiler, hardware, OS, and middleware. In this proposal, we show how we plan to transform a large cluster using augmented hardware (demonstrated via FPGA implementation), smart compiler capable of extracting application-specific hardware assists, and a configurable OS/middleware to execute the application. The key aspect of our approach is to perform these transformations in a manner that is largely transparent to the applications.

The creation of the Trusted ILLIAC platform envisaged in this document is in itself a major intellectual endeavor. The design of linker, loader, OS, and compilers to provide seamless access to programmable hardware from programs is a research challenge. The proposed initial system has more than four years of research work as its basis. It will result in direct and significant technological advances for trusted computing, an increase in the number and type of researchers able to perform significant research work in the area of trusted computing, and enable more frequent trusted computing research and technology transfer breakthroughs in the future. More specifically, the project will bring together the following unique capabilities in one clustered computing system:

1. A flexible computing cluster, which incorporates the extensive use of programmable hardware (FPGA boards) and a novel software stack to enable rapid design, deployment, and experimentation with future highly trusted systems and applications.
2. Tools and well-defined procedures to support the configurability of both hardware and software components and to enable cost efficient customization of system services exposed to end users, including error and intrusion detection, security vulnerability masking, and recovery.

3. A rich collection of APIs (or gateways) to facilitate efficient communication and information flow between levels in the trusted architecture while preserving specified reliability and security guarantees. A key objective in the creation of these APIs will be to enable seamless integration of configurable hardware and software stack including the operating system, the middleware, and the application. Determining the appropriate API's is not obvious. If too much is exposed, the APIs will be too complex to be usable. If too little is exposed, not enough information will be accessible to be useful. Research work conducted using the trusted ILLIAC will determine the proper balance for particular application classes.
4. A validation framework employing fault, error and security attack injection and security modeling to enable rapid quantitative assessment of system designs and implementations.

The creation of this trusted computing cluster infrastructure will enable a large number of researchers to work with a common tool set rather than building their own. By augmenting the cluster computing environment to handle hardware specification and integrating FPGA synthesis tools into the cluster tool chain, we will open up the possibility of hardware-based trusted computing to a body of researchers well beyond the hardware specialist. Likewise, our integration and standardization of a validation environment that includes fault and attack injections and detailed monitoring will broaden the pool of researchers who will be able to participate in trusted computing research with a solid validation component.


## 4. Infrastructure and Facilities

Figure 1 shows an overview of the proposed Trusted ILLIAC hardware infrastructure for one of the 256 nodes. Each node is a ProLiant server, which consists of a rack-optimized chassis incorporating a two-way 2.2 GHz 64-bit processor pair augmented with a large field-programmable gate array (FPGA). The primary interconnection method to other nodes is via hierarchical 4X InfiniBand with a secondary gigabit Ethernet functioning primarily as a monitoring network, but capable of supplemental data transfer. A low bandwidth management network is also incorporated, and a tertiary experimental shadow network is planned for security/reliability evaluations and directly accesses the programmable elements. Both serve the interconnect needs of the control stack.

An infrastructure of this large size is needed for 3 reasons. First, there are real research issues related to the scalability of the application execution stack and system control stack while providing integrated protection that can only be addressed using a large-scale testbed. While programmable hardware has great potential to help address the scalability issue, there are uncertainties in the final outcome, especially due to the fact that programmable hardware must be incorporated with the software stacks in order to achieve full benefit. The proposed scale of the system is the bare minimum that will be required for us to address the scalability issues.

Second, the proposed scale gives the infrastructure the ability to test large-scale applications that require a high-level of trust. It will allow us to more meaningfully interact with industry data centers and utility computing centers. It will give us the critical mass to attract industry participation in our activities. Third, the proposed scale of the system will be large enough to accommodate multiple projects to use the system at the same time while interacting with each other. Over all, we feel that the proposed size of the infrastructure will likely be a sweet spot as we optimize the cost-benefit of the system by taking the hardware cost, staffing cost, and potential impact into consideration.

The reconfigurable component consists up to two large SRAM-based FPGAs mounted on a PCI-X, to be upgraded to PCI-Express card later upon availability. The VirtexPro FPGA will incorporate two "hard" (i.e. silicon-fabricated as opposed to IP core) PowerPC processors operating at clock rates of up to 450 MHz. The PowerPC processors will execute enhanced Linux to perform management and supervisory functions related to accelerator operation. Local expertise and research associations with the primary vendor have resulted in extensive Linux development for platform FPGAs, and this model will continue to be reflected in security/reliability accelerators. The selected FPGA will have 50K-100K standard logic elements (equivalent to approximately 5 million gates by traditional ASIC measurement methods). This FPGA component also contains 4 to 8 megabits of fast distributed SRAM adjacent to, and tightly coupled with, 200 to 400 dedicated 18x18 multipliers, which are anticipated to perform critical functions in security and encryption operations.



**Figure 1: Trusted ILLIAC node architecture**

Experience indicates that a sufficient quantity of fast scratch memory is crucial to hardware acceleration; therefore the PCI card will incorporate a minimum of 16 MB of dual-bank 166 MHz SRAM. Depending upon the specific task, it may also contain an additional 256 MB of SDRAM for intermediate result storage. A modular approach wherein one of three daughter cards can be mounted on the PCI board is anticipated. One card will have a single FPGA and SRAM/SDRAM, another type combines two identical FPGAs with SDRAM, and a third incorporates a single VirtexPro, memory, and a demountable InfiniBand connection to enable the above-described shadow security network. This scheme additionally permits future FPGA component upgrade.

The cluster system uses the software and hardware building blocks to implement control and execution stacks that enable the trustworthy execution of applications. The application execution stack uses the building blocks to present trust features as well as efficiency enhancing mechanisms to the application. The control stack, which also uses the building blocks, is a separate organization that vertically spans the software and hardware layers as the backbone of dependability and security features. The control stack also enables validation and assessment of activities in the application execution stack through its fault injection, attack generation, and instrumentation facilities.

## 5. Project Description

The Trusted ILLIAC project seeks to design, develop, and deploy a unique computing facility capable of here-to-fore impossible levels of performance, security, reliability, and adaptability. As the sub-45 nanometer era of silicon feature size is entered, traditional models of sequential computing are approaching their inherent limits in both frequency and power scaling. At the same time, burgeoning application complexity and vastly increased security and database interactions demand ever-increasing levels of performance. What makes this an even more challenging problem is that we must be able to move the entire software base forward, rather than just trying to cover a few special applications.

Once a core of the Trusted ILLIAC infrastructure is developed, we will make the infrastructure software and hardware design available to the broader academic community beyond UIUC, so other institutions can create their own integrated trusted computing testbeds. We will use the NCSA/UIUC Open Source License to distribute our software stack work, as exemplified by our OpenIMPACT distribution ([www.gelato.uiuc.edu](http://www.gelato.uiuc.edu)). We will also make our Trusted ILLIAC system available to researchers and institutions that do not have the resources to build and maintain such a large physical installation.

Finally, the Trusted ILLIAC platform will benefit technology transfer efforts from research to the “real world” where the results will have a *real* effect. The testbed will enable researchers to collaborate with developers from the government and industry to iterate and determine how adaptive, configurable, integrated, and combined HW/SW trust mechanisms can be integrated into products.

## 6. Indicators of Success and Major Accomplishments

The primary focus of Year 1 is on hardware system construction and initial software stack development. Deployment of the main 256 node system, however, is dependent upon external factors such as shipping schedules, installation and validation requirements, facility availability, etcetera. In order to decouple installation from research activities, a plan was implemented wherein a prototype system would be obtained and deployed, in advance of and parallel to, the main hardware.

This machine, termed the *Trusted ILLIAC Prototype*, (obtained with non-NSF funds), was designed to serve three primary purposes: first to act as stand-in until the main hardware achieved full functionality, second, perform as a rapid prototyping area where test applications can be staged and quickly evaluated, and third, with project maturity, to serve as a testbed for emerging technologies for potential inclusion in the full system. The prototype computing cluster has been assembled in advance of the Trusted ILLIAC 512 processor cluster scheduled for delivery in 1H06. This development machine is comprised of 8 nodes of enterprise-class dual processor chassis, each of which is 4U in size to permit extensive external instrumentation, and equipped with a single PCI card-based FPGA board supplied by the collaborating vendor on the CRI grant, and identical to units slated for the main cluster. Expansion capability was incorporated for technology evaluation, including cooling and power well in excess of anticipated requirements. An in-house designed front panel monitoring system was incorporated, as were multiple networks and high-performance local switch.

The prototype cluster components were ordered, assembled, and became functional in mid-May. A separate domain was created for exclusive use by collaborating researchers, and the prototype cluster bound to that domain. A fully automated installation of Debian Linux was performed, and user logins enabled; acquisition of open-source clustering system management utilities will take place as required. All FPGA boards were installed, and the Nallatech FUSE connection software installed on each node. This environment was configured to be interrupt driven, and suitable kernel flags to support operation added to the standard booting method. With this step, users now have full access to the reconfigurable hardware and porting of existing hardware accelerators has commenced.

## 7. Publications/Patents

No publications or patents since we are only two months into the project.

## 8. Outreach Programs and Supported Students

On May 10, 2006, UIUC Chancellor Richard Herman hosted “The Trusted ILLIAC Cluster Launch” at The University Club of Chicago. With attendance of some 200 technology leaders from across the nation and beyond, the response to the announcement of the Trusted ILLIAC was unanimously hailed as the first step into a new era of computing. Highlight of the event was a “Corporate Technologists Roundtable,” whose participants included David Belanger, Chief Scientists and Vice President, Information and Software Systems Research, AT&T Labs; Allan Cante, President and Chief Executive Officer, Nallatech; Michael Goddard, Director of Advanced Technology Planning and Performance Computing, AMD; Ed Hammersla III, Chief Operating Officer, Trusted Computer Solutions; Matthias Kaiserswerth, Vice President of System Management in the Security, Privacy, and Compliance Unit, IBM; Patrick Lysaght, Senior Director, Xilinx Research Labs; Joe Pato, Distinguished Technologist, Hewlett-Packard Labs, and Paul Steinberg, Vice President of Technology, Networks, and Enterprise Business, Motorola, Inc.

The event was announced in The Chicago Sun-Times (5/8/06) and covered in a story by The Chicago Tribune (5/11/06), which was picked up by The Seattle Times (5/15/05), The Norman Transcript (5/14/06), and Dr. Dobb’s online (5/3/06).

## 9. Future and Strategic Directions

In an extension of the current investigation, Trusted ILLIAC researchers are examining next-generation hardware in collaboration with AMD and a Chicago-based startup for early evaluations of pre-release hardware; beta development is underway on a standalone experimental node. This effort will evaluate a significantly closer coupling of the configurable fabric based upon HyperTransport, which in turn will greatly enhance the responsiveness of security and reliability applications through reduced latency.

# RI: The Metaverse: A Laboratory for Digital Media Networks

Award Number EIA-0101242, 2005-2006 Report

Departments of Computer Science, Mechanical Engineering, Education, and Art  
University of Kentucky, Lexington, KY 40506

K. Calvert, Z. Fei, J. Griffioen, C Jaynes, D. Maloney, J. Mazur, J. McDonough, W.B. Seales  
{calvert,fei,griff,jaynes,Maloney,mazur,mcdonough,seales}@netlab.uky.edu

<http://www.netlab.uky.edu/metaverse>

## 1. Abstract

*The goal of the Metaverse Project is to develop self-calibrating, self-monitoring, collaborative, low-cost immersive environments. This past year, we completed the design and installation of a new metaverse portal and enhanced our existing portals. We developed techniques to improve the accuracy of our calibration system, enable digitization of embedded cultural artifacts, support a high-bandwidth transport service, automatically adapt to network or component failures, and enhance the security of the system.*

## 2. Introduction

The goal of the Metaverse Project is to develop and deploy low-cost, reconfigurable, collaborative, immersive networked environments that offer a visually compelling collaborative space in which people interact with virtual (simulated) environments and with each other. Our research in immersive environments emphasizes the importance of self-monitoring, self-configuring, and loosely configured sets of rendering elements. A rendering element is composed of a digital projector, a camera, a network card, and a graphics accelerator. Given a set of rendering elements, called a *Metaverse Portal*, we are developing methods that allow rendering elements to detect their relative positions within a display environment (calibration), to determine an appropriate scheme for cooperative rendering of a single visual display (distributed rendering), to continuously monitor the display environment for transient artifacts such as shadows or device movement.


## 3. Goals

The research goals and activities for this past reporting period included (1) improving the accuracy and utility of multi-projector displays, (2) enabling preservation and display of historical artifacts, (3) building a high-speed, reliable transport service, (4) enabling secure group communication, and (5) using the metaverse to implement an artistic, immersive environment – i.e., an immersive art exhibit.

## 4. Infrastructure and Facilities

Over the past several years we have developed multiple metaverse portals, each with a different architecture and research focus. The current set of deployed and actively used metaverse portals include, (1) the **Collaborative Rendered Environments (CoRE) lab** which now has upward of 24 projectors displayed on 3 walls and the floor and is used to test and evaluate our core calibration technologies, (2) the **Mechanical Engineering (ME) visualization lab** has 16 projectors that create a seamless display across 2 walls and the floor and can also run off-the-shelf open-GL software for scientific visualization, (3) the **UPR Metaverse Environment Remote Core Portal (MERC) lab** at the University of Puerto Rico Rio Piedras, a four projector display system used for remote experiments between UK and UPR and by Faculty in the College of Natural Sciences at UPR, and (4) the **Digital Object Media Environment (DOME) portal**, an inverted portal of tiled projectors placed inside a back-projected half-sphere.

This past year we completed the design and installation of a new portal called **the Digital Research Theater (DiRT) lab**. The DiRT lab achieves an immersive look-and-feel with a curved surface screen and the use of stereo rendering techniques. It is a 9-projector system that uses a back-projection screen to create crisp 3D images. It is being used to visualize medical data sets and scanned 3D historical artifacts. Figure 1 shows the small set of projectors needed by the DiRT lab (Figure 1a), the curved display surface prior to calibration (1b), and a stereo image after calibration (1c).



**Figure 1.** The DiRT Lab

## 5. Project Description

As stated in the introduction, our research on immersive environments emphasizes the importance of a self-monitoring, self-configuring, and loosely configured sets of rendering elements (PC, projector, and optional camera) into easily deployed metaverse portals that can be connected via networks to form a shared collaborative visualization environment. Not only are we exploring new visualization algorithms and network algorithms, we are also exploring how these inexpensive, easy to deploy, immersive environments can be used in, and can benefit, the areas of scientific visualization, teaching, education, and the visual arts.

## 6. Indicators of Success and Major Accomplishments

We made several advances this year, both in terms of new infrastructure, and in terms of scientific discoveries/application development. The following briefly highlights some of our research accomplishments.


### 6.1 Improved Accuracy of Multi-Projector Displays

We extended traditional auto-calibration methods to support arbitrary surface geometry using a combination of multi-view geometry and traditional close-range photogrammetry. With our algorithm we support fully-automatic alignment without making assumptions about underlying room/display geometry. The method combines traditional close-range photogrammetric techniques with recent results in computer vision (5 point solutions) to simultaneously recover surface shape and the position/intrinsics of all projectors.

Recently we have introduced a solution to the simultaneous estimation of radial distortion and linear multi-view geometry for projective (camera/projector) devices. The technique extends the work of Pollefeys et. al. by casting the estimation problem as a quadratic Eigenvalue problem (QEP) where matchpoints over each of the devices are used to compose the design matrix. Earlier restrictions the QEP assumed rectangular matrices and could only support a minimal set of constraints. We solved the case in which QEP involves rectangular matrices and showed that it improves calibration accuracy and robustness as more views are introduced.

In addition we have developed a new theoretical approach to intensity and color that will allow us to render on non-white surfaces and correct for projected overlap. The approach operates in the camera space to minimize the perceptual difference between any two regions on the display. We utilize a relaxation algorithm to discover the optimal set of projective transformations that modify the gamuts of each region on the display. Early results from the color balancing are promising and we have demonstrated the ability to correct for color artifacts due to illumination changes across the display, surface markings, and differences due to bulb characteristics.

Finally we have integrated our research into the VJ Juggler platform from Iowa State University. The platform will allow us to render data seamlessly across the commodity cluster in real time and supports a variety of different data formats expected to be encountered.



**Figure 2:** Seamless display across multiple surfaces with improved calibration/distortion/intensity blending. (a) Ft. Knox MOU training site observed from the air. (b) User interactively explores the model and approaches a building. (c) User within the building. The rendered geometry is warped based on our calibration solution to provide a geometrically correct view of the scene for the users eyepoint.

## 6.2 Preservation and Display of Cultural Artifacts

We developed agile methods (in software and hardware) for digitizing difficult-to-access cultural heritage objects. The primary indicator of success in this area is the acquisition of the Petroglyph models using a custom mobile digitization platform. These models are detailed and show the surface of the petroglyphs and the carvings to sub-millimeter accuracy. In particular, we have (1) scanned petroglyphs located in a remote section of Puerto Rico; (2) scanned the Parian Marble at the Ashmolean Museum in Oxford; (3) digitized a number of paintings and sculptures at the Instituto de Cultura Puertorriquena.

We also developed methods for efficiently organizing digitized models to support compelling visualization. Specifically we organized the datasets described above into a diverse set of presentation material, including interactive DVDs, OpenGL visualization programs in 3D, and documentary video.

Another major accomplishment was the development of the DiRT lab. The DiRT lab is now operational and is used to visualization video material, high-resolution imagery, and constructed models from our digitization projects. The DiRT portal is capable of displaying theater-quality video, high-resolution (tile-based) imagery and passive (polarized) 3D stereo (see Figure 1 above). These achievements have enabled us to support a number of scholarly and educational activities including promotional demonstrations to students, faculty and administrators as well as exploration of datasets with scholars.

## 6.3 High-Speed and Reliable Transfer Services

To enhance our high-bandwidth transfer service (which uses multiple paths across an overlay network), we focused on the problem of *efficient* path selection, where each “path” is a concatenation of virtual links. Prior methods had high state-exchange overhead and suffered from synchronization of state updates, resulting in selection of non-optimal paths. We developed a simple method based on knowledge of a node’s nearby neighbors and the neighbors of the other endpoint of the connection. We demonstrated substantially better throughput with substantially reduced overhead, and have validated our results through experiments on the University of Kentucky’s Emulab and on PlanetLab. We also designed an abstraction (API) that enables applications to make optimal use of the multipath transfer layer. The research issue is how to combine application knowledge about the relative importance and logical grouping of data, with the transport service’s knowledge of the state of the network. Our API allows the application to provide information to the transfer layer to enable it to achieve the application’s performance goals.

We also studied the problem of how to deal with failures and ungraceful departures of application end points in overlay multicast. Fast detection is key to minimizing the disruption of service. We proposed a cooperative failure detection mechanism to reduce the failure detection time. We quantified and studied the tradeoff among three important measures, i.e., the expected detection time, the probability of false failure detection, and the overhead. The analysis and simulations show that the proposed cooperative failure detection mechanism can significantly reduce the failure detection time while maintaining the probability of false positive at the same level, at the cost of slightly increased overhead. We also proposed a proactive approach to deal with the problem of fast recovery of overlay multicast tree after detecting a failure. The salient feature of our approach is that each (non-leaf) node can compute a recovery plan independently.

## 6.5 Secure Group Communication

A key problem in secure multicast is the reliable and timely distribution of key information. The transmitted information is encrypted with a group key that is known to all receivers; when the group membership changes, the key must be changed. A number of schemes have been developed for imposing structure on the set of receivers to make this “re-key” process more efficient. Our work focuses on improvements to this process. One aspect is the development of a new class of policies for determining when to re-key the group; in situations where the membership is changing frequently, it is not practical to re-key after every change. The typical approach is to change the transmission key after a certain number of membership changes, or after a fixed period of time has elapsed. Our policies instead take into account both the change in group membership, and the amount of data that has been “exposed” to non-members due to such changes. This approach has two advantages. First, it makes the tradeoff between efficiency and security more explicit. Second, it smooths out peaks in the demand for computing cycles (re-keying is a compute-intensive process), making it possible to support a larger number of secure multicast groups.

Another aspect is reliability: re-key messages must be transmitted reliably to group members, but the multicast service used for transmission is best-effort. Others have proposed the use of Forward Error Control (FEC) for reliable transmission of re-key messages. Since receivers are typically organized into a tree structure for the purposes of key management, we propose to have group members who are close to each other in this “key tree” assist each other with recovery of lost re-key messages. That is, if a member fails to receive re-key information it needs to continue receiving data, it first tries to get the missing information from its neighbors in the key tree. To ensure that neighbors are likely to have received the missing data, receivers are placed in the key tree in such a way that neighbors in the key tree are likely to be located in different parts of the network; since multicast losses are correlated according to network location, this increases the likelihood that some member of a cluster of neighbors in the key tree receives each packet. Simulation studies have shown that this techniques performs comparably to FEC, while requiring both less computation capacity and less bandwidth at the sender.

## 6.6 Artistic Media

To study the intersection of virtual reality, reality television, and video surveillance, Dr. Maloney (College of Fine Art) created a “Conversations” exhibit, an interactive multi-projector video experience that surrounds the user with eight different viewpoints of a series of conversations immersing the visitor in a virtual conversation.

The first viewing of “Conversations” was in Lexington in May 2005. The Lexington public received the piece well, and Dr. Maloney was encouraged by curators to show the piece in national venues. In the fall of 2005 “Conversations” was selected by jury into the ARC Invitational Gallery in Chicago, Illinois for exhibition during December 1-31, 2006. ARC is an internationally recognized alternative space and has built a reputation for consistently exhibiting innovative and experimental art forms. By exhibiting at ARC this work will be seen and reviewed by critics and the public at large from the Chicago area. It will expose the work to a much greater audience and will also provide a national public face to the Metaverse research here in Lexington.

## 7. Publications

G. Landon and W.B. Seales, “The Museum and the Media Divide: Building and Using Digital Collections at the Instituto de Cultura Puertorriquena”, D-Lib Magazine, 11(5), 2005, ISSN: 1082-9873.

G. Landon and W.B. Seales, “Building and Visualizing 3 {D} Textured Models for Caribbean Petroglyphs, 21<sup>st</sup> Congress of The International Association of Caribbean Archaeology (Trinidad and Tobago), July 2005.

L. Wang, J. Griffioen, and K. Calvert, “Estimating Achievable Throughput”, Short paper appearing in the poster proceedings of SIGCOMM 2005, August, 2005.

K. Calvert, J. Griffioen, and S. Wen\*, Scalable Network Management using Lightweight Programmable Network Services, Special Issue on Management of Active and Programmable Networks, Journal of Network and System Management, Vol 14, No 1, March 2006

Q. Zhang and K. Calvert, "A Peer-Based Recovery Scheme for Group Rekeying in Secure Multicast", to appear in International Journal of Network Security.

Zongming Fei, Mengkun Yang, "A proactive tree recovery mechanism for resilient overlay multicast," IEEE/ACM Transactions on Networking, accepted.

Zongming Fei, Mengkun Yang, "A segmentation-based fine-grained peer sharing technique for delivering large media files in content distribution networks," IEEE Transactions on Multimedia, accepted.

Boxuan Gu, Hu Wang, Zongming Fei, "An Efficient Update Algorithm for Supporting Mobility in Structured P2P Systems," Proceedings of International Conference on Communications in Computing (CIC'06), Las Vegas, Nevada, June 26-29, 2006.

Zongming Fei, Mengkun Yang, "Intra-session fairness in multicast communications," Telecommunication Systems, vol.29, no.4, pp.235-255, August 2005.

Mengkun Yang, Zongming Fei, "Cooperative failure detection in overlay multicast," Proceedings of IFIP Networking 2005, LNCS 3462, Waterloo, Ontario, Canada, May 2005.

## **8. Outreach Programs and Supported Students**

We have plans to make the Digital Research Theater (DRT), which is the production Metaverse environment, a crucial part of our outreach and recruiting effort for the Computer Science Department at the University of Kentucky. The Chair has allocated continued resources for its operation and for a recruiter to develop materials that explain research and educational objectives. We believe these outreach activities will have a positive effect on our enrollments and on our visibility within the community. Two students from the University of Puerto Rico (summer exchange program) will work in the summer 2006 on Metaverse-related projects.

The grant supported one student, Cindy Lio, through matching funds provided by the Center for Computational Science at UK. Upwards of 10 additional students (RAs) supported through other funding sources participated in the development of the metaverse infrastructure or used it for their research.

## **9. Future and Strategic Directions**

We envision several future research directions, including more compelling algorithms for manipulating and transmitting 3D data. We have begun working with epigraphic models (carved text) to simulate processes such as "squeezing" (creating an inverse 3D cast) and "rubbing" (creating images from chalk and paper over the irregular surface). Natural and meaningful digital equivalents of these algorithms are surprisingly difficult to implement correctly and in real-time. A second research direction emphasizes the digitization problem and the equipment (software and hardware) that will enable non-experts to quickly and successfully acquire high-fidelity models without expert direction. Aspects of this process can be automated, but the variability in context and in object characteristics makes it a challenge. We are also interested in capturing and playing super high resolution video, where each frame of the video is at a resolution of about 8 megapixels. The detail available in such a video sequence is astounding; the computational challenges on both the acquisition and the display side are challenging.

To further improve network reliability, our future work will focus on extending adaptation techniques to improve the resilience of multimedia systems to various disruptions caused by mobility and other emerging network environments. Dealing with disruptions becomes a normal task rather than an exception in designing routing protocols and developing new services in these networks. In the area of new media, we will examine how much simultaneous information the viewer can absorb, or bombard them with the same or similar stories told by a multitude of anonymous people. This will be a cognitive investigation as well as a piece about identity.

# **RI: High-Performance and Visualization Cluster for Research in Coupled Computational Steering and Visualization for Large-Scale Applications**

**NSF CISE RI 0403313**

*PIs: Amitabh Varshney, Rama Chellappa, Ramani Duraiswami, Joseph JaJa, Dianne O'Leary  
Senior Investigators: Shuvra Bhattacharyya, Michael Cummings, Larry Davis, Leila De Florian, Howard Elman,  
Francois Guimbretiere, David Jacobs, David Mount, Hanan Samet, Alan Sussman, Fritz McCall*

*University of Maryland Institute for Advanced Computer Studies  
University of Maryland, College Park, MD 20742*

*<http://www.umiacs.umd.edu/research/GPU>*

## **Abstract**

We are building a high-performance computing and visualization cluster that takes advantage of the synergies afforded by coupling central processing units (CPUs), graphics processing units (GPUs), displays, and storage. The infrastructure is being used to support a broad program of computing research that revolves around understanding, augmenting, and leveraging the power of heterogeneous vector computing enabled by GPU co-processors. The driving force here is the availability of cheap, powerful, and programmable graphics processing units (GPUs) through their commercialization in interactive 3D graphics applications, including interactive games. The CPU-GPU coupled cluster is enabling the pursuit of several new research directions in computing, as well as a better understanding and fast solutions to several existing interdisciplinary problems through a visualization-assisted computational steering environment.


## **Introduction**

The recent rise in the capabilities and programmability of the graphics processing units (GPUs) has enabled them to rise above the threshold where they have now become powerful enough to be useful co-processors to the CPUs. The goal of this project is to study, understand, and harness the power of this newly emerging paradigm of a CPU-GPU combine for high-performance computing and scientific applications with visualization-assisted computational steering. A cluster of CPUs and GPUs offers a rich parallel computing environment – a controlled heterogeneity strengthened with the complementary capabilities of CPUs and GPUs with MIMD and SPMD dataflow paths. A CPU-GPU cluster has the added attractiveness of enabling a tightly-coupled interleaving of visualization and computation. Such a coupling permits exploring problems whose solution methodology itself can be profitably changed by using visualization to assist in computational steering and the computation to provide feedback for a richer, more expressive visualization.

The research groups that are using this cluster fall into several broad interdisciplinary computing areas. We are exploring visualization of large datasets and algorithms for parallel rendering. In high-performance computing we are developing and analyzing efficient algorithms for querying large scientific datasets as well as modeling complex systems when uncertainty is included in models. We are using the cluster for several applications in computational biology, including computational modeling and visualization of proteins, conformational steering in protein structure prediction, folding, and drug design, large-scale phylogeny visualization, and sequence alignment. We are also using the cluster for applications in real-time computer vision, real-time 3D virtual audio, efficient compilation of signal processing algorithms, large-scale modeling of neural networks, and computational radiology. An important aspect of this research is to ensure a high impact of the cluster towards educational and outreach goals. Towards this end, our graduate courses have started using the cluster. We have also enriched our current coursework with research results obtained on the cluster. The coupled cluster with a large-area high-resolution display screen is serving as a valuable resource to present, interactively explore, evaluate, and validate the ongoing research in

visualization, vision, scientific computing, human-computer interfaces, and computational biology with active participation of graduate as well as undergraduate students.

## Infrastructure and Facilities



The overview of our CPU-GPU Cluster architecture appears to the left. Our current cluster configuration is:

- 13 CPU/GPU display nodes driving 25 Dell 24-inch 1920 x 1200 LCDs
- 2 CPU/GPU display nodes driving 4 rear-projection DLP projectors
- 12 CPU/GPU compute nodes
- 1 CPU/GPU direct interaction node
- 3 storage nodes with 10 TB of storage
- 1 scheduler node

All nodes are connected by a combination Infiniband/Myrinet network as well as gigabit ethernet. Except for the file servers and scheduler, each node has dual Intel Xeon 3GHz CPUs, 8GB of RAM, and 100 GB disk. The LCD wall nodes have 512MB NVIDIA GeForce 7800 GTX GPUs, while the other CPU/GPU nodes have 256MB NVIDIA GeForce 6800 Ultra GPUs. We are currently looking into acquiring motherboards with 4 PCIe slots- this would allow us to drive more display wall panels per display node, thus freeing more nodes to serve as dedicated CPU/GPU compute nodes. We have installed the GPFS cluster file system, and 64-bit Linux OS on all the nodes. Our cluster and the tiled display wall are shown below.



We have installed several software packages on the pilot cluster. In anticipation of the larger cluster, our work on the operating system has focused on developing an automated installation of the customized components and automated mechanisms for updating them across a much larger cluster. We have configured and tested a number of software packages including: Chromium, DMX, BrookGPU, NVIDIA Cg, the NVIDIA SDK, and the NVIDIA Scene Graph SDK, GNU Compiler Collection, NAGware FORTRAN Compiler, the Fastest Fourier Transform in the West (FFTW), and MatLab. We have integrated our CPU-GPU cluster with our standard resource allocation mechanisms: the Portable Batch System, the Maui Scheduler, and Condor. Users can request immediate allocations for batch or interactive jobs or they can reserve nodes in advance of a demo. Unused nodes are donated to the condor pool and may be used by any researcher in the Institute.

We have also started deploying an infrastructure for interacting with the cluster using standard desktops, cluster-connected workstations, and large-scale display technologies. We have installed the cluster nodes and cluster-connected workstations in a secure data center and their displays and terminals in the graphics lab and a neighboring public lab. We are using digital KVM as a practical technology for interacting with the CPU-GPU nodes from a data center console or for remote IP clients. Although this technology does not support high resolution video or suitable refresh rates, it is invaluable for troubleshooting and managing the nodes. We are also evaluating tools to provide software-based remote display access based on VNC and the HP Remote Workstation Solutions. These solutions

allow remote users to monitor the cluster displays at reduced frame rates over local area networks. This is a key tool for developers, who may wish to work from their own desk while interacting with the cluster.

## Major Accomplishments

A number of research groups are using the cluster and making excellent progress towards their goals. We outline just a few of these below:

### *Visual Analysis of Large-Scale Time-Varying Data*

Faculty : Joseph Jaja and Amitabh Varshney

Research Scientist: Qingmin Shi

Graduate Students: Qin Wang and Jusub Kim

We have worked on the problem of isosurface extraction and rendering for large scale time varying data. Such datasets have been appearing at an increasing rate especially from physics-based simulations, and can range in size from hundreds of gigabytes to tens of terabytes. . As an example of such a dataset, consider the fundamental mixing process of the Richtmyer-Meshkov instability in inertial confinement fusion and supernovae from the ASCI team at the Lawrence Livermore National Labs. This dataset represents a simulation in which two gases, initially separated by a membrane, are pushed against a wire mesh. These are then perturbed with a superposition of long wavelength and short wavelength disturbances and a strong shock wave. This simulation took 9 days on 960 CPUs and produced about 2.1 terabytes of simulation data. The data shows the characteristic development of bubbles and spikes and their subsequent merger and break-up over 270 time steps. Each time step is simulated over a  $2048 \times 2048 \times 1920$  grid, has isosurfaces exceeding 500 million triangles with an average depth complexity of 50. Such high resolution simulations allow elucidation of fine scale physics; in particular, when compared with coarser resolution cases, the data allows observations of a possible transition from a coherent to a turbulent state with increasing Reynolds number. Isosurface extraction and rendering is one of the most widely used visualization techniques to explore and analyze such datasets. A common strategy for isosurface extraction involves the determination of the so-called active cells followed by a triangulation of these cells based on linear interpolation, and ending with a rendering of the triangular mesh. We have developed a new simple indexing scheme for parallel out-of-core processing of large scale datasets, which enables the identification of the active cells extremely quickly, using more compact indexing structure and more effective bulk data movement than previous schemes. The highlights of our work are:

- Our serial algorithm uses a smaller indexing structure and a more effective bulk data movement than the best known previous algorithms while achieving similar asymptotic bounds. In particular, the size of our indexing structure for each step of the Richtmyer-Meshkov data set is only 6KB.
- Our scheme can be implemented on a distributed storage multiprocessor environment such that the data distribution across the local disks of the different processors results in a provably balanced workload irrespective of the isovalue. Moreover, the total amount of work across the different processors is about the same as that required by our efficient serial algorithm.
- Our experimental results show that we can generate and render isosurfaces at the rate of 3.5 – 4.0M triangles per second on the Richtmyer-Meshkov dataset using our algorithm on a single processor. On an 16-node cluster, we achieve scalable performance across widely different isovalues with a performance of up to 60 million triangles per second. The experimental results also show that our algorithm achieves load balancing across the processors for a wide range of isovalues.

### *Large Multi-particle Systems Dynamics using the Fast Multipole Method*

Faculty : Ramani Duraiswami and Nail Gumerov

Graduate Student: Fei Xue

Large multi-particle systems are challenging for simulation as their cost rises with the problem size. The bottleneck is often the prohibitive computation of pairwise forces, which costs  $O(N^2)$ . The fast multipole method (FMM), first proposed by Greengard, has become one of the most powerful methods since 1980's to speed up the computation by reducing the cost to  $O(N \log N)$ . It is worthwhile to emphasize that the FMM has wide applicability, and our work on this model problem should apply to other problems where it can be used. These problems include ones in acoustics, electromagnetics, elasticity, statistics, machine learning and computer vision. In principle, FMM can be used whenever there is a summation of a function of the distance of two point sets

$$f(x_i) = \sum_{j=1}^N C_j \Phi(x_i - x_j) \quad i = 1, \dots, M$$

where  $x_i$  and  $C_i$  are the position and the intensity of  $i$ -th point, respectively.  $\Phi$  is some function (e.g. fundamental solution of Laplace equation) centered at  $x_j$  to be evaluated at  $x_i$ . Straightforward evaluation of all  $f(x_i)$ , which is in fact a multiplication of  $M \times N$  matrix by an  $N$  vector, requires  $MN$  operations. Our contributions in this field involve three components: parallelization of the fast multipole method (FMM), multi-scale time stepping integrators, and mapping the computation of nearby particles forces to GPUs. Possible applications of the target software include, but not limited to, efficient simulation of classic N-body problems with fairly uniform distribution and finding minimum Coulomb potential in modern sphere equidistribution. We have finished a parallelized FMM with multi-scale time stepping. Our initial results on a parallel CPU implementation on the cluster are highly promising; we can achieve ideal speed up ratio by using the parallelized code for fairly large  $N$ . We have developed a mapping of FMM to CPU/GPU cluster. The work to be done by GPU is the direct evaluation of accelerations given by particles in the neighboring boxes, which is a sparse matrix-vector product and has to be processed block by block as a sequence of dense matrix-vector product. We are currently extending this work to take advantage of the highly efficient texture processing capabilities of GPUs for dense matrix-vector multiplies.

### ***Cholesky Decomposition and Linear Programming on GPUs***

Faculty : Dianne O'Leary

Graduate Student: Jin Hyuk Jung

In this work we present an efficient algorithm for solving symmetric and positive definite linear systems using triangular update on a GPU. Using the decomposition algorithm and other basic building blocks for linear algebra on the GPU, we demonstrate a GPU-powered linear program solver based on a Primal-Dual Interior-Point Method. A system of linear equations,  $Ax = b$ , where  $A$  is a large, dense  $n \times n$  matrix, and  $x$  and  $b$  are column vectors of size  $n$ , can be efficiently solved using a decomposition technique, LU for instance. If the matrix is symmetric and positive definite, Cholesky decomposition is the most efficient in solving the system. The decomposition algorithm involves 3 routines (square rooting, normalizing, and updating submatrix) at each iteration. Each routine can be well matched to a GPU kernel. The most time consuming part among the three is the submatrix update routine. The update must take place only in the triangular region to be most efficient, but specialized languages such as BrookGPU don't support designating a triangular domain as an output target. A column-wise multi-pass update can be used to mimic the triangular update, but this approach cannot fully utilize the parallel architecture because each pass has several processors unused. However, single pass update over a triangular domain can be implemented with the OpenGL API as drawing a triangle is truly a native feature of the GPU. An oversized triangle, the vertices of which are at  $(k + 1, k)$ ,  $(k + 1, n)$ , and  $(n + 1, n)$ , is drawn to cover the lower triangular matrix at the  $k^{\text{th}}$  iteration. Two sets of texture coordinates can be assigned to the vertices to generate indexes for the active pivot and the active neighbor. In fact, the index pairs passed to a fragment processor are the same as in the LU algorithm. However, the index pair for the pivot must be treated differently, since our Cholesky decomposition doesn't update the upper triangular part. The index pair needs to be swapped, which can be handled on the GPUs by a swizzle operator at no cost. The swizzle operator reorders the coordinates of a multidimensional variable. This technique reduces the number of instructions from 6 to 4. In addition, the two index pairs share an invariant index  $k$ . Thus, it can be put at a  $z$  coordinate which is invariant to the  $x$  and  $y$  coordinates. The pairs can be packed into, interpolated from, and restored (using swizzle) from a single 3D texture coordinate to reduce the rasterizer's work. Using our algorithm and other BLAS kernels, we have also shown how to build a GPU-powered primal-dual interior-point method for linear programming with minimal feedback to the CPU.

### ***Parallel Rendering***

Faculty : Amitabh Varshney and Joseph JaJa

Graduate Students: Derek Juba, Youngmin Kim, and Chang Ha Lee

Recent trends in parallel computer architecture strongly suggest the need to improve the arithmetic intensity (the compute to bandwidth ratio) for greater performance in time-critical applications, such as interactive 3D graphics. At the same time, advances in stream programming abstraction for graphics processors (GPUs) have enabled us to use parallel algorithm design methods for GPU programming. Inspired by these developments, we have explored the interactions between multiple data streams to improve arithmetic intensity and address the input geometry bandwidth bottleneck for interactive 3D graphics applications. We have introduced the idea of creating vertex and transformation streams that represent large point data sets via their interactions. We have shown how to factor such point datasets into a set of source vertices and transformation streams by identifying the most common translations amongst vertices. We accomplish this by identifying peaks in the cross-power spectrum of the dataset in the Fourier domain. Our scheme can achieve a factor of 2X to 5X reduction of geometry bandwidth requirements.



We are also exploring better load-balancing strategies with Kitware's ParaView system to enhance parallel rendering performance on the CPU-GPU cluster using MPI. For better load balancing we are exploring techniques that work directly on the raw volumetric scalar fields as well as on compact representations using anisotropic radial basis functions (RBFs). Representing volumetric scalar fields with anisotropic RBFs has allowed us to trade off compressed data size with error control by varying the number of basis functions used in the projection. The compactly supported nature of these basis functions eases both compression and rendering, since any given point in space is affected by only a limited number of basis functions. Encoding to and rendering directly from this format using a GPU-accelerated approach is one of the areas of research we are

currently exploring. The figure above shows the parallel rendering of the LLNL Richtmyer-Meshkov data with each node's rendering represented by a different color.

## Education and Outreach

One of the important contributions of this infrastructure has been in the education of graduate and undergraduate students at the University of Maryland. During Fall 2006, PI Varshney taught his graduate computer graphics class (CMSC 740) on the cluster. Co-PI Duraiswami has been using the cluster for projects in his Advanced Scientific Computation courses (AMSC 663/664). A number of undergraduates have been involved in various stages of the building, deployment, and use of the CPU-GPU cluster. The CPU-GPU cluster has also been serving a pivotal role in bringing together a community of scientists from different parts of the campus together. Professor Bill Dorland and Dr. George Stantchev (Center for Multiscale Plasma Physics, UMD) have been using the cluster to explore gyrokinetic data for the study of turbulence in high temperature, ionized gas using GPUs for stochastic volume visualization. Professor Raj Shekhar (Department of Diagnostic Imaging, UMD Medical Center) is using the cluster for registration of volumetric data from multiple modalities, and is using the GPUs for rapidly solving the non-linear warping equations. Professor Jim Reggia (Computer Science, UMD) is involved in a simulation of large-scale neural networks on the GPU cluster. The cluster is also accessible by Life Scientists throughout the University of Maryland using *Lattice* – a Globus-based grid developed by Mike Cummings, one of the senior investigators on this project.

## Conclusions

A CPU-GPU combine offers a very promising platform for a number of high-performance computing research areas. This unique computational infrastructure is enabling us to pursue several new research directions in computing, as well as enable a better understanding and fast solutions to several existing interdisciplinary problems through a visualization-assisted computational steering environment. In addition, the research infrastructure resource is also greatly assisting in education for undergraduate and graduate students through its impact on courses as well as through better research opportunities.

## Publications

Following publications have resulted from this grant thus far:

1. Q. Wang, J. JaJa, A. Varshney, An Efficient and Scalable Parallel Algorithm for Out-of-Core Isosurface Extraction and Rendering, *Journal of Parallel and Distributed Computing* (submitted May 2006).
2. J. H. Jung and D. P. O'Leary, Cholesky Decomposition and Linear Programming on a GPU, *Workshop on Edge Computing Using New Commodity Architectures (EDGE)*, Chapel Hill, North Carolina, May 2006.
3. Y. Kim, C. H. Lee, and A. Varshney, Vertex-Transformation Streams, *Graphical Models*, (accepted April 2006)
4. X. Hao and A. Varshney, Geometry-guided Computation of 3D Electrostatics for Large Biomolecules, *Journal of Computer-aided Geometric Design*, (accepted May 2006)
5. Q. Wang, J. JaJa, and A. Varshney, An Efficient and Scalable Parallel Algorithm for Out-of-Core Isosurface Extraction and Rendering, *Proceedings of the 20th IEEE International Parallel and Distributed Processing Symposium (IPDPS'06)*, April 25 - 29, 2006, Rhodes Island, Greece

**Collaborative Research:**  
**CRI: A Testbed for Research and Development of Secure IP  
Multimedia Communication Services**

**Grant # 0551694, 2006**

*PI: Ram Dantu, Co-PIs: Elisa Bertino, Sonia Fahmy, Dipak Ghosal, Henning Schulzrinne*  
[rdantu@unt.edu](mailto:rdantu@unt.edu), [bertino@cerias.purdue.edu](mailto:bertino@cerias.purdue.edu), [fahmy@cs.purdue.edu](mailto:fahmy@cs.purdue.edu), [ghosal@cs.ucdavis.edu](mailto:ghosal@cs.ucdavis.edu),  
[hgs@cs.columbia.edu](mailto:hgs@cs.columbia.edu)

## **ABSTRACT**

The next generation Internet needs to support multimedia services such as Voice over IP (VoIP) and IP-based television (IPTV) and media distribution. These services will have to be deployed on large-scale networks including a large variety of nodes, such as residences, small, medium, and large enterprises (with several branch offices), and service providers. It is estimated that VoIP is going to reach critical mass within the next 5 years. Our survey results show that many federal agencies are already putting migration strategies in place. We anticipate that several security holes and vulnerabilities will be discovered during the deployment and usage, especially in view of the fact that VoIP will have to interoperate with the conventional Public Switched Telephone Network (PSTN). It is essential that we investigate the vulnerabilities proactively and develop algorithms and techniques to secure VoIP from security threats due to interoperability problems, lack of standards, attacks by hackers, script kiddies, spammers, corporate espionage, and terrorism. The goal of the proposed project is to develop a testbed that will enable research on understanding and analysis of vulnerabilities of VoIP networks. This project is a collaborative proposal from University of North Texas, Columbia University, University of California at Davis, and Purdue University, with University of North Texas serving as the lead institution. In particular, we plan to investigate voice spam prevention, Denial of Service (DoS) attacks, 911 emergency management, and Quality of Service (QoS) under attacks and with defenses. We plan to translate our research results to engineering guidelines for preventing security breaches during development and deployment of VoIP networks. It is important to note that VoIP infrastructure can be reused for different multimedia services like video and instant messaging. Therefore, the impact of our work is far reaching, and not limited to VoIP.

## **Introduction**

Over the past few years, there has been a rapid development and deployment of new strategic services based on the IP protocol, including Voice over IP (VoIP) and IP-based media distribution (IPTV). These services operate on private and public IP networks, and share their network with other types of traffic such as web traffic. Forecasts concerning the growth of VoIP services are particularly impressive. It is estimated that in a few years most enterprises and residences will be transitioning from circuit-switched to VoIP services. For example, the Department of Defense (DoD) and the Navy have announced the deployment of 400,000 and 300,000 IP phones, respectively. Similarly, Bank of America, American Airlines, and Ford Motor Company (with 50,000 employees in 110 locations) have announced plans for migration. Voice and video services over Wireless LAN (WLAN) and cellular networks are also on trial in several hospitals.

As Voice over IP (VoIP) technology penetrates worldwide telecommunications markets, advancements in performance, cost reduction, and feature support make VoIP a compelling proposition for service providers, equipment manufacturers, and end users alike. Not only will VoIP reduce communication costs and provide enhanced and more flexible communication experience to people, but it will also pave the

way for innovative, value-added, highly-personalized services. We can expect that interactive multimedia and broadcast video services will be reusing the infrastructure that is being deployed for VoIP. Such trends will result in what we can call *IP-based multimedia communications infrastructure*, encompassing, both the equivalent of conventional phone conversations and advanced communication and content distribution services.

In light of this growing interest, security in voice communications is evolving into a key requirement for VoIP solutions. Packet-based communication is particularly vulnerable to security risks including voice “tapping” by sniffing packets, unpaid service usage by falsification of network ID, and service disruption by packet manipulation. Mobility and multimedia services of VoIP add another dimension of complexity to the problem of security. Convergence of the two global and structurally different networks (Public Switched Telephone Network and VoIP) also introduces new security weaknesses. Voice services over wireless LANs (VoWLAN) create additional vulnerabilities. Moreover, VoIP networks are prone to virus and worm spreading through their data network elements. Finally, VoIP is expected to support high availability (e.g., 5 minutes of downtime in a year) for 911 emergency services.

Because of the increased penetration of the new services, IP-based multimedia communication services will become *a critical infrastructure* for which high assurance security is crucial. Although academic and commercial labs have been conducting studies on the security of next generation networks and VoIP, and these studies have offered fruitful preliminary results in understanding the threats and vulnerabilities, but the size and scope have been limited to the available resources at the respective labs. The goal of this *testbed* is conducting research, development, and testing of inter-domain security and QoS mechanisms for new services such as voice, multimedia, video, and 911 emergency services. *It is important to note that the SIP-based (Session Initiation Protocol) infrastructure can be used for setting up multimedia, video, and voice sessions.*

#### *Why security for VoIP*

Because Public Switched Telephone Network (PSTN) voice calls are not typically protected, so a question that naturally arises is whether security for VoIP is really necessary. The answer is twofold. First, the packet nature of IP networks makes it much more susceptible to security threats than PSTN. With the current network technology, it is easier to probe into voice information on a packet network than to physically tap into the circuit switched network. Additionally, with the new security concerns posed by current socio-political conditions, it would be beneficial to both service providers and end users to include security features on our voice networks.

From the service provider’s perspective, adopting security safeguards can prevent a variety of malicious actions that may result in theft of service and significant loss of revenue. By accessing network databases and IP addresses, fraudulent service subscription can be obtained and used without payment, or could be charged to another actual customer. Additionally, telephony end equipment might be implemented and configured such that it appears as a clone of a valid end device. Network hackers pose a threat if they can successfully access network equipment, modify the databases, or replicate the equipment, resulting in a shutdown, “jam,” or takeover of the voice network. Finally, packet network protocols, such as SIP, H.323, and MGCP (Media Gateway Control Protocol) can be manipulated by accessing the packets, modifying the protocol information, and subsequently altering the packet destination or the call connection.

Other threats concern end-user privacy. Again, by using simple packet network “snooping,” hackers can “listen” to the voice bearer channel, or “see” call setup (signaling) information, and subsequently derive call detail information. The extraction of personal information, behavior, and habits of subscribers, for illegal use can result in personal information theft. This can be accomplished by end telephony equipment clones configured to masquerade as another legitimate subscriber, by the network protocol manipulation described earlier, or by the “tapping into” or collecting the voice and related signaling traffic that is then used for off-line analysis.

## Background

The proposed testbed will support experiments that can be used for research and development of security mechanisms for new triple-play services such as voice, data, and video. A typical VoIP network includes (but is not limited to) soft IP phones, hard IP phones, IPPBX, Proxy Call Servers, Registration Servers, Location Servers, Signaling Gateways, Media Gateways, and Media Gateway Controllers. We now define each of these elements.

**IP Phone:** An Internet protocol phone (IP phone) is a device (a telephone set) that converts audio signals and telephony control signals into Internet protocol packets. These standalone devices plug into (connect to) data networks (such as Ethernet), and operate like traditional telephone sets. Normally, a standalone phone is called a *hard phone* and a PC-based phone is called a *soft phone*. In addition, we shall be testing our algorithms with Voice over WLAN phones.

**IPPBX:** IPPBX systems use an IP telephone server to provide call processing functions and switching between IP phones in a LAN segment.

**Proxy Server:** A proxy server forwards a communication session request to more than one device on behalf of the communication connection request. This server has call routing capabilities.

**Signaling Gateway (SGW):** A signaling gateway (SGW) is used to interface a signaling control system (e.g., SS7) and an IP network device (e.g., router, firewall, IPPBX, Proxy). The signaling gateway may convert message formats, translate addresses, and allow different signaling protocols to interact.

**Media Gateway Controller:** A media gateway controller is the portion of a PSTN gateway that acts as a surrogate call management system (CMS). The MGC controls the signaling gateway and the media gateway (MG).

**Media Gateway:** A device which converts IP streams (such as audio) to the TDM or analog equivalent.

**Session Border Controller (SBC):** In an enterprise network, the SBC "controls" the communication "session" as it crosses the "border" from the LAN to IP. Conventional firewalls support the secure traversal of data streams, but for IP networks, SBCs are needed to facilitate secure, real time, multimedia communication.

We plan to investigate several sample configurations, such as:

**Sample Configuration 1: Enterprise peering:** This setup typically emulates four branch offices of an enterprise. Calls are generated randomly from each enterprise, and it is expected that the PBX as well as proxy server can simultaneously handle several hundred calls. This setup can be used for the research projects related to DoS and SPAM

**Sample Configuration 2: Hosted IP Centrex:** In this configuration, all the end-users are connected to a centralized facility through an IP trunk. No IPPBX or VoIP firewall or proxy servers are used for switching or routing the calls. All the calls are routed through a central facility and a VoIP firewall and NAT are maintained by this central facility. This configuration emulates a service provider, residences, and SOHO (small office and home office) users. This configuration can be used for the research projects related to DoS, Spam, and threats on threats on 911 service

**Sample Configuration 3: Carrier peering:** Carriers will build trust relationships with their peers and subsequently end-to-end trust can be established with pre-arranged relationships. These relationships are important in preventing spamming and denial of service attacks. One of the important aspects of our research is analyzing vulnerabilities across VoIP and PSTN networks. In this context, we will be studying several adaptation protocols such as M2UA (MTP2 User Adaptation Protocol), M2PA (MTP2 Peer to Peer Adaptation Protocol), M3UA (MTP3 User Adaptation Protocol) and SUA (SSCP User Adaptation Protocol). We plan to study threats and vulnerabilities in heterogeneous networks. In particular, we will be investigating the behavior of the converged network in terms of DoS and spam attacks.

## Research Activities

### *IP-PSTN convergence vulnerability analysis*

The Internet is an open collection of interconnected networks (e.g., educational institutions, commercial and government organizations) that communicate using the Internet Protocol (IP). The TCP/IP protocol has inherent vulnerabilities that can be used to perform various attacks including, but not limited to, eavesdropping, service disruption, and unauthorized access. Since access to the Internet is not restricted (as it is in PSTN), it is easier for attackers to exploit weaknesses.

As discussed in the previous sections, the PSTN has been maintained as a closed network, with access limited to carriers and service providers. Entry to the PSTN has traditionally been protected by price, as the annual costs exceed \$100,000 per month. These two characteristics of the PSTN (closed network and very high cost of access) have established the false perception that the PSTN is a secure network. In fact, many people believe that PSTN is more secure than the Internet! This claim is quickly discredited once someone starts to analyze the security controls, or lack thereof, that are available in the PSTN. The PSTN is comprised of thousands of interconnected network elements over dedicated circuit-switched facilities that use the SS7 (Common Channel Signaling System No.7, SS7 or C7) which relies upon a model of trusted neighbors. The circuit switch facilities include ISDN, X.25, and tunneled variants of both.

Network convergence requires that the networks communicate together using common protocols that allow interoperable mapping of the inter-exchanged messages. In the absence of common protocol translators, gateways must exist. These gateways have the ability to provide protocol translation between circuit and packet based networks (i.e., SS7 and IP respectively). Vulnerabilities that exist in IP-based networks may be propagated or affect SS7 networks, since there are no concrete security controls to prevent malicious messages traversing the SS7 network. Furthermore, the convergence of the two networks and the complexity of the protocols and associated components provide a fertile ground for new types of vulnerabilities, as well as modified versions of existing vulnerabilities associated with IP-based networks or the PSTN.

### *Preventing voice spamming*

VoIP introduces a whole new set of problems for the network operators and service providers. We believe that one of the biggest risks with VoIP is its vulnerability to spamming attacks. The Internet is wide open for tapping and intrusion. In theory, anyone who can locate an IP phone via scanning the Internet, can call in, and with the cost of an Internet call near zero, this kind of vulnerability invites voice spamming. The purpose is to research various VoIP spam filters and their performance in terms of false alarms, and conduct a sensitivity analysis of various parameters. *It is not possible to collect spam without a large subscriber base. Until we reach a critical mass of VoIP subscribers, we need to research the filters using the realistic spam generated in the lab. Hence the proposed spam generator is a critical component for spam research. Moreover we need a platform for analyzing spam from VoIP network to PSTN and vice versa.* As part of the research, we will encourage students to act as black hats and white hats. This project will thus provide a unique experience for graduate and undergraduate students on a real network using hands-on experience with real protocols such as Session Initiation Protocols (SIP) and call processing software.

### *Denial of Service (DoS) and Distributed DoS (DDoS)*

The current Internet was designed under the premise that most users and traffic sources are cooperative. DDoS attacks and Internet worms have, however, become increasingly prevalent. The DDoS attacks in February 2000 brought down several e-businesses; a more massive attack in October 2002 brought down eight of the thirteen root Domain Name System (DNS) servers. Despite the multitude of products and research papers on defending against DDoS and worms, there is little work on evaluating these solutions in realistic environments.

A DoS or DDoS attack designed to render a VoIP system useless is relatively easy to launch. *Unlike traditional DDoS attacks that use mostly use high rate UDP, ICMP, TCP SYN or TCP RST floods, attackers can cripple an IP phone by creating very few signaling messages per minute. Therefore, an attacker with very limited capabilities can bring down an enterprise VoIP system (consisting of several thousand employees).* Since attack traffic is being sent at a low rate, traditional DDoS mitigation solutions that detect attack traffic by its high rate are ineffective. Simple anomaly based techniques are also ineffective since the attack traffic can resemble normal traffic. Traceback of the attacker becomes significantly more challenging. More sophisticated solutions that combine several mechanisms thus need to be researched and tested in a realistic setting. We propose to study the DDoS threat models specific to VoIP networks. *In particular, we need to precisely characterize the damage that can be done with different attacks, and the effort/cost required on the part of the attacker.* For example, we can measure the usability of an enterprise VoIP system at different attack intensities.

#### *Quality of Service (QoS) and security mechanisms*

As multimedia services are real-time, it is important to ensure that the deployed security mechanisms do not impact the quality of service (QoS) of a session when the system is not under attack, while ensuring the QoS is maintained at an acceptable level when the system is under attack. *The purpose of this research is to study the scalability (several thousands of sessions) of security mechanisms in terms of quality of service and performance degradation of the session.* These experiments can be conducted with and without attacks, and repeated with and without different security solutions. *We propose to develop instrumentation tools to measure the QoS of a session in terms of call setup time, end-to-end delay statistics, jitter, etc.* Visualization tools will be designed to aid in understanding and summarizing the results from a suite of experiments performed under different network conditions. These tools will make it easy for a researcher to evaluate the efficacy of a particular solution in a completely automated fashion. This project will be integrated into our networking and security courses. The development and use of the instrumentation and visualization tools is ideal for students to understand the tradeoffs that must be made between security and performance.

#### *Denial of 911 emergency service*

Highlighted by recent FCC mandates, emergency calling (“9-1-1”) is likely to change fundamentally within the next few years. Infrastructure for routing calls, deriving location information, and for coordinating emergency response are being developed by several vendors. *In an emergency services context, denial of service attacks can impact the availability of three types of resources, namely (1) PSAP (Public Safety Answering Point) network facilities, both at the network layer and for call signaling (2) call taker resources; (3) first responders.* Roughly, scarcity and ability to deflect attacks is largest for network facilities, using standard DoS detection and mitigation techniques. Call taker resources are scarce, but other PSAPs may be able to assist with call filtering if an attack is localized. Both of these types of attacks can be automated and are similar to the DoS attacks more typically discussed in security considerations for protocols. DoS attacks on first responders are launched by having human callers send these first responders to bogus emergencies, by providing false location or incident information. If multiple attackers are coordinating their DoS attack on first-responder resources, they do not need to remain anonymous or provide false location information, although doing so allows such attacks to be launched from a much wider geographic area.

## **Work-in-progress**

We are currently receiving and processing quotations for soft/hard phones, video phones, call generators, media gateways, media gateway controllers, and session border controllers. We hope to soon receive the infrastructure for the first phase of experiments.

# **MII: WALSAIP: wIDE Area large scale automated information processing**

*Domingo Rodriguez (PI and Research Group Leader)*  
*Yi Qian, Wilson Rivera, and Manuel Rodriguez (Research Group Leaders)*  
*Ismael Pagan and Nestor Rodriguez (Development Group Leaders)*

**CISE-CNS Grant No. 0424546 – Year II**

University of Puerto Rico at Mayaguez

<http://www.walsaip.uprm.edu>

## **1. Abstract**

This report describes the work being conducted under the WALSAIP project, a CISE-CNS grant awarded to the Electrical and Computer Engineering Department of the University of Puerto Rico at Mayaguez. The WALSAIP project concentrates on the development of an infrastructure for the treatment of signal-based information arriving from physical sensors in a wide-area, large scale setting, as it pertains to environmental surveillance and monitoring (ESM) applications. ESM deals with the gathering and processing of appropriate environmental information to aid in the process of effective decision making. The significant importance of this proposed infrastructure relies upon its unified and integrated capability to treat the distributed acquisition, storing, processing, and rendering of information in a timely and mostly automated manner.

## **2. INTRODUCTION**

This project is developing a new conceptual framework for the automated processing of information arriving from physical sensors in a generalized wide-area, large-scale distributed network infrastructure. The project is focusing on water-related ecological and environmental applications, and it is addressing issues such as scalability, modularity, signal representation, data coherence, data integration, distributed query processing, scheduling, computer performance, network performance, and usability. This new framework treats signals as elements in prescribed sets and their associated structures. The project is constructing a computing and information processing (CIP) environment to deal with the algorithmic treatment of signal-based large scale content in order to extract information relevant and important to a user. It is also developing new theories and algorithms for computational signal processing to gather, process, and represent data obtained from physical sensors. Further, it is also developing new concepts in middleware integration, distributed query optimization, distributed query processing, and distributed scheduling algorithms to adapt to an ever changing network infrastructure and provide a pathway between a physical world sensory reality with its associated physical sensors, and a user with network and database infrastructure applications.

## **3. WALSAIP GOALS AND OBJECTIVES**

This section describes the goals and objectives of the WALSAIP project. A first goal is to develop modular, reconfigurable, and scalable working prototypes in the form of systems, tools, and applications to aid and support a network infrastructure for the automated processing of signal-based information acquired from array sensors in heterogeneous, wide-area, large scale, and distributed systems. A second goal is the utilization of working prototypes in specific scientific and engineering applications dealing with identification, monitoring, assessment, and management of regional water resources in tropical areas, with particular emphasis given to the Caribbean region. We have as third main goal to use acquired experience, expertise, and developed works through this project as vehicle to enable our research groups a very discernable presence in research activities pertaining to the evolving worldwide cyberinfrastructure research thrust. The *main objectives* of the WALSAIP project are: 1) continue the infrastructure development of laboratories associated to the project such as the Advanced Data Management Lab, Automated Information Processing Lab, Computing Research Lab, and Parallel and Distributed Computing Lab; 2) foster, encourage, and promote the integration of our research and development areas to address complex, large scale, multidisciplinary problems in *information network infrastructure* associated with environmental monitoring and surveillance in tropical areas; 3) provide a formal characterization of a computing and information processing (CIP) environment over a distributed network infrastructure; 4) strengthen our research and development groups with the recruitment of outstanding graduate students at the doctoral and master degree levels and continuing this

way to increase the number of well qualified Hispanic women and underrepresented groups pursuing graduate education.

#### 4. WALSAIP PROJECT DESCRIPTION AND CONCEPTUAL FRAMEWORK

This project is tailoring its work for strategic environmental and ecological applications in the identification, monitoring, assessment, and management of hydrological events in tropical areas. Our conceptual framework is focused primarily on the manner in which systems, tools, and applications are being integrated, under a computing and information processing (CIP) environment infrastructure to deliver end-to-end information relevant to users through tailored requests. Through this proposed work we expect to provide fundamental insights on concepts, theories, rules, and techniques for the automation and integrative processes of systems, tools, and applications needed in a generalized distributed network infrastructure in order to provide the required services, features, and options requested by users. The WALSAIP Project conceptual framework centers on a computing and information processing environment as an aggregate of seven basic components: A set of input entities, a set of output entities, a database infrastructure, a set of generalized computing and information processing (CIP) operators, a set of composition rules for CIP operators, a set of actions rules for CIP operators to act on input entities in order to produce targeted and desired output entities, and a user interface.

##### 4.1. Stream Model Conceptual Framework

Figure 1 depicts a model for a wide-area, large scale, CIP environment infrastructure, with the necessary entities to be characterized as a signal-based automated information processing infrastructure. This model is commonly used in the processing of signal-based information. The model structure clearly accentuates its *stream* processing format, where signals are acquired from sensors, through a sensor array structure (SAS), and sequentially treated until the desired information is extracted. The WALSAIP project aims at improving at this standard model structure with the formulation of a new conceptual model which accentuates a *distributed space-time* processing format.




Figure 1: Standard Signal-based Automated Information Processing System

##### 4.2. Distributed Space-time Conceptual Framework

In Figure 2 we present our design of a conceptual model for the WALSAIP infrastructure, emphasizing the distributed space-time processing nature of the model structure. It is important to point out that this distributed space-time format permeates all other system substructures such as distributed sensor networks for signal acquisition, distributed databases for data management, distributed signal processing for sensor-based systems, as well as overall distributed scheduling for network workflow. A *network communications architecture* is being develop to enable this distributed space-time model structure consistent with the realization of a CIP environment in a generalized *cyberinfrastructure* framework.




Figure 2: Distributed Space-Time Automated Information Processing System

### 5. WALSAIP RESEARCH GROUPS

Currently, there are four research groups and two development groups actively working as part of the WALSAIP project. Each group is conformed by one or more professors, associated research students, and varied degree of external collaboration. A main objective of the coming third year of the WALSAIP project is to expand and strengthen these external collaborations. The research groups are: Advanced Data Management Group (ADMG), Automated Information Processing Group (AIPG), Network Communications Infrastructures Group (NCIG), and Parallel and Distributed Computing Group (PDCG). The development groups are: Hydro-ecological Research Group (HERG) and Human Computer Interfaces Group (HCIG). Below, we provide a description of the activities conducted by each research group for this past year 05/06 of the WALSAIP project.

#### 5.1 Advanced Data Management Group

The ADMG conducts research on issues pertaining to distributed databases, heterogeneous databases, Web services, e-Government, and Paper-equivalent electronic forms. Specific research issues related with WALSAIP include:

1. Design, implement and test a Database Middleware System for Wide-Area Networks in which mobile devices are treated as bona-fide data sources.
2. Develop and study algorithms for efficient query processing, including novel schemes with Quality-of-Services (QoS) guarantees.
3. Develop and study algorithms and protocol for efficient peer database selection in distributed networks.
4. Develop and study algorithms for client-side and server-side query execution recovery and result re-delivery in Wide-Area Networks.

In the NetTraveler framework, middleware system that we have begun to prototype, we model the WANs as having a collection of applications  $H = \{h_1, h_2, \dots, h_n\}$ , each having a specific role in helping the user solve a given query.

The collection of applications  $H$  is running on host computers spread over a group of LANs that form the entire WAN environment. These LANs can be made of wired or wireless technologies, such as Ethernet, DSL, IEEE 802.11b, and 3G networks. From the set  $H$  we have a subset of applications  $C = \{c_1, c_2, \dots, c_i\}, i < n$ , which have client capabilities to submit queries.

#### 5.2. Network communications Infrastructures Group

The research goal of NCIG is to study issues for distributed wireless sensor networks, wireless ad hoc networks and local area networks, and wide area infrastructure networks, and the internetworking issues. The specific objectives include:

1. To design and implement a new energy aware medium access control (MAC) protocol for wireless sensor networks with wide-area large scale environmental monitoring applications

2. To study and analyze the performance of retransmission schemes for high-data-rate MAC protocol in wireless local area networks
3. To propose and study a framework for key management schemes in distributed wireless sensor networks with heterogeneous sensor nodes
4. To apply these research ideas to the framework being designed in the WALs-AIP project

Wireless sensor networks have recently come into prominence because they hold the potential to revolutionize many segments of our economy and life, from tasks as surveillance, to widespread environmental monitoring, to manufacturing and business asset management, to automation in the transportation, to security, and health-care industries. They can be used in virtually any environment, especially in those where wired connections are not possible, the terrain inhospitable, or physical placement difficult. Wireless sensor networks consist of battery-operated sensor devices with computing, data processing, and communicating components. Energy conservation is a critical issue in wireless sensor networks since batteries are the only limited life energy source to power the sensor nodes. Like in all shared-medium networks, medium access control (MAC) protocols enable the successful operation of the network. The MAC protocol in a wireless sensor network must achieve two goals. The first is the creation of the network infrastructure. Since thousands of sensor nodes are densely scattered in a sensor field, the MAC scheme must establish communication links for data transfer. This forms the basic infrastructure needed for wireless communication hop by hop and gives the sensor network self-organizing ability. The second objective is to fairly and efficiently share communication resources between sensor nodes.

### 5.3. Parallel and Distributed Computing Group

The research goal of PDCG is to study issues for distributed processing in wide area networks. The specific objectives include:

1. To design and implement adaptive quality of services (QoS)-based scheduling strategies for distributed processing
2. To study replication schemes for distributed sensor data
3. To study adaptive service provision and orchestration in wide area networks
4. To apply these research ideas to the framework being designed in the WALs-AIP project

The WALs-AIP project aims at developing an infrastructure for the treatment of signal-based information arriving from physical sensors in a wide-area, large scale environment. The proposed model accentuates a distributed space-time processing format. This approach demands efficient data and resource management mechanisms in terms of scheduling and data replication. Furthermore, next generation of applications will be increasingly dynamic. This implies that the current static infrastructures will not be adequate unless adaptive functionalities are provided. We assume that the resources are connected via two-level hierarchical networks. The first level is a wide area network that connects local area networks at the second level. Users submit job specifications along with QoS requirements. The scheduler then discovers appropriate resources for processing the job and schedules the tasks on the resources. In order to discover suitable resources, the scheduler has to predict execution responses on the available resources and verify QoS capabilities and availability of the resources. Re-scheduling mechanisms are then implemented to adapt scheduling to service dynamics. Our scheduling strategy focuses on providing high priority to jobs with low probability of failure. To achieve this, an urgency criterion is introduced to account for relevance, laxities and probability of failures of incoming jobs. The proposed urgency criterion is a linear combination of one static parameter and two dynamic parameters. These parameters are defined as follows.

- i. **Criticality (Relevance).** This static factor is initially established by the user according to experience and/or customer importance.
- ii. **QoS (Quality of Service).** This dynamic factor applies to the matching of job needs with resource availability and capability. Different QoS metrics can be defined such as the desirable bandwidth for the application or the required speed of processors. Since we are more interested in stochastic online scheduling, we add into the accounting the probability of failure of jobs.
- iii. **Laxity.** This dynamic factor is defined as  $Laxity = Jd - (t + Jlat)$ ; where  $Jd$  is the Job deadline,  $t$  is the actual time of calculation, and  $Jlat$  is the expected latency of the job.

#### 5.4. Automated Information Processing Group

This group has the following main objectives: 1) the study of principles, concepts, methods, and rules for the measurement of observable physical quantities; 2) the analysis, synthesis, and representation imaging and visualization of distributed sensors array-based signals; and 3) the development of computational operator methods for the treatment of sensor-based large signal data sets. It concentrates on the formulation of computational methods and implementation techniques in four broad areas addressing environmental surveillance monitoring (ESM). The first of these areas is *distributed sensor networks signal processing*, where the main objective is to treat signals, using appropriate algorithm partitioning techniques, from a collection of large number of heterogeneous sensors, distributed logically or spatially over a target environment, which can be connected to a cyberinfrastructure in such a manner that measurements can be obtained from observables of physical entities. The second of these areas is *content analysis of multimedia data*, which deals with the processing of data from different media over a cyberinfrastructure in order to provide proper imaging representations. The third area deals with the *modeling and simulation for active sensor image formation*, where array processing techniques associated with detection, estimation, and modulation of signals in a stochastic environment are utilized. The last area is *aperture synthesis computational interferometry* and it deals with the fundamental study of information generation from change detection through novel multidimensional signal interferometry technique.

#### 5.5. Hydro-ecological Research Group

The HERG group has the main objective of identifying and developing a test case framework for surveillance of hydro-ecological systems. WALSAIP will permit to identify, represent, monitor, assess, and manage a regional hydro-ecological system (H-E) in the tropics. Puerto Rico H-E system representing the macro-scale has been selected as the test case. Jobos Bay National Estuarine Research Reserve (JBNERR) has been selected as the test site for being a national reserve under NOAA administration and sponsorship. This site represents a micro-scale system test case. The complexity of the system can be described as multi-variate, multi-dimensional, heterogeneous, non-stationary, highly dynamic and sensitive to extreme events.

#### 5.6. Human-Computer Interface Group

One of the main issues when dealing with satellite images is the variety of sites maintaining images repositories with different formats and different visualizations. The tools typically used to handle images target specific formats and visualizations. This problem limits the accessibility of images to prospective users worldwide. Another issue with satellite images is the lack of tools for finding images with particular characteristics. In order to get a desired image a user needs to be acquainted with the database of images and rely on their recollection of the image geographical coordinates. Thus, searching for an image requires previous knowledge of the image. One of the objectives of this work is to develop a Web-based application that can facilitate the access and manipulation of images independently of their site, format or visualization.

### 6. REFERENCES

1. Kejie Lu, Dapeng Wu, Yi Qian, Yuguang Fang, and Robert C. Qiu, "Performance of An Aggregation-Based MAC Protocol for High-Data-Rate Ultra-Wideband Ad Hoc Networks", accepted for publications, to appear in *IEEE Transactions on Vehicular Technology*.
2. Rose Q. Hu, Robert Best, Yi Qian, and Mingzhou Jin, "A Scheduling Algorithm in a Core Optical Router with Heterogeneous Traffic", *International Journal on Wireless and Optical Communications*, pp.1-22, Vol.3, No.1, April 2006.
3. Yi Qian, Rose Q. Hu, and Catherine Rosenberg, "Integrated Connection Admission Control and Bandwidth on Demand Algorithm for a Broadband Satellite Network with Heterogeneous Traffic", *IEICE Transactions on Communications*, pp.895-905, Vol.E89-B, No.3, March 2006.
4. Yi Qian, Rose Q. Hu, and Hsiao-Hua Chen, "A Call Admission Control Framework for Voice over WLANs", *IEEE Wireless Communications*, pp.44-50, Vol.13, Issue 1, February 2006.
5. N. G. Santiago, D. T. Rover, D. Rodríguez, "A Statistical Approach for the Analysis of the Relation between Low-Level Performance Information, the Code, and the Environment," *Proceedings of the SIAM Journal of Parallel and Distributed Computing Practice*.
6. Elliot Vargas-Figueroa and Manuel Rodríguez, "Design and Implementation of the NetTraveler Middleware System based on Web Services", in Proc. 2006 IEEE International Conference on Internet and Web Applications and Services, Guadalupe, France, February 23-25, 2006.
7. Edilberto Garcia and Manuel Rodríguez, "WAMDAS: A Web Service-Based Wireless Alarm Monitoring and DataAcquisition System for Pharmaceutical Plants", in Proc. 2006 IEEE International Conference on Internet and Web Applications and Services, Guadalupe, France, February 23-25, 2006.
8. D. Arias and W. Rivera, "Using Grid Computing to Enable Distributed Radar Data Retrieval and Processing" IEEE International Conference on Network Computing and Applications, 2006.
9. W. Lozano and W. Rivera, "An Adaptive Quality of Service Based Scheduling Algorithm for Wide Area Large Scale Problems." IEEE Workshop on Adaptive Grid Computing, 2006.
10. A. H. Diaz-Perez, D. Rodriguez, "One Dimensional Cyclic Convolution Algorithms with Minimal Computational Complexity," IEEE ICASSP 2006, Toulouse, France, May 2006.

# RI: Mastodon: A Large-Memory, High-Throughput Simulation Infrastructure

CISE Research Infrastructure Grant EIA-0303609, Year 3

Risto Miikkulainen, Chandrajit Bajaj, Doug Burger, Kathryn McKinley, Vijaya Ramachandran  
Department of Computer Sciences, The University of Texas at Austin  
risto,bajaj,dburger,mckinley,vlr@cs.utexas.edu  
<http://facilities.cs.utexas.edu/htdoc/documentation/condor/localcondor.html>

## Abstract

In this project, the Department of Computer Sciences at the University of Texas at Austin is assembling an infrastructure for performing large-scale, memory-intensive simulations. Called Mastodon, this infrastructure consists of a large number of rack-mounted Linux/x86 servers, each with multiple processors and several gigabytes of physical memory. Leveraged by industrial and university matching funds, we anticipate scaling Mastodon to over 500 processors and over a terabyte of DRAM over the five years of the grant. The cluster utilizes both distributed batch scheduling using the Condor job management software, as well as soft partitioning for timing and parallel experiments. The cluster is widely utilized by faculty, staff, students, and guests in the department, supporting research in several areas such as computer systems, artificial intelligence, algorithms, and computational approaches to neuroscience, biology, linguistics, and mathematics.

## 1 Introduction

Digital simulation is becoming dominant in not only many fields of physical sciences, but computer science as well. Despite the great improvements in processing speed over the past decades, advanced research on large-scale systems requires simulation engines that are more powerful than the average desktop workstation.

In this project, the Department of Computer Sciences at the University of Texas at Austin is assembling an infrastructure for performing large-scale, memory-intensive simulations. Called *Mastodon*, this infrastructure consists of a large number of rack-mounted Linux/x86 servers, each with multiple processors and several gigabytes of physical memory. Leveraged by industrial and university matching funds, we anticipate scaling Mastodon eventually to over 500 processors and over a terabyte of DRAM. The cluster utilizes both distributed batch scheduling using the Condor job management software, as well as soft partitioning for timing and parallel experiments.

This infrastructure services the needs of a majority of our department's faculty. The projects utilizing Mastodon include the areas of systems (e.g. architecture, networking, and distributed information management), algorithms, machine learning (e.g. reinforcement learning, evolutionary computation, transfer learning, game playing, robotics), and computational approaches to neuroscience, linguistics, biology, and mathematics. In each of these areas, Mastodon makes multiple projects feasible, whereas with conventional computing infrastructure, most of these projects cannot even be attempted.

## 2 The Mastodon Cluster

As of May 2006, the Mastodon cluster consists of 346 processors with 1036 GB of RAM, running Ubuntu Breezy version of Linux. It currently includes the following components (figure 1:




Figure 1: **The Mastodon cluster and its environment in the UT Austin Department of Computer Sciences.** The cluster currently consists of 131 nodes with 346 processors and 1036 GB of RAM, a 1.4TB file server, two checkpoint servers with 1.3TB NetApp checkpoint storage, and a Gigabit switch and ports on all components. Through this switch, Mastodon is connected to the smaller scale Scout cluster as well as a pool of 200 Linux desktops, all managed transparently to the user through the Condor job management system.

- Central manager and file server (morgoth): IBM xSeries 360, 4x Pentium IV Xeon @ 2GHz, 4GB RAM, 1.4TB of file storage, Bullfrog backup system.
- Two checkpoint servers (sauron and gothmog): IBM xSeries 335, 2x Pentium IV Xeon @ 3.06GHz, 4GB RAM.
- Checkpoint storage (udun): Network Appliance FAS940, approximately 1.3TB storage.
- Cluster backbone switch (switch57): Cisco Catalyst 6509 with 240 1-Gbps ports.
- 56 compute nodes (uruk-N): IBM xSeries 335, 2x Pentium IV Xeon @ 3.06GHz, 4GB RAM.
- 16 compute nodes (uruk-hai-N, donated by IBM): IBM xSeries 360, 2x Pentium IV Xeon @ 2GHz, 4GB RAM.
- 35 compute nodes (balrog-N): IBM eServer 326, 2x Opteron 250 2.4GHz (64 bit), 4GB RAM.
- 10 compute nodes (balrog-N): IBM eServer 326, 2x Opteron 250 2.4GHz (64 bit), 16GB RAM.
- 14 compute nodes (lhug-N): HP DL585, 4x dual-core Opteron 2.2GHz (64 bit), 32GB RAM.
- Gigabit Ethernet ports for all cluster components

The Mastodon system is integrated with the rest of the CS department computing resources, including the Scout cluster and about 200 Linux desktops. All these resources are jointly managed through the Condor job management system. A researcher can submit his or her job to Condor and it will run on the most appropriate resource available, depending on load and the job profile. The condor system keeps track of the utilization of each of these three components separately, and also tracks each job based on research area.

We designed the cluster to maximize large-memory job throughput while keeping costs as low as possible. To that end, we initially purchased 2-way industry standard IBM servers, with a 1u rack footprint, to minimize the space and number of racks needed. Each of these machines came with 4GB, allowing a single processor to access all the memory or two processors to access 2GB each. In the second year, 64-bit IBM Opteron servers became available with little extra cost; some of these machines were configured with 16GB of RAM to support jobs with large memory demands.

In the third year, we continued this trend by purchasing several 64-bit Opteron machines with 32GB of RAM, now from HP which provided the best performance/price ratio.

The local disk on the machines holds only the OS and swap, because disk space is more effectively managed under the common file server rather than across hundreds of individual file systems. The dedicated central manager, morgoth, is also the file server. It has about 1.4 TB of storage available for user files, and it also shares the /lusr (local user) filesystem, which includes the single installation of Condor mounted by all compute nodes. The existing department-wide Bullfrog backup system was upgraded and expanded to accommodate the growing cluster. Leveraging existing resources in this manner allowed us to spend less on backup and to buy more compute nodes.

In addition to the file server, a Netapp FAS940 storage device with approximately 1.3 TB of storage is dedicated to checkpointing running processes. It was scaled to accommodate the total projected memory image of the cluster at the time of its completion, but storage can be added (up to 12TB) to accommodate increased per-system memory, if needed. The cluster uses two dedicated checkpoint servers, sauron and gothmog. In the Scout cluster (a smaller cluster predating Mastodon) we found that checkpoint loads could at times bring down the checkpoint server so we wanted to divide the much greater load of the Mastodon cluster over two machines. This configuration reduces the likelihood of overloading the checkpoint server and ensures that at least half of the cluster retains full functionality even in the event that one checkpoint server goes down.

In addition to the local storage within the cluster, all user home and project directories are available to the compute nodes over 1 Gbps ethernet, and the file systems local to Scout are available over 100 Mbps ethernet links. These are automounted by the compute nodes as needed. The Condor jobs are serviced using a dedicated machine, as are the file system requests.

The current 59 64-bit nodes are running the 32-bit version of Ubuntu Linux; they will be upgraded to the 64-bit version as soon as our testing is complete. Future expansion will include more 64-bit AMD Opteron processors with large main memory sizes.

### **3 Research Highlights**

The Mastodon cluster is available to the faculty, staff, students, and guests of the UT Computer Sciences Department under the Condor job management system (which also includes the Scout cluster and the public desktops). During the last year (May 2005 - April 2006), 83 different users have used a total of 981,222 hours of CPU time under Condor, the vast majority of them on Mastodon. In addition, 26 users have had direct access to a subset of nodes in Mastodon, using approximately 105,000 hours of CPU time, mostly for parallel computing and Java applications that Condor does not currently support, and for ntiming runs. Three projects that have made most extensive usage of Mastodon are described briefly below, representing the areas of systems, algorithms, machine learning, and computational biology; other projects are briefly mentioned as well.

#### **3.1 Computer Systems Research**

In the TRIPS project (Desikan et al. 2004), we are building a prototype microprocessor and system based on a new type of instruction-set architecture, called Explicit Data Graph Execution (or EDGE). Each prototype processor chip contains two 16-wide issue processor cores, each of which supports up to 1,024 instructions in flight simultaneously. This design is intended to scale to new levels of single-thread performance while not placing any burdens on the programmer, since the architecture supports compilation of conventional sequential languages such as C and FORTRAN. The prototype chip itself is a 170-million, 330 square millimeter ASIC being manufactured by IBM in their 130 nanometer process technology. The Mastodon cluster was essential for completing this aggressive design. Our team, which consists on 29 cluster users, accumulated over 200,000 simulation hours this year on the Mastodon cluster. These hours were used broadly for three categories of simulations: performance evaluation of the architecture space (tuning the design), measuring the efficacy of compiled code, and verifying the physical design against a series of reference simulators. The verification was by far the most simulation-intensive part, taking place at six different levels of abstraction (sub-unit, unit, tile, processor, cache network, chip, and low-level chip design). Verification runs included randomized tests at all levels, measuring state space coverage, event counter coverage, input variance coverage, and lines of Verilog code hit, to ensure a reasonable coverage of an intractably large space.

Another major focus in systems research is distributed information systems (Jain et al. 2006). The goal of the Scalable Distributed Information Management System (SDIMS) project is to develop a "distributed operating systems

control plane” that will serve as the backbone for a large-scale distributed services. SDIMS aggregates information about large-scale networked systems to provide detailed views of nearby information and rare events and summary views of common, global information. In parallel with SDIMS, we are developing a distributed monitoring system, PRISM, for constructing large-scale data aggregation and continuous event monitoring applications, such as IP traffic monitoring (DDoS attacks), network anomaly detection (Internet worms), accounting and bandwidth provisioning (hot spots, flash crowds), sensor monitoring and control, and grid resource monitoring. At the core of these applications is a distributed query engine that aggregates information and performs continuous tracking of queries over collections of physically-distributed and rapidly-updating data streams. The underlying aim is to provide a global view of information in the system at a reasonable cost and within a specified precision bound. PRISM leverages Distributed Hash Tables (DHT) for scalability and fault-tolerance. PRISM provides four key properties: (1) scales to a large number of streams and query attributes, (b) incurs minimal communication overhead for aggregating query results, (c) is time responsive for quickly identifying anomalies, and (d) bounds the inaccuracy of the computed value for the aggregate function. Our initial results are promising: For a network monitoring application, PRISM is able to identify the global ”worst-offender” flows out of millions of ongoing flows at a modest communication overhead. The primary task for which the cluster proved highly useful was to set up a distributed system which can be easily controlled. This distributed system was then used to perform a scalable aggregation on a cluster of about 50 nodes. We then used this system to simulate different scenarios and compare the performance of different policies in SDIMS system.

Other computer systems projects include those in networking (bandwidth allocation in wireless networking, network measurement, traffic engineering, routing, and channel assignment), memory allocation, and compiler-generated checkpointing.

### **3.2 Algorithms Research**

In the cache-oblivious algorithms project (Chowdhury and Ramachandran 2006), we experimentally evaluate running time and cache performance of algorithms and data structures we have designed for various graph and dynamic programming problems on massive data sets. We have already obtained timing and caching data for several of our new cache-oblivious dynamic programming algorithms (longest common subsequence and the cache-oblivious Gaussian elimination paradigm) that have many practical applications including sequence alignment in bioinformatics, and also for a new cache-oblivious priority queue data structure we developed called the ”buffer heap” that has applications in shortest paths computations. In our experiments all our algorithms and data structures performed significantly better than traditional ones for large inputs.

Other algorithms projects include developing an analytic model of the bandwidth cost in ray tracing and evaluating random graph theory in distributed sensor networks.

### **3.3 Machine Learning Research**

Funded by the DARPA program on ”Transfer Learning,” we study how learning knowledge in a source task can be transferred to a related target task to improve learning performance in the target task (Kalyanakrishnan et al. 2006). In simulated RoboCup soccer domain, we obtained results showing transfer learning is possible and effective using several related but distinct methods. In particular, we developed a transfer learning algorithm for a new probabilistic graphical model, Markov Logic Networks (MLNs), that combines first-order logic with probability. The algorithm works by successfully diagnosing previous knowledge and relearning only the portions of the MLN that are incorrect for the new domain. This process significantly accelerates learning while maintaining the accuracy of the current state-of-the-art MLN learning algorithm. We need the cluster to perform a significant number (20-50) of independent parallel experiments to obtain statistically significant learning results.

Another significant focus of our machine learning research has been evolutionary algorithms, in particular those that evolve neural networks in game playing domains (Stanley et al. 2005). In the visible intelligence project, we evolve neural networks as controllers for intelligent agents embedded in games and simulators. The goal is to obtain visibly intelligent behavior, i.e. behavior that conforms to a recognizable doctrine rather than to some abstract optimization. To date we have been able to produce agents that organize a division of labor in situ when deployed for a task, and have been able to refine that behavior by using human-generated examples to guide evolution to solutions similar to the style of play used by human players when generating the examples. This technology will not only provide intelligent agents that make commercial games more engaging, but will also be useful for simulators and training

games, since it will make it easy for subject matter experts to train embedded agents by example, bypassing the need for expensive analysis, specification, and coding of the required patterns of behavior in the agents.

Other machine learning projects include kernel-based reinforcement learning, evolving a neural-network value function approximator, evolving sequence predictors, developing evolvable representations, coevolutionary memory, developing a neuroevolution interface to TIELT gaming middleware, robot arm control, reinforcement learning in high-diameter environments, information extraction of biomedical texts, classification through random projections, and electronic auctions.

### 3.4 Computational Approaches to Neuroscience, Linguistics, Biology, and Mathematics

In the LISSOM project (Miikkulainen et al. 2005), we are simulating the development and interaction of different areas in the visual cortex. The main focus is to understand the role of feedback connections (i.e. from a higher level area to a lower level area), which is still a mystery, using a computational model. Preliminary results from the simulation show that meaningful connections within and between different areas in the visual cortex may be formed by following a single principle. The model also predicted the signature of feedback signal in biological data. The prediction led to a biological experiment, in which the signature was found, demonstrating for the first time how feedback activation could be observed directly in voltage-sensitive dye imaging. The large memory and fast processing speed of the machines in Mastodon are essential to my simulations because each run needs about 2GB of memory and runs for more than 10 hours.

Mastodon has also been used to develop and test the computational neuroscience theory that complex brain structures self-organize prenatally, driven by internally generated activation patterns. In other areas, it is used in building neural network representations of narratives, in model schizophrenic story processing, in phylogenetic tree reconstruction, and in automated theorem proving.

## 4 Conclusion

Mastodon cluster has already had an enormous impact on research in several levels, and we expect it to continue to do so for several more years. At the departmental level, it serves the majority of our faculty, allowing them to make significant leaps in their research. It creates synergy among complimentary projects, permitting them to tackle more ambitious goals with larger, synthesized simulated systems. Its powerful capabilities allows it to serve many different communities: researchers who need a dedicated partition of the machine, students in large classes running jobs as part of their undergraduate education, researchers elsewhere in campus who need to use large-memory platforms, and collaborating researchers in other departments and institutions. In many cases, Mastodon makes possible projects that could not be completed with conventional computing infrastructure.

## References

- R. Desikan, S. Sethumadhavan, D.C. Burger, and S.W. Keckler (2004). "Scalable Selective Re-execution for Speculative Dataflow Architectures." *11th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- N. Jain, D. Kit, P. Mahajan, P. Yalagandula, M. Dahlin, and Y. Zhang (2006). PRISM: Precision-aware Aggregation for Scalable Monitoring. Technical Report TR-06-22, UT Austin Department of Computer Sciences.
- R. A. Chowdhury and V. Ramachandran (2006). "Cache-Oblivious Dynamic Programming." *Proceedings of the 17th Annual ACM-SIAM Symposium of Discrete Algorithms (SODA 2006)*, Miami, Florida, pp. 591-600.
- S. Kalyanakrishnan, Y. Liu, and P. Stone (2006). Half Field Offense in RoboCup Soccer: A Multiagent Reinforcement Learning Case Study. *The 2006 Robo Cup International Symposium*
- K. Stanley, B. D. Bryant, and R. Miikkulainen. (2005). The NERO Real-time Video Game. *IEEE Transactions on Evolutionary Computation*, 9:653-668.
- R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh. (2005). *Computational Maps in the Visual Cortex*. New York: Springer.

# MII: UTEP Distributed Computing Lab: Creation of a Pipeline for Minority Participation in Research

Project EIA-0325024 Report, 09/01/04 – 05/31/05

Research, Education and Outreach Activities

Patricia A. Nava, David H. Williams, and Joseph Pierluissi  
([pnav@ece.utep.edu](mailto:pnav@ece.utep.edu), [williams@ece.utep.edu](mailto:williams@ece.utep.edu), [pier@utep.edu](mailto:pier@utep.edu))

Department of Electrical & Computer Engineering

University of Texas at El Paso

El Paso, TX 79968-0518

<http://academics.utep.edu/Default.aspx?alias=academics.utep.edu/ee>

## 1. Abstract

*The UTEP Distributed Computing Lab (DCL) project, in addition to enriching the curriculum and facilitating interdisciplinary research, has assisted in the recruitment and support of underrepresented minorities, with the objective of creating a pipeline of students from high school through graduate school that are trained in the design, operation and applications of high-performance computers. The 40-node student-built cluster is the hub of the DCL. The focus of this document is to summarize accomplishments to date, which include: (1) fine-tuning of the cluster, which has resulted in 139 GFLOPS performance speed, and which serves as the hub for educational, research and outreach activities; (2) migration of the pilot research projects onto the cluster and evaluation of performance; (3) creation of cluster-based experiences for students in the curriculum; (4) engagement of graduate and undergraduate students in the DCL; and (5) conducting outreach to middle- and high-school students, particularly underrepresented minorities and women, to inspire their interest in CISE. Work to date has exceeded the objectives, with advancement of the research, new collaborative efforts and the production of four student posters and 25 technical papers.*

## 2. Introduction

The University of Texas at El Paso (UTEP) is a major regional university serving a rapidly growing, bi-national, bi-cultural population on the U.S.-Mexico border. The U.S.-Mexico border is a few hundred yards from the UTEP campus, creating a border community of around 2 million people. As the only major state university in Texas within 350 miles of El Paso, UTEP is central to providing educational opportunities and fostering the human and economic development of the region. El Paso is at the forefront of demographic trends that are rapidly changing the face of American society. The population of El Paso is estimated to be over 70% Hispanic. El Pasoans have a median age of 25, compared to a median age of 28 years for the state as a whole. Since the younger age brackets are more heavily Hispanic, public schools in El Paso County range from 75% to 98% Hispanic.

The PIs from UTEP sought to establish a Distributed Computing Lab (DCL), which would serve as a hub for educational, research and outreach activities that create a pipeline of minority and women students from high school through graduate school. UTEP has a long and successful record of educating minority students, making it the top producer of Hispanic graduates in engineering in the continental U.S. 69% of UTEP students are Hispanic and an estimated 60% are the first in their families to attend college. This program involves the recruitment and support of underrepresented minorities with the objective of creating a pipeline of students from high school through graduate school that are trained in the design, operation and applications of high-performance computers.

## 3. Goals, Objectives, and Targeted Activities

The major goal of the project is to create a pipeline of students in CISE, and create more professionals experienced in distributed computing, architecture and applications. Specifically, it was proposed to train students, and have them construct a *32-CPU Cluster System*, and utilize it in three interleaved efforts: (1) to perform specialized outreach to intelligent minority and female high-school students labeled “at risk;” (2) to provide enrichment of educational activities in CISE; and (3) to involve students in research and use the DCL as a catalyst to promote multi-disciplinary collaborations that would benefit from this form of computing. The project has these three objectives, each of which contributes to creating a pipeline of minority students that creates a larger number of motivated entering students, leads to higher retention and attainment of degrees, and provides qualified professionals experienced in CISE research activities.

The targeted activities stem from affording students the opportunity to carry out hands-on projects in courses. Additional targeted activities include multiple research projects, in particular: (1) neuro-fuzzy systems modeling, simulation, and applications, (2) cluster architecture and networking, (3) general distributed computation techniques, and (4) enhancement of electrocardiography.

#### 4. Infrastructure and Facilities

**Physical Infrastructure.** The Distributed Computing Lab (DCL) consists of a student-designed, student-built, student-maintained 40-CPU *Beowulf Computer Cluster*. The cluster is housed in the renovated CPU room of the ECE-UNIX Laboratory. This room was enlarged, and renovated, with university cost-share funds. The room is secured by a deadbolt lock and pushbutton code entry, resulting in the ability to ensure that access to the physical machine is limited to DCL-related faculty and students.

**Educational and Research Infrastructure.** A very important part of the infrastructure built by this project deals with the human capital developed by the DCL. There are three facets to this infrastructure; (1) the development of professional skills by the DCL student members; (2) the exposure of UTEP students to distributed processing through the enrichment of courses; and (3) the development of research collaborations that have resulted and expanded because of the DCL. On the first facet, students participate in DCL group meetings, where they are required to make presentations, developing their presentation and public speaking skills. The students work with faculty, so that each member has the opportunity to develop their writing skills as well. Twenty-three of the twenty-five technical papers to date have at least one student author, and most conference papers have been presented by students. Additionally, we have graduated one Ph.D., 12 M.S. and 11 B.S.-level students (of these, two undergraduates have transitioned to the Master's program, and six Master's level students have transitioned to the Ph.D. program.) On the second facet, due to cluster exposure during coursework, we have identified more students interested in the DCL and distributed computing. On the third facet, the PIs have developed research collaborations with cardiologists, biologists, computer scientists, business researchers and biomedical engineers. These collaborations have led to proposals in related areas, and to other funding opportunities.

#### 5. Project Description

The project consists of constructing a 40-CPU *Beowulf Computer Cluster* (originally described as a 32-CPU, but upgraded to 40-CPU's upon its initial design) and its utilization in three interleaved goals: (1) to provide enrichment of educational activities in CISE; (2) to perform outreach to minority and female high-school students; and (3) to involve students in research and use the DCL as a catalyst to promote multi-disciplinary research collaborations that can benefit from this form of computing.

#### 6. Indicators of Success and Major Accomplishments

The indicators of success and accomplishments of the project are described in terms of three categories. The first category, research advances, are addressed in this section. The second success indicator is student impact, which is addressed in section 8, "Outreach and Supported Students." The third indicator of success is scholarly work, collaborative efforts and leveraged funding, and is reported in section 7.

**Systems Research:** The "student-constructed supercomputer" was such a novelty for UTEP, it was featured [1] in the campus newspaper. A follow-up article [2] entitled "Supercomputer Nears Completion," detailed the status of the cluster and plans for its use in research and education. Nito Gumataotao, a graduate student within the ECE Department at UTEP, was primary architect of the final design. Although the initial design consisted of a 32-CPU Beowulf cluster [3], an equipment grant of three 3750 (24-port gigabit Ethernet communications) switches by Cisco Systems allowed the leveraging of funds in order to expand the design to a 40-CPU system. Nito Gumataotao directed the construction of the cluster (a time-lapse movie of a build session is available for viewing at <http://www.ece.utep.edu/research/webdcl/DCL-2/Templates/index.html>) and currently administers its operation. The cluster employs Red Hat Linux [4], with ROCKS [5]. Monitoring is provided through "ganglia" which takes snapshots of node operation at user defined intervals.

System performance was tested using High Performance Linpack (HPL), the standard benchmark for high performance computers [6]. With various kernel configurations and different settings for HPL parameters, the performance has risen to 139 GFLOPS, using all compute nodes.

**Cluster Performance Research:** One of the research projects is neuro-fuzzy systems modeling, simulation, and applications. The portion associated with performance is reported here. A study on the efficiency of distributing the code for a neural network simulation (using three standard benchmark problems) was reported in [7], and had a speed-up of 6 to 8 when using 10 processors (in comparison to 1 processor). These results were reported in [8]. Additionally, studies on network topologies were performed and are slated for reporting in [9].

**Neuro-Fuzzy Simulation:** The neuro-fuzzy research has spawned four lines of investigation: (1) fuzzy inference neural networks in conjunction with genetic algorithms [10, 11], and their application to standard classification problems [8, 24]; (2) use of neural networks to extract knowledge for use in fuzzy systems [12], and their application to sleep classification [13, 14]; (3) use of neural networks to classify respiratory diseases [15, 16, 17, 18, 19, 20, 21, 22, 23]; and (4) use of neuro-fuzzy methods to classify human sleep via features extracted solely from an EEG signal [25, 26, 27, 28]. Collaborations involve DCL PIs, other UTEP faculty, and local physicians. These lines of investigation have spawned other questions worthy of research, which will be funded (as reported in section 7).

**Research on Distributed Solution of Lineality (or Linearity) Space Problems:** This project investigated a distributed approach to solve *lineality* (or *linearity*) *space* (**LS**) problems using the cluster. A distributed approach is useful for these computations because **LS** problems have large cardinalities and a large number of dimensions. Substantial speedups can be obtained by distributing the processing to multiple compute nodes [29]. The **LS** solution has many important applications in engineering, science, and business, and includes a subset of solutions of the more general *linear complementarity problem* (**LCP**). A partial list includes convex quadratic programming solutions for the contact problem, the porous flow problem, the journal bearing problem, and free boundary problems, among others. Other applications of **LCPs** include the solution of equilibrium problems for market, capital stock, or traffic equilibrium, the Nash equilibrium problem and the n-person Nash-Cournot equilibrium problem.

**Enhancement of Spatial Resolution in Electrocardiography:** This pilot project has generated several distinct research questions that are interleaved: (1) digitization of recorded electrocardiograms; (2) validated use of measured 12-Lead EKGs in the simulation of body surface potentials; and (3) an improved method for synchronization of electrocardiogram printouts. These investigations have led to collaborations with a biomedical engineer, a signal processing specialist, and a local cardiologist. A more detailed discussion of each of these studies is provided below.

**Digital Processing of Recorded Electrocardiograms:** Digitization of EKG Printout software has been developed to convert an electrocardiogram into digital form, which facilitates a detailed analysis of properties of the EKG signals. Furthermore, it allows a digital vectorcardiogram to be obtained and used in simulation of body surface potentials. This research was presented as at the 22nd Conference on Biomedical Engineering Research [30].

**Validated Use of Measured 12-Lead EKGs in the Simulation of Body Surface Potentials:** This subproject was initiated to validate the use of measured 12-Lead EKG in the simulation of body surface potentials. In order to simulate electrical activity of the human heart, the 12-Lead EKG in vector form, or vectorcardiogram, is needed. The technique used to obtain the vectorcardiogram was the Inverse Dower technique [31]. In addition, Frank's XYZ transfer coefficient matrix for deriving the 12-Lead EKG [32] was used to validate the vectorcardiogram produced by the Inverse Dower technique. The synchronization method implemented in this research is synchronization by cross-correlation of an unsynchronized EKG with a synchronized EKG. The synchronized 12-Lead EKG can then be used to obtain the vectorcardiogram required for simulation of body surface potentials. Validation of body surface potentials is achieved by comparing the measured precordial leads and the simulated precordial leads. This study was reported in [33], and has spawned further investigation, as reported in [34].

**Digitization and Synchronization Method for Electrocardiogram Printouts:** This research is an improvement of the Digital Processing of Recorded Electrocardiograms study published in [30]. This improved software is capable of synchronizing 24 EKG signals automatically and includes an improved method of synchronization than that presented in [33]. This study has been published [35], and has spawned several collaborations on signal processing of the EKG [34, 35, 36].

## 7. Publications, Proposals, and Leveraged Funding

The DCL has provided a venue for collaboration and publication efforts that have been very effective. Specifically, this work has led to four journal submissions, the publication of 15 conference papers, two conference posters, five Master's theses, and a doctoral dissertation. The grant has also provided the PIs and participating students the opportunity to collaborate with three physicians, two biomedical engineers, a biologist and a sleep researcher. This collaboration has led to the submission of proposals to various agencies, as shown in the bulleted list below. In addition, the grant has been leveraged with funding from other sources. For example: (1) DCL will receive one graduate student from the following grant: National Institute of Environmental Health Sciences – Advanced Research Cooperation for Environmental Health (NIEHS–ARCH) Subproposal, “Impulse Oscillometric Evaluation of the Effect of Air Quality on Respiratory Function in Normal and Asthmatic Anglo and Hispanic Children at the Border,” PIs: Nazeran (Lead), Nava, Diong; (2) “SHARE, the Southwest Center for the Advancement of Cyberinfrastructure in Research and Education,” submitted to NSF-CREST, PIs: Gates (Lead),

Velasco, Tweedie, Kreinovich, Chiu, Nava, Velasquez. (decision pending); (3) “Request for Three Cisco Gigabit Switches for a 32 CPU Beowulf Cluster,” Cisco Systems, Inc. (Higher Education Equipment Grant), PIs: Williams (Lead), Nava, Pierluissi, 2/1/04-1/31/05, \$22,305 (awarded); (4) Four students funded by other sources

## 8. Outreach Programs and Supported Students

**Supported Students:** The DCL group meets once a week. The meetings are used as a venue for discussion of lab issues, tutorials on professional skills, exercising professional skills, tutorials on cluster operation, team building activities, and schedule and management planning for outreach activities. The numbers of students that have been participating in the DCL group are detailed in Table 1: DCL Student Summary.

**Table 1: DCL Student Summary**

Academic Year	Total	Sponsored Students*	Gender		Under-rep. Minorities		Students by Degree			Conference Participation		# Students Participating in Outreach Activities
			M	F	M	F	UG	MS	PhD	Papers**	Posters	
2003-2004	24	19	15	9	12	5	8	15	1	2	1	12
2004-2005	29	21	18	11	16	11	15	11	3	13	2	16
2005-2006	32	26	21	11	19	11	14	14	4	10	1	4***

NOTES: \* – Number of students sponsored, directly or with leveraged funding.

\*\* – Student co-author, preparation of slides and/or conference presentation.

\*\*\* – Does not include 2006 summer camps, which will take place in July 2006.

**Outreach (ExciTES Summer Institute):** The 2006 Summer Camps will take place in July. We have projected the involvement of **100** local middle- to high-school students for this year. The ExciTES Summer Institute activities consist of a DCL module (which exposes the campers to CISE careers, allows them a “hands-on” experience in building a PC, and rewards them by playing a video game upon successful construction), as well and other science and engineering modules.

In summary: **300** students have participated in the camps (first two years); Questionnaires show high level of interest in DCL module; 22 DCL students served as role models, mentors, and advisors to pre-college students; DCL students participated in developing materials and delivering the modules; and Campers’ comments, as well as assessment results, indicate the module is effective at raising levels of awareness and interest in CISE majors and careers.

**Outreach (UTEP Engineering EXPO and high school tours):** Operation of the Virgo cluster is demonstrated to students and their teachers during the Engineering EXPO program. The groups are from various high schools in the El Paso metropolitan area. Emphasis is placed on the fact that Virgo was constructed and is operated by students. DCL students hosted a total of **165** students from local schools such as Chapin and Parkland.

**Outreach (Chapin High School Engineering Week):** Chapin High School is an engineering magnet school, and has collaborated with UTEP to celebrate Engineering Week by hosting speakers and presentations about engineering. Six DCL students participated in three different sessions. Approximately **120** students attended these presentations.

## 9. Future and Strategic Directions

The project’s outcomes (as described in sections 6 through 8), which fit with our objectives, would not have been possible without the grant. Indeed, all three explicit goals listed dovetail to enhance the overarching goal of creating a pipeline, from high-school to graduate school, of students interested in CISE studies. The award has already made possible at least several breakthroughs that would not have otherwise been practical: establishment of a distributed computing lab; increasing awareness of CISE studies and careers amongst pre-college students;

increasing the number of college students interested in distributed computing; and providing incentive for faculty with other expertise to join a collaborative research effort.

The future directions for this project are: (1) to obtain funding for the collaborative research being carried out on the cluster; (2) to continue to mentor students through the pipeline; and (3) find funding for the summer camp outreach activities from industry.

## References

NOTE: student authors designated with an \*.

- [1] Trejo, Erica, "Student[ysic] Build Supercomputer from Scrach[ysic]," *The Prospector*, UTEP, April 14, 2004, p 6.
- [2] Hall, John, "Supercomputer nears completion," *The Prospector*, University of Texas at El Paso, October 6, 2004, p 8.
- [3] <http://www.beowulf.org>
- [4] <http://www.redhat.org>
- [5] <http://rocks.npaci.edu/Rocks/>
- [6] [www.top500.org](http://www.top500.org)
- [7] Sagarnaga\*, Miguel A., David H. Williams, and Patricia A. Nava, "Distributed Processing of a Neural Network Application", *Proc 2004 Int'l Conf on Parallel and Distributed Processing Techniques and Applications*, pp. 119-124, 2004.
- [8] Cruz-Cano\*, Raul, Patricia A. Nava, and David H. Williams, "A Parallel Genetic Algorithm for Fuzzy Inference Neural Networks," *Proc. 2005 Int'l Conf on Parallel and Distributed Processing Tech. and Applications, (PDPTA '05)*, July 2005.
- [9] Zapata\*, S.N., D.H. Williams, and P.A.Nava, "Analysis of Cluster Interconnection Network Topologies," accepted for *Proc. 2006 Int'l Conf. on Parallel and Distributed Processing Techniques and Applications (PDPTA '06)*, July 2006.
- [10] Cruz-Cano\*, R., Creation and Optimization of Fuzzy Inference Neural Networks, Ph.D. Dissertation, UTEP, May 2005.
- [11] Cruz-Cano\*, R., P.A. Nava, and D.H. Williams, "Fuzzy Inference Neural Networks: A Novel and Efficient Method for Optimization," *Journal of Intelligent and Fuzzy Systems*, in review, 2005.
- [12] MartinDelCampo\*, E., Design of Fuzzy Systems using Knowledge Extracted via Neural Networks, MSEE Thesis, UTEP, August 2004.
- [13] Cruz-Cano\*, R., E. MartinDelCampo\*, P.A. Nava, and R. Cabeza, "Fuzzy Rule Extraction and Optimization for Rat Sleep-Stage Classification," *Proc 2005 Int'l Conf on Parallel and Distributed Proc. Tech. and Appl. (PDPTA '05)*, June 2005.
- [14] Varghese\*, C., Sleep Classification using Neuro-Fuzzy Intelligence, MSEE Thesis, UTEP, December 2005.
- [15] Barua\*, M., Classification of Pulmonary Diseases Using Artificial Neural Networks, MSEE Thesis, UTEP, December 2004.
- [16] Barua\*, M., H. Nazeran, P. Nava, and V. Granda\*, "Classification of Pulmonary Diseases using Artificial Neural Networks," *Proceedings Of ANNIE 2004*, November 2004.
- [17] Baswa\*, S., Evaluation of Respiratory System Models Based on Parameter Estimates from Impulse Oscillometry Data, MSEE Thesis, University of Texas at El Paso, December 2005.
- [18] Diong, B., N. Nazeran, P. Nava, and M. Goldman, "Analysis of the Extended RIC Model of Human Respiratory Impedance," *IEEE Engineering in Medicine and Biology*, slated for November 2006 issue.
- [19] Baswa\*, S., B. Diong, H. Nazeran, P. Nava, and M. Goldman, "Evaluation of Respiratory Systems Models Based on Parameter Estimates from Impulse Oscillometry Data," *Proc. IEEE-Engr in Medicine and Biol Soc (EMBS), 27th Annual International Conference*, Sept 2005.
- [20] Barua\*, M., H. Nazeran, P. Nava, B. Diong, and M. Goldman, "Classification of Impulse Oscillometric Patterns of Lung Function in Asthmatic Children Using Artificial Neural Networks," *Proc. IEEE-Engr in Medicine and Biol Soc (EMBS), 27th Int'l Conf*, Sept 2005.
- [21] Barua\*, M., H. Nazeran, P. Nava, B. Diong, and M. Goldman, "Implementation of Artificial Neural Networks to Classify Impulse Oscillometric Patterns of Lung Function in Asthmatic Children," *Proc. of the 5th Int'l Workshop on Biomedical Signal Interpretation*, Sept 2005.
- [22] Barua\*, M., H. Nazeran, P. Nava, V. Granda\*, and K. Behbehani, "Classification of Pulmonary Diseases by an Artificial Neural Network, Based on Measurements from the Impulse Oscillometry System," *Proc. Of the 26th Int'l Conf-IEEE Engr in Medicine and Biol Soc*, 2004.
- [23] Woo\*, T., B. Diong, L. Mansfield, H. Nazeran, M. Goldman, and P. Nava, "A Comparison of Various Respiratory System Models Based on Parameter Estimates From Impulse Oscillometry Data," *Proc. Of the 26th Int'l Conf-IEEE Engr in Medicine and Biology Society*, 2004.
- [24] Joshi\*, K., S. De La Cruz\*, B. Diong, D. Williams and P. Nava, "A Neural Network Implementation of Dual-frequency Output Control for Multilevel Inverters," *Proc. Of the 2004 ComPEL Conference*, 2004.
- [25] Barragan\*, J., Neuro-Fuzzy Sleep Classification using Features Extracted from EEG Signals, MSEE Thesis, April 2005.
- [26] MartinDelCampo\*, E., E. Estrada\*, J. Barragan\*, P. Nava, and H. Nazeran, "Automatic Sleep Stage Classification Utilizing Fuzzy Logic and Neural Network with the EEG Signal," *Proc. Of ANNIE 2004*, November 2004.
- [27] Barragan\*, J., E. Estrada\*, P. Nava, and H. Nazeran, "EEG-based Classification of Sleep Stages Using Artificial Neural Networks," *Proc of the 5th International Workshop on Biomedical Signal Interpretation*, Sept 2005.
- [28] Estrada\*, E., H. Nazeran, P. Nava, K. Behbehani, J. Burk, and E. Lucas, "Itakura Distance: A Useful Similarity Measure between EEG and EOG Signals in Computer-aided Classification of Sleep Stages," *Proc 26th Int'l Conf of the IEEE Engr in Medicine and Biol Society*, 2005.
- [29] Caire\*, M. E., F. J. Lopez, and D. H. Williams, "Distributed Identification of the Lineality Space of a Cone", *International Journal of Computers and Applications*, in review, 2005.
- [30] Sevilla\*, D., E. Morales\*, J.H. Pierluissi and Z. Abedin, "Digital Processing of Recorded Electrocardiograms," *Proceedings of the 22nd Annual Houston Conference on Biomedical Engineering Research*, 2005.
- [31] Edenbrandt, Lars and Olle Pahlm, "Vectorcardiogram synthesized from a 12-lead ECG: superiority of the inverse dower matrix," *Journal of Electrocardiology*, vol. 21, issue 4, pp. 361-367, 1988.
- [32] Frank, E., "An accurate, clinically practical system for spatial vectorcardiography," *Circulation*, vol.13, pp.737-749, 1956.
- [33] Morales\*, E.,D. Sevilla\*, J.H. Pierluissi, and Z. Abedin, "Validated Use of Measured 12-lead EKGs in the Simulation of Body Surface Potential", *Proceedings of the 22nd Annual Houston Conference on Biomedical Engineering Research*, 2005.
- [34] Pierluissi, J.H., E. Morales\*, D. Sevilla\*, R.F. von Borries, and Z. Abedin, "Invertible Algorithm for Patient Specific Electrocardiogram to Vectorcardiogram Transformation," *International Journal of Bioelectromagnetics*, in review, 2005.
- [35] Morales\*, E., D. Sevilla\*, J.H. Pierluissi, and Z. Abedin, "Digitalization and Synchronization Method for Electrocardiogram Printouts", *Proc 27th Annual Int'l Conf of the IEEE Engineering in Medicine and Biology Society*, 2005.
- [36] von Borries, R., J.H. Pierluissi, and H. Nazeran, "Wavelet Transform-based ECG Baseline Drift Removal for Body Surface Potential Mapping", *Proc 27th Int'l Conf of the IEEE Engineering in Medicine and Biology Society*, 2005.

# CISE MII: Research Experience for Minority Students in High Performance Computing and Communications

Award # 0117255

**Principal Investigators: Kleanthis Psarris and Rajendra Boppana**

Department of Computer Science  
The University of Texas at San Antonio  
San Antonio, TX 78249  
{psarris, [boppana](mailto:boppana@cs.utsa.edu)}@cs.utsa.edu

## 1. Abstract

The MII program is oriented towards academic achievement. We have recruited many talented full time students from under-represented groups in the program who participated in research projects under the supervision of their faculty mentors. The selection criteria included a high GPA, outstanding letters of recommendation and a commitment to pursue a thesis or project option for their degree. We enhanced the educational experiences of our students and increased the number of graduate degrees. We also improved the research productivity of the principal investigators and made several research contributions in the field of high performance computing and communications. We have been very successful in producing many minority students with BS and MS degrees who joined the high-tech industry workforce. We have also been successful in increasing the number of minorities who attend graduate school.

## 2. Research Activities

### *2.1. Program Analysis and Optimization for High Performance Execution*

Scientific source code for high performance computers is extremely complex containing irregular control structures with complicated expressions. This complexity makes it difficult for compilers to analyze the code and perform optimizations. In particular with regard to program parallelization, complex expressions are often not taken into consideration during the data dependence analysis phase. The problem results from non-linear expressions that appear in array subscripts and loop bounds, from loop regions with trapezoidal or symbolic bounds, from multi-dimensional arrays with coupled subscripts, and from if-statement constraints. Most data dependence tests have been designed to analyze dependence problems with linear constraints only, even though non-linear expressions are very common in practice. Therefore considerable amounts of potential parallelism remain unexploited. In this research we propose new efficient data dependence analysis techniques to handle such complex instances of the dependence problem and increase program parallelization.

### *2.2. Ad Hoc Wireless Network Security*

Ad hoc wireless networks are networks formed among mobile, wireless nodes without any infrastructure support. These networks provide multihop routing by letting each node seek and disseminate routing information. These are more flexible, but performance challenged than the widely used one-hop wireless networks such as cellular and WiFi hotspots. Previously, we investigated various design and performance aspects of ad hoc networks: adaptive routing protocols, analyzing and improving the IEEE standard 802.11 medium access control protocol, tuning transport protocols for better performance on ad hoc networks, and ad hoc networks with limited infrastructure support to improve reliability and performance of multihop routes.

In the current research, we investigated vulnerabilities and security threats to MANETs. We investigated denial of service (DoS) attack in which malicious nodes attempt to overload MANET with a large number of useless messages to cause congestion. In MANETs, these DoS attacks are highly leveraged owing to the fact that all nodes gather and disseminate routing information. For an  $n$  node network, a single message generated by a malicious node can result in up to  $n$  transmissions in the network. The other types of attacks are based on malicious nodes generating and propagating false routes, which causes data to be sent to them for forwarding them to appropriate destinations; but malicious nodes can drop these data packets and cause dramatic performance loss. Another research issue we addressed is energy drain caused by these attacks.

### 3. Research Results

#### *3.1. Program Analysis and Optimization for High Performance Execution*

We developed a dependence analysis tool based on a set of polynomial time techniques that can conclusively prove or disprove dependences in the presence of non-linear expressions, complex loop bounds, arrays with coupled subscripts, and if-statement constraints. In addition it can produce accurate and complete direction vector information enabling the compiler to apply further transformations. We implemented the dependence analyzer in the PLATO library, which can be ported in any optimizing and parallelizing compiler. To validate the proposed techniques we performed an experimental evaluation and comparison against the I-Test, the Omega test and the Range test, which represent the current state of the art in data dependence analysis. For our experiments we integrated the PLATO library into the Polaris compiler and we used the Perfect benchmarks, the SPEC benchmarks and the Lapack library. The experimental results indicate that our dependence analysis tool is accurate, efficient and more effective in program parallelization than the other dependence tests. The improved parallelization of key loops resulted into higher speedups and better program execution performance in several benchmarks.

#### *3.2. Ad Hoc Wireless Network Security*

The impact of DoS service attacks can be devastating on the ad hoc network performance when an on demand routing protocol such as AODV is used. Since each control packet transmitted by a malicious node can result in  $n$  repeat transmissions of the same in an  $n$ -node network, even a low-rate attack by one or more malicious nodes can render the network unusable. Using simulations, we showed that with 2 control packet transmissions/second by a single malicious node in a 100-node network can reduce peak throughput by 25%. With higher attack rate, the network is unusable except for very low traffic loads. It is hard to detect low-rate attacks since even a normal node may generate several control packets in a second to repair a broken route. We developed a statistical profiling technique to aggregate the control packet rates of all nodes over time. This can be used to identify malicious nodes from normal nodes regardless of current control packet activity. This technique can also be used to handle consistent false route propagation by malicious nodes in blackhole and other types of attacks. However, this technique cannot address the so called stealth attacks in which malicious nodes propagate false routes only for a small fraction of time. Our simulation analysis indicates that such attacks can be highly effective on certain routing protocols such as AODV; five nodes, each sending false routes 1% of the time, can reduce network performance by as much as 30%.

## 4. Education Activities

### *4.1. Improve the curriculum and the teaching*

We have worked on a standardized curriculum for all sections of the introductory courses. The curriculum carefully integrates the laboratory material with the lectures. All of the lectures and laboratories are offered in computer classrooms. Students are provided with many hands-on activities during the course of the lectures and laboratories. All of the lectures and laboratory materials are on the web. All sections of the first two Computer Science courses for majors are taught by tenure-track faculty. We have also implemented small laboratory and recitation sessions to accompany each lecture, with typically 12-15 students in each section, where the students can get individual attention.

### *4.2. Peer tutoring and other support to help students feel part of a community*

Peer support provides a learning community in which students can collaborate. Several studies have reported that a collaborative environment with peer support is crucial to attracting women and minorities to computer science. When students have software problems, account problems or program bugs, they need help when the problem occurs. A delay of hours or days to see course instructors substantially decreases productivity and increases frustration. With tutors providing lab assistance and handling routine technical problems, course instructors can concentrate on concepts and improve the overall quality of instruction. Each semester we select our top undergraduate students including the MII scholars to act as peer tutors for the freshman and sophomore level courses.

### *4.3. Advising and career services support*

We have provided advising and career services support so that students begin to understand where they are going and how to get there. An academic advising session has replaced one of the 10 required course laboratories in the first introductory Computer Science course. The academic advisors produce individualized programs of study for each student in the class. We have nearly 100% participation and regard this initiative as a success. We have also implemented a second milestone in the second Computer Science course. As one of the 10 required course laboratories, a Career Services representative registered the students with career services and gave a presentation about career paths and internship opportunities.

## 5. Education Results

We achieved a significant increase in minority student participation in computer science. Eight Masters and twenty four Bachelors degree students from under-represented minority groups in Computer Science, supported with scholarships from this award, have already completed their degrees since the inception of the MII program in the Fall of 2001. The majority of the graduates have secured positions and pursued careers in high tech industry and many of them went on to graduate school and research careers.

The MII program had a very positive effect in the overall retention and graduation rates of minorities in Computer Science. The percentage of Hispanics with BS degrees who graduated in 2004-05 was 41% compared with only 13.5% in 2000-01. It is clear that the MII program has been a tremendous success at the undergraduate level. The percentage of Hispanics with MS degrees who graduated in 2004-05 was 13.6% compared with only 7.1% in 2000-01. The pool of Hispanic graduate students has substantially increased in the last couple of years and it is expected that the number of Hispanic MS degrees will further increase in the near future. The MII program has also been a success at the graduate level, which is our current focus with the

expectation of an even greater success. We have achieved this success at a time when undergraduate and graduate enrollments in Computer Science are significantly declining nationwide.

## **6. Training and Development:**

### *6.1. Student Research Training*

Training and development activities on the MII Program included research training for graduate and undergraduate students. Eight PhD and one Masters students directly participated and conducted research in the Principal Investigators' research projects. Several other graduate and undergraduate students from underrepresented groups in Computer Science have worked with faculty mentors on various research projects with support from this award.

### *6.2. Dual-Mentor Partnership Program*

Training and development activities on the MII Program included internships for students interested in participating in industry-based research projects. These internships are formally organized under the MII Dual-Mentor Partnership Program in collaboration with the Southwest Research Institute (SWRI) in San Antonio. In this program, qualified students are assigned to active 'real world' research projects according to their skills and background and are supervised by a SWRI mentor responsible for the project and a UTSA faculty mentor. Students gain excellent research experience and a good understanding of the cross-disciplinary science and engineering technology involved in the projects. They have an opportunity to contribute directly to the research projects tasks and goals. Six students from underrepresented groups in Computer Science have participated in this program so far.

## **7. Outreach Activities**

The Summer Computer Science Institute (SCSI) offered an entry-level course in computer science to 24 high school students. The demographics of the class were 22 male and 2 female students, 15 hispanic and 2 black. The course was organized and taught by the Assistant Project Director, Dr. Robert Castaneda and spanned over ten weeks during the summer of 2004 for 2 hours twice a week with an additional lab once a week for 1.5 hours. Tutoring was also available to the SCSI students by undergraduate NSF Scholars. The purpose of this course was to stimulate interest in computer science among high school minority students and gear them towards pursuing a computer science degree at UTSA or elsewhere. The grant provided the financial support to SCSI students for tuition and fees. The outcome of this course was that 9 of the 24 students joined UTSA as freshmen in the Fall of 2004.

## **8. Selected Recent Publications**

Konstantinos Kyriakopoulos and Kleanthis Psarris, "Addressing the Issues in Data Dependence Analysis," *Proceedings of the ISCA Eighteenth International Conference on Parallel and Distributed Computing Systems*, pp. 199-205, Las Vegas, Nevada, September 2005.

Konstantinos Kyriakopoulos and Kleanthis Psarris, "Efficient Techniques for Advanced Data Dependence Analysis," *Proceedings of the Fourteenth International Conference on Parallel Architectures and Compilation Techniques*, IEEE Computer Society Press, pp. 143-153, St. Louis, Missouri, September 2005.

S. Chalasani and R.V. Boppana, "Adaptive Multimodule Routers for Multiprocessor Architectures." *Information Systems Frontiers*, Volume 7, Issue 3, Pages 317-327, July 2005.

S. Chalasani and R.V. Boppana, "Software Architectures for E-Commerce Computing Systems with External hosting." *International Journal of Computers & Applications*, Vol. 27, no. 3, pp. 190-198, 2005.

S. Desilva and R. V. Boppana, "Mitigating Malicious Control Packet Floods in Ad Hoc Networks," IEEE Wireless Communications and Networking Conference (WCNC), New Orleans, March 2005.

R.V. Boppana and Z. Zheng, "On the Performance of Metro-Scale Ad Hoc Networks," IEEE Wireless Communications and Networking Conference (WCNC), New Orleans, March 2005.

R. V. Boppana and Z. Zheng, "Designing Ad Hoc Networks with Limited Infrastructure Support," IEEE Consumer Communication and Networking Conference (CCNC), January 2005.

Konstantinos Kyriakopoulos and Kleanthis Psarris, "Data Dependence Analysis Techniques for Increased Accuracy and Extracted Parallelism," *International Journal of Parallel Programming*, Vol. 32, No. 4, pp. 317-359, August 2004.

S. Desilva and R. V. Boppana, "On the Impact of Noise Sensitivity on Transport Layer Performance in 802.11 based Ad Hoc Networks," IEEE International Conference on Communication (ICC), June 2004.

Kleanthis Psarris and Konstantinos Kyriakopoulos, "An Experimental Evaluation of Data Dependence Analysis Techniques," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 15, No. 3, pp. 196-213, March 2004.

R.V. Boppana, R.B. Boppana, and S. Chalasani, "Designing SANs to Support Low-Fanout Multicasts," International Conference on High-Performance Computing, Dec. 2003.

Kleanthis Psarris and Konstantinos Kyriakopoulos, "The Impact of Data Dependence Analysis on Compilation and Program Parallelization," *Proceedings of the Seventeenth ACM International Conference on Supercomputing*, ACM Press, pp. 205-214, San Francisco, California, June 2003.

T.D. Dyer and R.V. Boppana, "On Routing Web and Multimedia Traffic in Mobile Ad Hoc Networks," IEEE Hawaii International Conferences on System Sciences-36, January 2003.

# **CRI: Experimental Research in High Performance Computing and Wireless Networks**

Proposal # 0551501, First Year 2006-2007

University of Texas at San Antonio

PI: G.V.S. Raju

Co-PIs: R.V. Boppana, T. Korkmaz, A. Grigoryan, and O. Moreira-Tamayo

## **Abstract**

The project research covers various aspects of wireless networks and systems, high speed networking, parallel computing and image processing. The key characteristics of this project can be summarized as: multidisciplinary infrastructure development, multi-program sharing of research resources, enhancing current research efforts and creating new synergy among various research groups involved, leading to scholarly publications and additional research ideas, improving gender and ethnic diversity in graduate programs at UTSA. The research contributions are new techniques to solve various problems in computing and networking fields, experimental evaluation, the software and testbed configurations we develop, and technical know-how and our experiences with the testbeds. We believe that successful completion of the project on persistent wireless networks can have significant impact on the cost and accessibility to Internet for users, particularly those in remote areas. Our results on energy-efficient and secure use of ad hoc and sensor networks will facilitate making these networks more pervasive. We contribute to the state of the research by sharing our code, testbed setup and techniques. Our research results will be incorporated into our teaching curriculum via seminars, special topics courses and, eventually, regular courses on networks.

## **Introduction/Background**

The University of Texas at San Antonio (UTSA) is currently implementing its strategic plan to become a Research One university. UTSA was recently designated by the University of Texas System Board of Regents to become the system's next flagship institution. UTSA is a Hispanic Serving Institution (HSI) and is aggressively committing resources and efforts to recruit and maintain the minority pipeline of students interested in electrical and computer engineering (ECE), computer science (CS), and biotechnology. In the past academic year, enrollment at UTSA was over 27,000 with 59% minorities and 54% women. CS and ECE programs have 50 doctoral students, out of which 10% are women and minority students.

This is a collaborative project from the faculty in the ECE and CS programs at UTSA. We have several researchers working in highly active areas of national importance such as wireless networks and their security, high-speed networking, parallel and distributed computing, and image processing.

## **Goals, Objectives, Targeted Activities**

The overall goals and outcomes of the project are:

- Improve physical (computational and networking) infrastructure in ECE and CS programs to enable the faculty and students to conduct more experimental research.
- Improve human infrastructure through PhD student support and faculty development to contribute to the UTSA's strategic plan for becoming a Research One university.

The targeted activities are: multidisciplinary infrastructure development, multi-program sharing of research resources, enhancing current research efforts and creating new synergy among various research groups involved, leading to scholarly publications and additional research ideas, improving gender and ethnic diversity in graduate programs at UTSA and expanding pool of minority students electing to pursue academic positions or high tech jobs.

## Infrastructure and Facilities

**Networking and Computing Cluster:** An existing 20-node cluster of Linux workstations connected by Ethernet will be expanded and used to support our research on parallel computing and computer networking. The workstations will be equipped with Wi-Fi (IEEE 802.11) network interface cards (NICs) to provide both wired and wireless networks among the workstations. The Ethernet interconnect will be used for parallel computing and to emulate a miniature high-speed Internet – to study high-speed Internet protocols and quality of service issues. In addition, three layer 3 switches will be used to emulate mobile ad hoc networks and multidomain network for quality of service experiments in computer networking. The WiFi based network will be used to emulate an ad hoc wireless network for our experiments in wireless networks. For additional flexibility, each node will be configured with two wireless network interface cards. Using firewall software to filter packets appropriately and different channels, we can emulate mobile ad hoc network (MANET). This provides an easy to use platform to implement and test various network and transport protocols. For more rigorous testing of ad hoc network protocols, the ad hoc wireless network testbed described below will be used.

**Ad Hoc Wireless Network:** An existing 8-node ad hoc network testbed will be expanded with laptops and PDAs for a total of 25+ nodes. This equipment will supplement the computing cluster, in which can only emulate mobility. The MANET testbed will be used to evaluate the network software developed using the computing cluster.

**Sensor Network:** The sensor testbed will consist of two developer kits each with 25+ radios and MICA/MICA2 motes from Xbow or some other vendor. We plan to implement security protocols on tiny OS platform. To demonstrate the practicality of these protocols and to enable their widespread use for securing sensor networks, we plan to implement these protocols on the tiny OS platform, which has become the defacto operating system for the wireless sensor networks. The MICA and MICA2 motes will provide us with a testbed for developing and testing our protocols. Prior to deploying code on actual sensors it is necessary to debug this code on simulators such as TOSSIM. The simulations are very CPU-intensive, and the proposed computational cluster will remove a bottleneck for our research.

## Project Description

Several research projects are described below which will be pursued under this grant by the investigators and Ph.D. students.

### Ad Hoc Metropolitan Area Networks

We are currently designing and evaluating mixed wireless networks (MWNs) that are not tied to a particular communication link technology such as cellular, wired, and 802.11 or organizational category such as ad hoc and infrastructure. Instead, MWNs will intermix infrastructure and ad hoc networking concepts and can use all commonly used link technologies such as cellular, point-to-point (p2p) wired, and wireless links based on IEEE 802.11 or 802.16. To see the performance benefits mixed link technologies, we simulated an MWNs with up to 1000 nodes. Our simulations indicate that a few reliable links can increase the robustness of network greatly. In this project, we will investigate design issues such as optimal number of infrastructure nodes for a given ad hoc network to improve performance and route stability. We will also investigate secure routing techniques that leverage the infrastructure nodes to mitigate malicious attacks. This wireless network testbed will be used to develop software and verify our solutions in a controlled environment.

### Power-Aware Multi-Radio Wireless Networks

To improve capacity in multihop wireless networks, researchers have been proposing various approaches such as utilizing multiple channels and/or multiple radios. In typical multihop wireless networks, nodes are often powered by batteries with limited energy. In such networks, power control is essential for extending the lifetime of the underlying network. Moreover, power control helps to increase throughput by allowing better use of spatial spectrum. With this in mind, we are currently designing a power-aware multi-radio unification protocol (PMUP). We also plan to study the routing problem, particularly in case of heterogeneous networks where each node has different number of radios with different wireless link

quality. Finally, we will investigate how to make the operation of the proposed solutions more effective by considering cross-layer designs between network and data-link layers.

#### **A Security Layer for Mobile Ad-hoc Networks Based on Location Technologies**

The wireless network will be instructed to determine whether the GPS (Global Positioning System) position supplied by the node (self-positioning) is consistent with the network's own internal signal measurements. In most circumstances unauthorized nodes would not identify their true position to the network for fear of apprehension. If the network can identify that the actual position claimed by a node is statistically unlikely, that node could be blocked by the system. Currently we are developing software indoor positioning methods for nodes self-positioning. It is focused on developing software solutions based on novel algorithms with substantially reduced computational load. In this project, we will develop self-positioning methods for MANETs based on our current work on A-GPS. We will also verify the performances of the existing network-based positioning methods. We will use both methods for verifications in the proposed security layer.

#### **Mobile Ad-hoc Network Security**

The desired security goals are a) dynamic access control of nodes in the network, b) non-centralized detection of local, global or propagating attacks, c) thwarting attacks before they spread too far, and d) behavior based threat analysis and access control. In order to meet these goals we have proposed two systems, an authentication system and an intrusion detection system (IDS) to control access to the MANET and to combat possible attacks from adversaries.

Currently we are developing an authentication system supported by Air Force grant (# F306902-02-1-001). In this scheme, all packets are authenticated on every hop on their paths using distributed certification authority (CA) and threshold cryptography. A new node has to be initially authenticated by each of its neighbors to join the network. This project will show how a node will re-authenticate itself when the topology of the network changes due to node mobility.

In this project, we plan to develop a Distributed Intrusion Detection System (DIDS). The scheme is divided into two phases: 1) detection of malicious or compromised nodes in the network 2) defense against these nodes. The underlying principle behind detection of malicious nodes is that all the nodes on the MANETs also function as monitor nodes, passively monitoring the behavior of neighboring nodes and the network traffic that passes to or from them. Each node has an IDS running on it that has components of monitor, storage, decision/rule set and action. Rogue nodes will be detected by neighbors, who capture and filter traffic through certain rule-sets. When a node starts to exhibit anomalous behavior, the neighbor nodes will send out no confidence vote and if enough nodes send no confidence votes against a particular node, it is black listed. We will test our authentication and IDS systems through emulations on requested cluster. We will implement our design on the wireless network testbed. We will study denial of service, impersonation, replay, message tampering, man in the middle, and wormhole attacks. We will measure the performance metrics: throughput, latency for authenticating new node and blacklisting an intruder, and defense against attacks mentioned above.

#### **Parallel Processing of Large Optimization Problems in Biophysics**

Spectroscopic applications from fields such as astronomy, physics, and biochemistry involve experiments generating large amounts of data and subsequent analyses of the experimental data. A crucial and time-consuming computation task in these analyses is to identify the most relevant parameters to model the experimental data from a large number of parameters. This is achieved using standard optimization techniques, which fit experimental data to a linear combination of models. One such application from the field of biophysics is the analytical ultracentrifugation (AUC). AUC is a powerful technique for determining hydrodynamic properties of biological macromolecules and synthetic polymers in solution, and one where very large datasets are frequently encountered. Hydrodynamic parameters determined from such experiments allow the investigator to identify shape, molecular weight, and partial concentration of each solute in the mixture.

The computational task of analyzing the experimental data (denoted as *data* henceforth) can be organized into three phases: building appropriate parameterized models, eliminating noise effects on the data, and selection/identification of the parameters whose models best fit the data. A parameterized model is built

using partial differential equations (PDEs). High-performance computing techniques for this phase of the analysis are well-known and are not the focus of this paper. The resolution of the model is often dictated by the time-constraints and available memory limitations of a high-performance workstation. The computational complexity of time invariant noise removal and parameter selection/identification has received very little attention from the parallel computing community (at least for biophysics problems). In this project, we will address efficient and high-performance computational techniques for time invariant noise removal and parameter selection techniques. In particular, we will develop empirical models to predict the degree of parallelization that yields optimal execution time and resource usage.

#### **Tuning Parallel Linear Algebra Applications**

Dr. Whaley, a new faculty at UTSA, is the lead architect of the ATLAS and iFKO research projects, both of which are aimed at producing optimized compute kernels which are widely used in the high performance community. He was one of the key developers of the ScaLAPACK parallel linear algebra library, and the main author of its associated communication library, the BLACS. Additionally, he is a coauthor of the widely used HPL benchmark. He will use workstation cluster to develop new optimization techniques for these packages.

#### **New Algorithms in Medical Image Processing**

In this project, we will work on developing new efficient methods of image reconstruction by projections that are based on the concepts of non traditional vector and paired representation of images with respect to the two-dimensional discrete Fourier transform (2-D DFT). Such approach leads to effective implementation of image reconstruction on the Cartesian grid without interpolation of analytical formulas derived for the continuous transforms, such as the Radon transform. The discrete model of image reconstruction will be solved without referring to polar coordinates or any transformation defined for continuous functions. This is a key point. Such 2-D unitary discrete transforms will be constructed that are completely defined by linear projections. Such transforms have been developed by Grigoryan and are known as tensor and paired transforms. The unitary property of paired transforms allows us to derive inversion formulas that can be used in practical applications. The following tasks will be carried out in the course of this research: 1) Find optimal schemes of collecting attenuation measurements for calculating the splitting-signals representing uniquely the image (on the original Cartesian lattice) from projections. 2) Find the minimum number of projections to be used for image reconstruction, as well as the required number of measurements to be performed for each projection. 3) Analyze the discrete and continuous models of image reconstruction, and solve the 3-D problem of image reconstruction by paired transforms. 4) Develop algorithms, MATLAB based codes, and Graphic User Interface (software) for image reconstruction by proposed methods.

#### **Indicators of Success and Major Accomplishments**

We will implement an evaluation plan with goals: (1) To assess the extent to which the proposed activities are developed and implemented; (2) To assess their effectiveness and impact on the overall project goals and aims; and (3) To identify areas where improvement is needed; and (4) To assess the overall accomplishments of the project at the end of each budget cycle and end of grant period. Our evaluation plan incorporates both formative and summative evaluation components, and employs both quantitative and qualitative measures. Formative evaluation methods will be used to a baseline understanding of the current program components and activities as they relate to project goals. In this manner, program components and activities that are most predictive of a “successful” outcome will be identified and enhanced. Summative evaluation methods will be used after program interventions to assess the degree to which program goals have been realized for purposes of accountability. Quantitative measures will include academic performance of participating students, retention rate, numbers of presentations and publications, number of graduates taking up faculty ranks. Quantitative and qualitative data will be collected on an ongoing basis. Both “proximal” and “distal” outcomes will be assessed.

- Learning Outcomes: Fundamental knowledge, Research Skills.

- Performance Outcomes: Publications, Timely completion of Ph.D., High percentage of Ph.D. graduates going to Academic and Research Positions, Establishment of testbeds to continue experimental research.

### **Publications/Patents**

This is the beginning of the project. In the coming annual reports we will inform our publications/patents.

### **Outreach Programs and Supported Students**

UTSA is a Hispanic Serving Institution (HSI) and is aggressively committing resources and efforts to recruit and maintain the minority pipeline of students interested in electrical and computer engineering (ECE), computer science (CS), and biotechnology. In the past academic year, enrollment at UTSA was over 27,000 with 59% minorities and 54% women. CS and ECE programs have 50 doctoral students, out of which 10% are women and minority students. We have active recruiting program to get high quality women and minority students into CISE fields and will support all qualified students.

### **Future and Strategic Directions**

Since this is the beginning of the project all the activities and strategies indicated above indicate future and strategic directions.

# **CRI: DETER Cyber Community of Researchers (DECCOR)**

CNS-0454381

July 2005-June 2006

University of Southern California  
University of California Berkeley  
Clifford Neuman and Terry Benzel  
<http://www.isi.edu/deter>

## **1. Abstract**

The DECCOR project provides for infrastructure extensions and operations support required to support an increasing community of users of the DETER. This expansion allows research teams to develop experiments stressing concepts of scale, realism, and repeatability. The expansion also enables increased opportunities for experimentation by researchers beyond the initial base of DETER researchers (i.e. beyond the EMIST community). An important component of the DECCOR effort is to actively engage in community building and outreach through workshops and meetings that include researchers from many government funded programs.

In the first year of the effort we added 64 Dell 1850s at ISI, together with 11 Nortel 5510 switches to interconnect them and 2 Nortel 5530 switches to connect the nodes into the rest of the testbed. Also at ISI, we are adding another 32 nodes which were donated to us from the decommissioning of a cluster on campus. Of these 32 nodes, 12 will be setup as a mini-testbed, allowing us to reduce downtime for our user community by staging changes on this mini-testbed before committing them for use on the main testbed. At Berkeley we added 32 Dell 1850s for production, and we received a donation of 32 Compaq DL360G2 dual 1.4 GHZ Xeon servers, of which 16 have been put in service as a second mini-testbed. The intent of the Berkeley mini-testbed is the development of enhancements to Emulab and DETER.

## **2. Introduction and Background**

Despite many years of investment in network security research by government, academia, and industry, the current deployment of security technologies is inadequate to protect our critical Infrastructure. There are presently few metrics and little scientific basis to assess the value of emerging network security technologies, making it difficult to determine where new advances can truly protect our infrastructure. This project will improve the experimental infrastructure available to a broad community of academic, government and industry researchers.

Under prior funding by the NSF and DHS HSARPA, the DETER testbed project was established as a research testbed to serve researchers in cyber-security. The DETER testbed serves as a center for interchange and collaboration among security researchers, and as a shared laboratory in which researchers, developers, and operators from government, industry, and academia experiment with cyber security technologies under realistic conditions. The testbed provides tools and resources that enable secure, repeatable scientific experiments, allowing researchers to duplicate, analyze, and compare experiments.

With its strong security, containment, and usage policies, the testbed fills a role not currently met by existing testbed facilities, such as PlanetLab and by other instances of Emulab. Deter is based on Emulab, but is configured in a way that provides stronger containment and isolation from the external internet, at the cost of somewhat more constrained interactions between external researchers and experiments running on DETER.

With the increased hardware and support funded through the DECCOR effort, the testbed is now able to support an increasing community of users with needs for a larger testbed capability. This expansion has already resulted in significant growth of the DETER user community.

## **3. Goals, Objectives, Targeted Activities**

The goals of the DECCOR project were to provide for infrastructure extensions and support for operations that are required to support an increasing community of users with needs for a larger testbed capability. This expansion allows research teams to develop experiments stressing concepts of scale, realism, and repeatability. The expansion enables increased opportunities for experimentation beyond the initial base of DETER researchers. An important

component of the DECCOR effort was to actively engage in community building and outreach through workshops and meetings that include researchers from many government funded programs.

The hardware expansion initially planned for the DECCOR effort was to include the addition of 64-80 testbed nodes at ISI (proposal listed 64, but a table in the proposal showed 80) and 32 nodes at UC Berkeley. The effort also called for the addition of a full time system administrator to meet the needs of the growing community of DETER users, 1 GRA to support testbed operations and to improve tools, and tutorial materials for users. Finally, the proposal provided support for community outreach.

#### 4. Infrastructure and Facilities

As planned, in the first year we added 64 Dell 1850s at ISI, together with 11 Nortel 5510 switches to interconnect them and 2 Nortel 5530 switches to connect the nodes into the rest of the testbed. In addition to the 64 nodes at ISI, we are in the process of adding another 32 nodes at ISI, which were donated to us from the decommissioning of a cluster on campus. Of these 32 nodes, 12 will be setup as a mini-testbed, allowing us to reduce downtime for our user community by staging changes on this mini-testbed before committing them for use on the main testbed.

At Berkeley we added 32 Dell 1850s for production, and we received a donation of 32 Compaq DL360G2 dual 1.4 GHZ Xeon servers, of which 16 have been put in service as a second mini-testbed. The intent of the Berkeley mini-testbed is the development of enhancements to Emulab and DETER. The remaining 16 donated nodes will be brought up soon after we address some power issues.

	ISI nodes	Berkeley nodes	Mini-bed ISI	Mini-bed Berkeley
Original nodes	72	32	0	0
Planned DECCOR nodes	64	32	0	0
Added through DECCOR	64	32	0	0
Added from other sources	20 in progress	16 in progress	12	16 + 16 in progress
Total	136 + 20	96 + 16	12	16 + 16

With the addition of the nodes and switches, the testbed currently consists of the following:

- 64 IBM Netfinity 4500R in 3U cases (Dual 733MHz Pentium III processors, 1 GB memory, one 18 GB SCSI disk, 2 10/100 Ethernet interfaces)
- 8 Sun V65x in 2U cases (dual 2.8GHz Pentium 4 Xeon processors , 2 GB memory, Six 36 GB SCSI disks, 5 10/100/1000 Ethernet interfaces)
- 64 Dell PowerEdge 1850s in 1U cases (Dual 3.0GHz Pentium 4 Xeon processors, 2 GB memory, 1 36 GB 15K RPM SCSI disk, 5 10/100/1000 Ethernet interfaces)
- 62 Dell PowerEdge 1850s in 1U cases (Dual 3.0GHz Pentium 4 Xeon processors, 2 GB memory, 1 36 GB 15K RPM SCSI disk, 6 10/100/1000 Ethernet interfaces)
- 2 Dell PowerEdge 1850s in 1U cases (Dual 3.0GHz Pentium 4 Xeon processors, GB memory, 1 36 GB 15K RPM SCSI disk, 10 10/100/1000 Ethernet interfaces)
- 4 Juniper M7is (with 1.5GB of memory)
- At Berkeley 30 Sun V60x in 1U cases (Dual 2.8GHz Pentium 4 Xeon processors, 2 GB memory, Two 36 GB 10K RPM Ultra320 SCSI disks, 5 10/100/1000 Ethernet interfaces, 2 on-board Intel 10/100/1000 interfaces)
- 32 Dell PowerEdge 1850s in 1U cases (Dual 3.0GHz Pentium 4 Xeon processors, 2 GB memory, 2 36 GB 15K (10K?) RPM SCSI disk, 5 10/100/1000 Ethernet interfaces, 2 on-board Intel 10/100/1000 interfaces)

#### Switches and Routers

- One Cisco 6509 high-end switch. This functions as the *testbed backplane* (programmable patch panel) for the pc733s and pc2800s at ISI. Equipped with a 720 supervisor blade and seven 6748 48-port 10/100/1000 Ethernet modules, giving 336 1 Gbps Ethernet ports in total.
- Six Nortel Networks 5510-48T switches and one Nortel Networks 5530-48T, with an 80 Gb/sec interconnect. These serve as the *testbed backplane* ("programmable patch panel") for the pc3000s at ISI. Six (soon to be seven) identical switches provide the *testbed backplane* for the bpc3000s at UC Berkeley. Eight similar switches handle the pc3060s and pc3100s.

- The control network switches at ISI for pc3000s, pc2800s, and pc733s are 96 port Foundry FastIron Edge Switch 9604s. These provide “control” interfaces for the test nodes as well as regulating access using VLANs to the testbed servers.
- The control network switches at ISI for pc3060s and pc3100s are two Nortel Networks 5510-48T switches with 96 ports total. These provide "control" interfaces for the test nodes as well as regulating access using VLANs to the testbed servers.
- One Foundry FastIron 1500 Switch at UC Berkeley. This acts as both the control network switch and the *testbed backplane* ("programmable patch panel"), with 15 slots and 10 16-port 1Gbps line cards, providing 160 1 Gbps Ethernet ports.
- Connecting the Foundry VLANs together is a Dell P4 2.8GHz FreeBSD running PC as a router. It is used to provide security between the various VLANs on the control switch.

In addition to the nodes and switches, there are several pieces of special hardware connected to the testbed, including intrusion detection appliances and NetFPGA boards at UCB and ISI .

## 5. Project Description

The DECCOR project has contributed significantly to the computing infrastructure that is available for research and education in computer security. The size of the DETER testbed has almost doubled, and we have seen significant growth of the community using the DETER testbed. The effort has provided many opportunities to students working on the testbed deployment and operation, and to a larger set of graduate students (from many institutions) using the testbed to conduct research and complete laboratory assignments and has supported the development of skills within the community of users. The availability of the DETER testbed as a resource that is used by the computer security community has significantly improved the ability of researchers to study computer security attacks and countermeasures in a controlled environment that is scaled appropriately and sufficiently realistic to provide meaningful results.

## 6. Indicators of Success and Major Accomplishments

We have generated statistics on the use of the DETER testbed which show significant expansion of our user community beyond the initial EMIST experimenters. We attribute this growth in part to our outreach activities (described later) and because many government program announcements from DHS, NSF, and ARDA, have cited the availability of the DETER testbed for use in validating the effectiveness of proposed approaches. Of the current 197 users of the DETER testbed, only 34 are from the EMIST project, and the remainder are from a broader community of users. We can see the number of users that have logged in during different time periods in the following table:

Type	1 Month	3 Months	6 months	9 Months	1 Year
EMIST	10	14	17	19	21
Non-EMIST	43	62	93	135	145
Total	53	76	110	154	166

We can break the statistics down further by project within general category – EMIST, other academic, industry, government, and operational overhead. Note that the total number of users listed here exceeds the number of users because many users appear in more than one category/project. These statistics are broken down as follows:

Category	User count	Project count
EMIST	34	4
Other academic	97*	27
Industry	32	9
Government	12	5
Operational	39	9

In this table, the user count is the number of unique users. An individual project (a set of experiments) will often have multiple team members in the user count, including the PI, and graduate students, and other team members.

## Data on Swap-Ins

The data below is collected information on experiments swapped in to the testbed since the beginning of 2006. The table lists the number of times an experiment was swapped in, the total days of swapped in experiments (number of experiments, times number of days), and a load figure which is the number of nodes swapped in for an experiment times the number of days in use (e.g. a project that used 10 nodes for 10 days will show in the load column as 100 node-days). These figures are broken down by category of use.

Category	Exp Time (exp days)	Swap-ins	Load (node-days)
EMIST	121.47	52	636.58
Government	1.02	32	4.51
Industry	254.03	271	4341.85
Academic (not EMIST)	340.28	1100	5254.16
Testbed	311.89	536	1153.88
Total	1028.69	1991	11390.98

### Use of the DETER testbed for classes

In fall 2005, Anthony Joseph organized a new undergraduate Computer Security course (CS 161) at UC Berkeley. This new course covers the most important features of computer security, including topics such as cryptography, operating systems security, network security, and hardware-and language-based security. After completing this course, students should be able to analyze, design, and build secure systems of moderate complexity. The philosophy behind the course is to allow students to think both as designers and analyzers of security systems. This means equipping them with theoretical tools required to understand security issues, but also using copious exercises, projects, and case studies to allow students to see how security succeeds (and fails) in practice. The first offering of the course in fall 2005 attracted 125 undergraduate students. DETER plays an important role in the course. It offers a place where the students can safely experiment with and test their project software.

### Use of the DETER testbed for U.S. Homeland Security CyberStorm Exercise

The U.S. government sponsored the Cyber Storm national cyber security exercise, which was conducted by the Department of Homeland Security (DHS) on 6-10 February 2006. Among other things, the exercise demonstrated the capabilities of DETER to an extended community of cyber security stakeholders. The Cyber Storm experience provided insight into simulation requirements for national scale exercises and provided opportunities to expand the scope of DETER's current experimentation objectives, broadening our community of users.

## 7. Publications

The following paper describing the DETER testbed was written and presented in the past year:

- T. Benzel, B. Braden, D. Kim, C. Neuman A. Joseph, K. Sklower, R. Ostrenga, and S. Schwab, *Experience with DETER: A Testbed for Security Research*. Second IEEE Conference on testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom2006), March 2006.

The following papers were published describing work by users of the DETER testbed:

- R. Chertov, S. Fahmy, and N. B. Shroff, *Emulation versus Simulation: A Case Study of TCP-Targeted Denial of Service Attacks*, In Proceedings of 2nd International IEEE/CreateNet Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom), March 2006.
- L. Li, I. Hamadeh, S. Jiwasurat, G. Kesidis, P. Liu, C. Neuman, *Emulating Sequential Scanning Worms on the DETER Testbed*, In Proceedings of 2nd International IEEE/CreateNet Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom), March 2006.

## 8. Outreach Programs and Supported Students

During the first year of the DECCOR effort, there were many presentations made by DETER and DECCOR personnel explaining the capabilities of the DETER testbed to members of the computer security research community. While space limitation prevent us from listing each of these presentations, they numbered 30 and were made to audiences in industry, academia, and government. In addition to the presentations to members outside the DETER community, three workshops were organized where DETER users could discuss their requirements for the testbed and describe their experience and research results.

### **September 2005 Community Workshop**

This workshop was held in conjunction with the September 2005 CyberTrust Principal Investigators meeting in September 2005 in Newport Beach, California. This workshop provided an overview of the capabilities of the testbed and tools provided by the EMIST effort. A sampling of some of the projects already completed and in progress by the current testbed users gave potential new testbed users an idea of what they could accomplish using the testbed.. Coordination with the CyberTrust PI meeting ensured that the workshop could be attended by many potential new users of DETER and many of our recent new users were present at the workshop.

### **January 2006 Community Workshop**

An additional community workshop was held in January 2006 in Marina del Rey. This workshop focused on the needs of existing users of the DETER testbed, including new users of the testbed since the September meeting. An important agenda item for this workshop was soliciting additional information on testbed requirements from the users, so that the DETER and DECCOR team could prioritize the addition of new hardware types and software features to best meet the needs of our growing community of users.

### **June 2006 Demonstrations**

A DETER Community Workshop will be held on June 15 -16 in Arlington VA. This event will demonstrate the capabilities of the DETER Testbed and highlight the research of experimenters. Presentations at the workshop will cover current capabilities of the DETER testbed, as it has been extended through DECCOR, and the methodologies and tools developed (under the EMIST grant and other research efforts) as well as highlight the activities that have used the testbed. Attendees are expected from a wide community of researchers, industry, and government program representatives. The workshop will include sessions on lessons learned from the first three years of testbed usage, national exercises, and future plans for the DETER Security Experimenters Workbench. The workshop will also provide an opportunity for the user community to exchange information on tools and methodologies. Finally, the workshop will include live demonstrations and posters from current testbed users. Short extended abstracts have been submitted by users of the testbed and these papers will be presented and published subsequent to the workshop.

### **Undergraduate research training and development opportunities**

In fall 2005, two undergraduates participated in DETER and DECCOR projects. At UC Berkeley an undergraduate used the DETER testbed to collect Windows system call information. As a part of the research, the student designed and implemented a Windows device driver that intercepts all system calls and records them for anomaly detection analysis. At USC an undergraduate student was involved with the management of system scheduling, and helped to maintain and improve documentation for new users of the DETER testbed. In spring 2006, an undergraduate participated in DECCOR research by designing and implementing an analysis tool to collate and analyze the traces collected using the Windows device driver and an alternate DLL-based trace collection tool.

### **Graduate student research training and development opportunities**

The project has provided significant opportunity for training and development of graduate students who have been involved as student workers in the deployment of the new hardware and management of the testbed. Students at USC, Berkeley, and elsewhere have also conducted research using the testbed. In fact, most the actual research conducted on the DETER testbed is carried out by graduate students working for faculty members. In the statistics cited previously, the majority of DETER academic projects will have a single faculty member, with the rest of the accounts associated with the projects belonging to the graduate students that are conducting the research. We estimate that well over 50 students have used the DETER testbed in research.

## **9. Future and Strategic Directions**

The primary focus in the first year was the expansion of the testbed. We added 64 nodes at ISI and 32 nodes at Berkeley. In the second year we still have additional nodes to bring online. In particular, the effort will focus on mini-beds's at Berkeley and ISI, which will allow us to stage system upgrades to minimize downtime for testbed users. Much of our second year effort is focused on support for users in our growing user community, including development and improvement of testbed tools, and on providing full time dedicated support to testbed users.

# CRI: Developing the Next-Generation Open-Source Network Simulator (ns-3)

## CRI Proposal #0551686

Project Year: 2006

Thomas R. Henderson and Sumit Roy (University of Washington)

Sally Floyd (ICSI Center for Internet Research)

George F. Riley (Georgia Institute of Technology)

thenders, roy@ee.washington.edu, floyd@icir.org, riley@ece.gatech.edu

## Abstract

We are organizing a software development program to comprehensively re-design, enhance and maintain the popular Network Simulator (*ns*), to address research and educational challenges for the next generation of data networks. In our four-year program, we will i) refactor the simulator's architecture, ii) develop new networking protocol models, iii) provide new opportunities for software encapsulation, and iv) integrate the tool with virtual network testbeds. We will also introduce proven open-source practices that should enable *ns* development and software maintenance to become self-sustaining in the future, based on a large community of developers and power users. This project will advance the state-of-the-art in simulator design of the identified areas, along with supporting the mundane yet critical activities of code maintenance, documentation, integration, validation, and educational script generation. Since the resulting simulation code will be freely available to any individual or organization, *ns-3* will facilitate a significant increase in use for new simulation-oriented research as well as integration into courseware. Further, the project will emphasize a software development model that encourages and incorporates contributions from the user community.

## Introduction/Background

*ns-2* is the second major iteration of a discrete-event network simulation platform programmed in C++ and Object Tcl (OTcl). *ns-2* was first released in 1996, and derives from earlier work on the REAL simulator by Keshav and *ns-1* simulator by McCanne, Floyd, and Fall. *ns-2* is a major architectural change from *ns-1*; the simulator became entirely based on the blend of OTcl and C++.

The core of *ns-2* is written in C++, but the C++ simulation objects are also linked to shadow objects in OTcl. Simulation scripts are written in the OTcl language (an extension of the Tcl scripting language). This structure permits simulations to be written and modified in an interpreted environment without having to resort to recompiling the simulator each time a structural change is made. In the timeframe that *ns-2* was introduced (mid-1990s), this provided both a significant convenience in avoiding many time-consuming recompilations, and also allowing potentially easier scripting syntax for describing simulations. *ns-2* has a companion animation object known as the Network Animator (*nam*), used for visualization of the simulation output and for (limited) graphical configuration of simulation scenarios.

Presently, *ns-2* consists of over 300,000 lines of source code, with probably a comparable amount of contributed code that is not integrated directly into the main distribution but is maintained elsewhere. Figure 1 illustrates the past funding history of *ns*.

The **DARPA VINT** (Virtual InterNetwork Testbed) was a collaboration between USC/ISI, Xerox PARC, Lawrence Berkeley National Laboratory, and UC Berkeley. The project ran from 1997-2000, and was focused on developing the core simulation code for *ns-2* and *nam*. During this period, *ns-2* saw its most intense core development. During the course of building out *ns-2*, a number of research aims were also met, including the study of composable simulation frameworks,

abstraction techniques and tools, visualization techniques, real-time network emulation, and network topology and traffic generators.




Figure 1. Previous ns funding and results.

From 2000-2004, *ns-2* development was no longer funded directly by a dedicated project, but instead used two programs **DARPA SAMAN** (Simulation Augmented by Measurement and Analysis for Networks) and **NSF CONSER** (Collaborative Simulation for Education and Research) to fund code maintenance activities. During this period, core development of *ns-2* and *nam* slowed, although a number of contributed modules were ported in annually.

Since 2004, no funding for *ns-2* has been in place. Development and extension of the simulator has slowed further still; although there is still a quite active user community, there is less support to integrate contributed code into the project, and little available manpower for addressing core architectural issues. In 2005, one of the PIs (Henderson) began to maintain the *ns-users* and *ns-developers* mailing lists, coordinated an effort to streamline the source code licensing (to conform all licenses to a free software license model), and initiated the move of the *ns-2* project to the Sourceforge project hosting web site.

## Goals, Objectives, Targeted Activities

Despite *ns-2*'s popularity, there is a critical need to do core refactoring, software maintenance, and extension of the simulator. Experience has shown that issues like software maintenance (e.g., tracking compiler and operating system evolution), documentation, educational development, and code integration are lesser priorities for individual contributors focusing their limited time primarily on producing research output for publications. In particular, we have identified the following key areas that require development of a new *ns-3* simulator:

**Software core:** The current *ns-2* core often limits the scalability (both speed and memory footprint) to tens or perhaps a hundred or so nodes. The software is not 64-bit clean, is not modular enough, does not make use of more recently established object-oriented design principles, and does not have native support for multi-processor or distributed, federated simulations. The emulation capability (ability to interoperate with live networks) is outdated and of limited utility.

**Internal composition:** The software architecture places restrictions on how simulation objects may be combined in new ways. Many pieces of the software do not correctly interoperate (for example, wireless simulations can not use the OTcl-based routing protocols). This is repeatedly a reported source of frustration for *ns* users.

**External software integration:** *ns-2* is currently missing out on opportunities to leverage the wide number of advanced, externally developed programs such as open-source routing daemons (quagga, XORP, OpenBGP), traffic generators (iperf, tcplib), network virtualization testbeds (IMUNES, PlanetLab), network analysis software (Ethereal, tcptrace), etc. These tools can be integrated by the

development of better *ns-3* APIs and approaches to syntactically port application- and kernel-level code into the simulator's object-oriented, event-driven architecture.

**New models:** The protocol and traffic model library needs updating to cover more recent growth of IEEE 802-based wireless network and channel variants, IPv6 protocols, and newer popular applications and services (e.g., peer-to-peer, messaging, voice over IP, BitTorrent).

**Documentation:** The *ns-2* documentation has not been significantly updated for many years, is no longer consistent with the software core, and does not include more recently added modules.

**Visualization:** The *nam* animator is not well integrated with much of the code (wireless, Ethernet, satellite) and needs updating or replacement.

**Educational use:** More work is needed to build a comprehensive set of educational scripts, such as tracking the contents of popular undergraduate networking texts, providing protocol models that map better to actual implementations, and providing “newbie” support for simulator operation, visualization, and simple model development.

## Infrastructure and Facilities

Our software will be developed as application software using widely and freely available software tools and general purpose desktop and server computers. The infrastructure needs of our project are lightweight; namely, a reliable and maintained software source code repository server, tinderboxes or virtual machines for development and regression testing, mailing lists, infrastructure (mirrors) for software dissemination, and a web site. The institutions involved in *ns-2* and *ns-3* will support these infrastructure needs.

## Project Description

The *ns-3* development project is divided into multiple tasks. Below, we provide brief description and justification for three major tasks: i) core refactoring, ii) integration, and iii) maintenance.

### 5.1 Core refactoring

This task will design and refactor the core of the *ns-2* simulator. The output of this task will be a cleaner, more scaleable, and more extensible core. George Riley (architect of the Georgia Tech Network Simulator (*GTNetS*) and parallel, distributed extensions to *ns-2*) will lead this task. In this section, we outline a set of design objectives for the *ns-3* simulator. We see these as basic requirements for the design, implementation, and overall success of this tool.

**Reuse:** First, the *ns-3* design and implementation should leverage large parts of the code base of existing tools, such as *ns-2*, *GTNetS*, and others. These existing tools have hundreds of thousands of lines of code, and hundreds of existing network models. While we fully expect some modifications to these code bases will be necessary to fit within our overall design, we will strive to avoid complete re-design or re-implementation of existing modules as much as possible. To achieve this goal, we plan to design an automated tool to analyze and convert of much of the existing *ns-2* OTcl code to enable inclusion in our design. We expect to be able to reuse much of the existing *ns-2* and *GTNetS* code.

**Scalability:** One of the major concerns about *ns-2* cited by its users is scalability. *ns-2* is a sequential execution simulator with a single event processing loop running on a single processor. Although such a simulator can scale to hundreds of nodes when underlying communications models are heavily abstracted, the memory and processing resources of a single machine become a bottleneck when more sophisticated channel models (e.g., wireless) or higher-rate links (e.g., 10 Gbps) are included. Researchers have taken various approaches to improve *ns-2* scalability, including the caching of redundant computations and function calls (the “Staged NS (SNS)” project at Cornell), use of on-demand route computation, and partitioning the simulation into wireless clusters (the *ns-2* *gridkeeper* and similar structures).

There are a number of factors contributing to the scalability limitation, but the fundamental bottleneck is the execution on a single processor. We believe that the *ns-3* simulator should be designed from the outset to support parallel and distributed simulation. We have already applied these concepts to *ns-2* to create a parallel, distributed version of *ns-2* (PDNS). In previously reported results, PDNS achieved a speedup of a factor of 80 and a simulation speed of 6 million packet hops per second, on 128 processors).

**Other Design Considerations:** Our *ns-3* design will also consider the following: better application of modern C++ design such as the use of C++ templates, polymorphism, and type-casting facilities; replacement of the OTcl scripting capability with a new scripting interface that avoids the split-object implementation paradigm that has proven to be problematic for *ns-2*; better match of simulation model design to the design of real networks and real network elements; memory-efficiency considerations; tracing and statistics logging facilities; emulation capabilities; network topology creation facilities; and visualization.

## 5.2 Integration

A key step forward of our proposed *ns-3* project will be the level of integration that we will obtain with the vast amount of free, open-source software and research projects available on the web. We have three specific goals in mind:

- 1) Extension of the simulation capability via integration with open-source software (e.g., Ethereal packet analysis, Click/XORP routing).
- 2) Abstraction layers and interfaces for porting implementation code into the *ns* environment; and
- 3) Interfaces to allow users to easily migrate between simulation and network emulation environments.

**Open source integration:** An opportunity that most every simulation environment has missed is the opportunity to leverage the extensive amount of free, open-source networking code within operating systems and applications. Typically, simulators reimplement protocols from scratch, leading to a costly software effort and divergence from actual implementation code. There are limited exceptions (the TCP code in QualNet, for example, is ported from BSD, and the NCTUns simulator uses a kernel-reentrant programming paradigm to use actual Linux stack code), but predominantly, protocols are reimplemented for the simulation environment. It is also difficult in general to write simulation software that can run both in simulation and implementation environments. The Naval Research Laboratory's *protolib* toolkit is an example of a publicly-available library that allows newly developed models to be written to a software abstraction library supported both in a simulation and implementation environment. This approach works well if a protocol implementation is written from scratch; however, it generally does not work so well when existing software, often written in a lower-level, non-object-oriented language such as C, is used.

In the *ns-3* project, we intend to focus on simulator design that facilitates the reuse of existing software and applications. Such an approach helps to meet our educational goals as well, since the simulation models mimic how the software is run in real implementations. Specifically, we see the following key opportunities for software integration and reuse: i) ported *application code* using sockets API, ii) *routing protocols* such as XORP, quagga, and OpenBGP, iii) *network stack code* such as the Network Simulation Cradle, and iv) *tools to parse output data* such as tools that work on pcap format traces such as tcpdump and Ethereal.

Finally, we recognize the significant research infrastructure advances of the past few years, with such projects as PlanetLab, Emulab, and WHYNET. These testbeds provide opportunities to explore protocol interactions in less controlled environments (such as cluster-based, remotely executed testbeds including Utah's Emulab) and to deploy long-running experimental services and overlays on the existing Internet (PlanetLab). Presently, Emulab uses *ns* scripting syntax to describe its experiments, and offers a version of the *ns* emulation environment to experimenters. Our project will coordinate with testbed projects to ensure that *ns* can successfully execute in those environments and is well documented enough for ease of use.

## 5.3 Maintenance

Some activities planned for this task include:

- h. develop functional and technical specifications for the new software design,

- i. establish a source code control system, coding styles, and software development model, including tinderboxes for daily builds on different platforms,
- j. *ns-2* has a validation facility for ensuring that simulation output in one module is not changed by a code change in an unrelated module. Create new validation test suites for the new simulator and models. Port validation scripts that do not work in backward compatibility mode,
- k. update the documentation and create tutorials,
- l. establish a web presence for collaborative software development, wikis, mailing lists, and dissemination of the software, and
- m. attend and participate in annual program review meetings, at sites to be determined.

Once *ns-3* is established, we intend to get a broader community involved in ongoing maintenance activities, by encouraging contributors and power users to take up portions of the maintenance activities, as in other open-source projects.

## Indicators of Success and Major Accomplishments

Broadly, we intend for our simulator to become a preferred simulation tool for research and education on data networking. The current version of *ns* (*ns-2*) enjoys widespread use in the research community; the simulation code has been contributed by over one hundred individuals and organizations, there are several thousand logged downloads per month, and use of the simulator is consistently referenced in 10 to 25% of the papers in the top conferences in the field. We intend to maintain or improve these statistics with our next-generation simulator.

## Publications/Patents

We have collectively authored one publication under submission:

T. Henderson, S. Roy, S. Floyd, and G. Riley, “*ns-3* Project Goals,” (*submitted to*) *Workshop on ns-2*, Pisa Italy, October 10, 2006.

## Outreach Programs and Supported Students

Presently, since our project has not yet officially started, we do not have outreach programs in place. Our project intends to develop and discuss the simulation software in open forums on the web. We also intend to collaborate with other projects. We have one collaboration already in place with the Planete research group at INRIA Sophia-Antipolis, who is funding software development for the next-generation simulator.

We have budgeted for partial funding of one student at the University of Washington, pending the funding of our program.

## Future and Strategic Directions

Our project intends to hold a software design review and kickoff meeting in the mid-August timeframe. Future directions will be discussed on the *ns*-developers mailing list.<sup>1</sup>

---

<sup>1</sup> <http://mailman.isi.edu/mailman/listinfo/ns-developers>

## **CRI: An Open Linguistic Infrastructure for American English**

**Nancy Ide (PI)**  
**Vassar College**  
ide@cs.vassar.edu

**Christiane Fellbaum (co-PI)**  
**Princeton University**  
fellbaum@princeton.edu

**Charles Fillmore (sub-contractor)**  
**Collin Baker (sub-contractor)**  
**International Computer Science Institute**  
{fillmore, baker}@icsi.berkeley.edu

This NSF/CRI grant supports planning and preparation of the Open Linguistic Infrastructure (OLI) for American English, in anticipation of full funding at a future date.

The project is a collaboration between the American National Corpus (ANC, Ide and Suderman, 2004, 2006) and the originators of the WordNet (WN, Fellbaum 1998) and FrameNet (FN, Ruppenhofer et al., 2002) databases, for the purpose of creating a sub-corpus annotated with WN senses and FN frames. One of the primary activities of the collaboration is to work toward harmonizing all three resources and maximally exploiting their respective strengths. Another goal is the continued development of the ANC, which currently includes 22 million words of written and spoken data across a variety of genres and is distributed through the Linguistic Data Consortium.

If fully funded, the project will use the American National Corpus as a core around which to build an "Open Linguistic Infrastructure" (ANC-OLI) for American English, including a wide variety of linguistic resources derived from and/or linked to the ANC data. These resources will include frequency, bi-gram, and tri-gram lists for lexical items and grammatical categories (with sub-lists by genre), collocation data, etc., together with multiple alternative annotations of the ANC data for various linguistic phenomena, including part of speech, syntax (both chunking and full parse), and semantics (sense tags and frames). The centerpiece will be the 10 million word gold standard corpus, to which we will add hand-validated annotation for syntax and named entities together with semi-automatically produced annotations for WordNet senses and FrameNet frames. All annotations will be freely downloadable from the ANC website as stand-off XML documents linked to the original ANC data.

The ANC is a representative corpus, including texts from a wide range of genres. Therefore, it reflects the use of modern American English by a broad community of speakers. Up to now, computational linguistics research has relied primarily on very small, outdated corpora such as the 1961 Brown Corpus, newspaper data, or domain-specific corpora such as the Wall Street Journal. Consequently, the resulting language models are biased and incomplete. One of the primary motivations for constructing the ANC is to provide both a corpus representing a wide variety of genres and a representative sample of newly emerging language types such as blogs, chats, and e-mail.

The major activities within the project are the following:

## 1 Automated annotation

We will use freely-available or contributed annotation software to automatically generate multiple alternative annotations for a broad range of linguistic phenomena. These annotations will be provided "as is", with no validation; our task will be to render the output in the ANC stand-off format and link it to the ANC data. For example, for syntax, we will provide various types of syntactic analyses generated with freely available existing tools. Numerous syntactic parsers are freely available on-line, such as MINIPAR<sup>2</sup>, the CMU Link grammar parser<sup>3</sup>, the XTAG parser<sup>4</sup>, Dan Bickel's<sup>5</sup> and Collins's<sup>6</sup> and Charniak's<sup>7</sup> statistical parsers, etc. Noun phrase and verb phrase chunking will be automatically generated using freely available systems such as the chunkers in GATE, LT CHUNK<sup>8</sup>, and Ramshaw and Marcus' chunker<sup>9</sup>. Multiple alternative annotations provide information suited to different schemes and linguistic theories, thereby comprising a valuable resource in their own right. They also enable researchers to merge annotations for comparative analysis and to disambiguate automatically produced tags.

## 2 Gold standard validation

The 10 million word "gold standard" portion of the ANC corpus (GSC) currently contains validated annotation for document structure (titles, headings, etc.), paragraph, sentence, and token boundaries, part of speech, and lemma. Materials in the GSC are distributed by genre in approximately the same proportions as the full ANC and cover both written and spoken data. This project will add hand-validation to the GSC for syntax and named entities (including at least persons, dates, organizations, and locations).

## 3 Semi-automatic WordNet and FrameNet annotation

There is a woeful lack of large-scale semantically annotated corpora available for development of language processing applications and linguistic research, especially covering a range of genres. A primary reason for this is that semantic annotation is inherently difficult, time-consuming, and costly, and most research groups are understandably reluctant to undertake such an endeavor.

This project will address this need by semi-automatically annotating a portion of the GSC with both WordNet senses and FrameNet frames, thus providing an unprecedented resource

---

<sup>2</sup> <http://www.cs.ualberta.ca/~lindek/minipar.htm>

<sup>3</sup> <http://www.link.cs.cmu.edu/link/>

<sup>4</sup> <http://www.cis.upenn.edu/~xtag/swrelease.html>

<sup>5</sup> <http://www.cis.upenn.edu/~dbikel/software.html>

<sup>6</sup> <ftp://ftp.cis.upenn.edu/pub/mcollins/PARSER.tar.gz>

<sup>7</sup> <ftp://ftp.cs.brown.edu/pub/nlparser/>

<sup>8</sup> <http://www.ltg.ed.ac.uk/software/chunk/>

<sup>9</sup> [ftp://ftp.cis.upenn.edu/pub/chunker/basenp\\_chunker\\_1/](ftp://ftp.cis.upenn.edu/pub/chunker/basenp_chunker_1/)

containing not only complete WN sense information, but also full annotation of FN lexical units (words and phrases that evoke semantic frames) and their frame elements (slot fillers). The collaborative annotation of the same texts will also lead toward harmonization of Wordnet and FrameNet, a desideratum frequently expressed by members of the NLP research community. As for all ANC annotations, the WN and FN annotations will be rendered into the ANC stand-off format and made freely available on the ANC website.

The WordNet team will manually assign WordNet senses to all content words (nouns, verbs, adjectives, and adverbs) in approximately 3 million running words of the GSC, producing about 720,000 sense-tagged tokens. This will include the 100K FN-annotated portion, which will be annotated first. The FrameNet team will manually annotate a subset of the WN-annotated data with FrameNet frames and frame elements.

## 4 Enhancing automatic annotation performance

The comparisons of the results of alternative annotations for the same linguistic phenomenon generated by different tools can be used to produce more accurate annotations. In addition, a variety of machine learning techniques have recently been explored to improve automatic annotation of a variety of linguistic phenomena, including parsing (e.g., Steedman, 2003; Hwa, 2004), semantic role labeling (e.g., Pradhan *et al.*, 2004), word sense disambiguation and sense assignment (see for example the report in Mihalcea, *et al.*, 2004), named entity recognition (e.g., Baluja *et al.*, 2000; papers in Ciravegna *et al.*, 2000; Zhang *et al.*, 2004), and others. One of the important aspects of our work will be to explore and expand these techniques to ensure that the quality of automatically-generated annotations in the ANC is optimized.

## 5 Summary

The vision for the ANC-OLI is to continually augment the infrastructure with contributed annotations from the research community, so that in the future annotations for other phenomena such as co-reference, discourse structure, additional entities, events, opinions, etc. will be added. We feel strongly that distribution of effort, together with integration of currently independent resources such as the ANC, WordNet, and FrameNet, will enable progress in resource development well beyond what can be accomplished at individual sites working independently (which is the model in operation at the moment), for considerably less cost, and achieving a greater degree of accuracy and usability.

The gold standard portion of the ANC with validated annotation for part of speech, syntax, named entities, and semantic annotation for WN senses and FN frames, will provide for the first time a large, representative corpus that can be used to train language processing software. Its availability should have a major impact on the speed with which similar resources can be reliably annotated.

The addition of semantic annotation for WN senses and FN frames will make the ANC the largest semantically annotated corpus of English in existence and provide a much-needed resource for computational linguistics research aimed at the development of robust language processing systems. Because both WN and FN are linked to corresponding resources in other

languages, WN and FN annotation of the ANC will immediately create a massive multi-lingual resource network, in which sense-tagged words are linked to their translation equivalents (and their synsets) in the 30 existing wordnets in other languages, as well as to frame information in the growing number of framenets throughout the world. The unprecedented nature and value of such a resource for machine translation and other multi-lingual NLP applications cannot be underestimated, as no existing resource even approaches this size and scope.

The funding provided will allow us to further explore and prepare for the proposed work, perform some preliminary work, and provide a solid basis for a future proposal. To that end, we undertake the following activities:

Creation and annotation with WN senses and FN frames and frame elements of a small corpus of materials drawn from the ANC, representing a range of genres. This will enable us to experiment with methods of annotation and coordination among the ANC, FN, and WN projects.

Further research into and experimentation with methods and software to enhance automatic annotation accuracy for semantic roles, syntax, sense tagging, etc. In addition, the process of harmonizing the two lexical resources (FN and WN) with each other will be further explored.

Outreach to the US computational linguistics community, to ensure that the development of the ANC and the surrounding infrastructure both serves the needs of computational linguistics research and development and employs cutting edge methodologies for the creation of the resources we intend to provide. This will include a workshop to present the plans for the ANC Open Linguistic Infrastructure (OLI) to the research community and solicit input concerning the community's needs and interests. The outcome of the workshop will be a list of specific priorities and methods, as recommended by the attendees, for both the ANC-OLI as a whole and the 10 million word hand-validated sub-corpus.

## References

- Baluja, S., Mittal, V., and Sukthankar, R. 2000. Applying machine learning for high performance named-entity extraction. *Computational Intelligence*, 16.
- Ciravegna, F., Basili, R. and Gaizauskas, R. 2000. Workshop on Machine Learning for Information Extraction. Held in conjunction with the 14th European Conference on Artificial Intelligence (ECAI), Berlin. Papers at <http://www.dcs.shef.ac.uk/~fabio/ecai-workshop.htm>.
- Fellbaum, C. (Ed.) 1998. *WordNet: A Lexical Database of English*. Cambridge, MA: MIT Press.
- Hwa, R. 2004. Sample selection for statistical parsing. *Computational Linguistics*, 30:3.
- Ide, N. and Suderman, K., 2004. The American National Corpus First Release. *Proceedings of the Fourth Language Resources and Evaluation Conference (LREC)*, Lisbon, Portugal, 1681-1684.

- Ide, N. and Suderman, K., 2006. Integrating Linguistic Resources: The American National Corpus Model. *Proceedings of the Fifth Language Resources and Evaluation Conference (LREC)*, Genoa, Italy.
- Mihalcea, R., Chklovski, T. and Kilgarriff, A. 2004. The Senseval-3 English lexical sample task. In Mihalcea, R., Edmonds, P. (Eds.), *Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*, Association for Computational Linguistics, 25-28.
- Pradhan, S., Hacioglu, K., Ward, W., Martin, J., and Jurafsky, D. 2003. Semantic role parsing: Adding semantic structure to unstructured text. *Proceedings of the International Conference on Data Mining (ICDM-2003)*.
- Ruppenhofer, Josef, Collin F. Baker and Charles J. Fillmore (2002): The FrameNet Database and Software Tools. In Braasch, Anna and Claus Povlsen (eds.), *Proceedings of the Tenth Euralex International Congress*. Copenhagen, Denmark. Vol. I: 371-375.
- Steedman, M., R. Hwa, S. Clark, M. Osborne, A. Sarkar, J. Hockenmaier, P. Ruhlén, S. Baker, and J. Crim. 2003. Example Selection for Bootstrapping Statistical Parsers. *Proceedings of the Annual Meeting of the North American Chapter of the ACL*, Edmonton, Canada.
- Zhang, L., Pan, Y. and Zhang, T. 2004. Focused Named Entity Recognition using Machine Learning. *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, Sheffield, UK, 281-288.

## Author Index

<p><b>Aazhang, B.</b> ..... 127</p> <p>Acharya, R. .... 112, 137</p> <p>Adjouadi, M. .... 57</p> <p>Agrawal, G. .... 107</p> <p>Ahalt, S. .... 107</p> <p>Alexander, W. E. .... 102</p> <p>Argamon, S. .... 132</p> <p><b>Bachnak, R.</b> ..... 142</p> <p>Bailey, D. .... 87</p> <p>Bajaj, C. .... 202</p> <p>Baker, C. .... 232</p> <p>Barreto, A. .... 57</p> <p>Becerra-Fernandez, I. .... 47</p> <p>Belding-Royer, E. .... 37</p> <p>Benzel, T. .... 222</p> <p>Bertino, E. .... 192</p> <p>Bestavros, A. .... 6</p> <p>Betke, M. .... 6</p> <p>Bhattacharya, A. .... 87</p> <p>Bieman, J. M. .... 27</p> <p>Blume, M. .... 167</p> <p>Boppana, R. .... 212,217</p> <p>Bramley, R. .... 77</p> <p>Brodersen, R. .... 147</p> <p>Buddhikot, M. M. .... 37</p> <p>Burger, D. .... 202</p> <p>Byers, J. .... 6</p> <p><b>Calvert, K.</b> ..... 182</p> <p>Catalyurek, U. .... 107</p> <p>Cavallaro, J. .... 127</p> <p>Chandra, N. .... 42</p> <p>Chandra, U. .... 42</p> <p>Chellappa, R. .... 187</p> <p>Cheng, B. H. C. .... 27</p> <p>Chien, A. A. .... 162</p> <p>Connelly, K. .... 77</p> <p>Cook, J. .... 87</p> <p>Cook, P. .... 117</p> <p>Crovella, M. .... 6</p> <p><b>Dantu, R.</b> ..... 192</p> <p>Das, C. R. .... 112</p>	<p>Doyle, J. .... 17</p> <p>Dua, S. .... 82</p> <p>Duraiswami, R. .... 187</p> <p><b>Esterline, A.</b> ..... 102</p> <p>Estrin, D. L. .... 157</p> <p><b>Fahmy, S.</b> ..... 192</p> <p>Fei, Z. .... 182</p> <p>Fellbaum, C. .... 232</p> <p>Ferhatosmanoglu, H. .... 107</p> <p>Fernandez, J. D. .... 142</p> <p>Fillmore, C. .... 232</p> <p>Flanagan, M. .... 72</p> <p>Floyd, S. .... 227</p> <p>Fox, G. .... 77</p> <p>France, R. .... 27</p> <p>Frantz, P. .... 127</p> <p><b>Gannon, D.</b> ..... 77</p> <p>Garcia, M. .... 142</p> <p>Garg, A. .... 137</p> <p>Ghosal, D. .... 192</p> <p>Girod, L. .... 157</p> <p>Griffioen, J. .... 182</p> <p>Grigoryan, A. .... 217</p> <p>Grosz, B. J. .... 62</p> <p>Guy, R. G. .... 157</p> <p><b>Hartley, R.</b> ..... 97</p> <p>He, J. .... 1</p> <p>Henderson, T. R. .... 227</p> <p>Huang, C. H. .... 172</p> <p>Hwu, W. .... 177</p> <p><b>Ide, N.</b> ..... 232</p> <p>Irwin, M. I. .... 112</p> <p>Itkis, G. .... 6</p> <p>Iyer, R. .... 177</p> <p><b>JaJa, J.</b> ..... 187</p> <p>Jayakar, P. .... 57</p> <p>Jaynes, C. .... 182</p> <p>Jeffery, C. .... 87</p> <p>Jin, H.-W. .... 107</p>
--	--

<b>Kaiser, W. J.</b> .....	157	Moreira-Tamayo, O.	217
Kar, D.	142	<b>Nahrstedt, K.</b> .....	177
Keeling, H. N.	67	Nava, P. A.	207
Kelly, J.	102	Neuman, C.	222
Kfoury, A.	6	Newman, H. B.	17
Kim, J.	102	Nikolić, B.	147
Kim, J. H.	102	<b>O'Leary, D.</b> .....	187
King, S.	142	<b>Pagan, I.</b> .....	197
Knightly, E.	127	Pai, V.	117
Kohler, E.	157	Palmer, M.	12
Kollios, G.	6	Panchanathan, S.	1
Korkmaz, T.	217	Panda, D. K.	107
Kriegman, D.	162	Parthasarathy, S.	107
Kurc, T.	107	Pasquale, J.	162
<b>La Porta, T.</b> .....	37	Pasztor, A.	57
Lauria, M.	107	Patterson, D.	152
Leach, R. J.	67	Peters, D.	87
Leake, D.	77	Peterson, L.	117, 122
Lee Giles, C.	112	Pfeffer, A.	62
Lee, D.	107	Pierluissi, J.	207
Lee, M.	22	Pincus, S.	82
Leung, H.	97	Plale, B.	77
Lewis, D. D.	132	Pontelli, E.	92, 97
Li, K.	117	Prietula, M.	47
Low, S. H.	17	Psaltis, D.	17
Lumsdaine, A.	77	Psarris, K.	212
<b>Machiraju, R.</b> .....	107	Pustejovsky, J.	12
MacQueen, D.	167	<b>Qian, Y.</b> .....	197
Madigan, D.	132	<b>Raghavan, P.</b> .....	112
Maloney, D.	182	Raghavan, V.	82
Marcus, M.	12	Raju, G.V.S.	217
Mark, D.	137	Ramachandran, V.	202
Martonosi, M.	117	Ranjan, D.	92, 97
Massey, D.	32	Reddy, G.B.	102
Matta, I.	6	Reddy, Y. B.	82
Mazur, J.	182	Reppy, J.	167
McBeath, M.	1	Reyes, M.	102
Mcdonough, J.	182	Reyzin, L.	6
McKinley, K.	202	Rikakis, T.	1
McMahon, C.	82	Riley, G. F.	227
Meyers, A.	12	Rishe, N.	52
Meyers, D.	32	Rivera, W.	197
Miikkulainen, R.	202	Rodriguez, D.	197
Miller, R.	137		
Miller, S.	37		

Rodriguez, M. ....	197	Trimble, J. ....	67
Rodriguez, N. ....	197	Tse, D. ....	147
Roig, G. ....	57	<b>Varshney, A. ....</b>	<b>187</b>
Roy, S. ....	227	Villaverde, K. ....	92
<b>Sabharwal, A. ....</b>	<b>127</b>	<b>Wang, L. ....</b>	<b>32</b>
Sadayappan, P. ....	107	Wawrzynek, J. ....	152
Saltz, J. ....	107	Weiss, S. ....	72
Sanders, W. ....	177	West, R. ....	6
Savage, S. ....	162	Wiebe, J. ....	12
Schulzrinne, H. ....	37, 192	Williams, D. H. ....	207
Sclaroff, S. ....	6	Williams, K. ....	102
Seales, W.B. ....	182	Winters-Hilt, S. ....	82
Shen, H.-W. ....	107	Wise, D. ....	77
Shieber, S. M. ....	62	Wolberg, G. ....	22
Sinha, P. ....	107	Wongsaroj, B. ....	52
Son, T.C. ....	92	Wyckoff, P. ....	107
Srinivasan, A. ....	42	<b>Xi, H. ....</b>	<b>6</b>
Srivastava, M. B. ....	157	Xiao, J. ....	22
Steidley, C. ....	142	Yoon, S. H. ....	102
Stredney, D. ....	107	<b>Zhang, A. ....</b>	<b>137</b>
Stutz, A. E. ....	107	Zhang, L. ....	32
<b>Teller, V. ....</b>	<b>72</b>	Zhu, Z. ....	22
Thompson, H. W. ....	82		
Tran, S. C. ....	97		