COMPUTER & INFORMATION SCIENCE & ENGINEERING

CISE NEWS

FALL 2009

WWW.CISE.UFL.EDU

# ASSISTANT PROFESSOR MY T. THAI RECEIVES THE PRESTIGIOUS DOD YIP AWARD

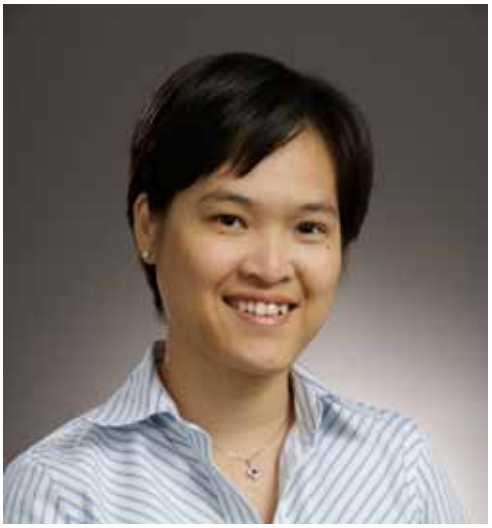

My T. Thai, an assistant professor in the CISE Department, received the DoD YIP Award for her research project "C-WMD: Models, Complexity, and Algorithms in Complex Dynamic and Evolving Networks." The Young Investigator research Program (YIP) recognizes and supports early career development of outstanding young investigators who show exceptional ability and promise for conducting basic research. In addition to this award, Thai's research group was supported by two other NSF grants.

Her group is working on combinatorial optimization and its application in networks, including computer networks, online social networks, and biological networks. More specifically, her group research focuses on the design and analysis of several new models and approximation algorithms for various practical problems, mostly optimization, arising from the networking fields mentioned above. These studies have led to many beautiful and challenging questions and results in theory, which enrich other computer science and mathematical areas, such as graph theory and advance optimization theory. Besides seeking theoretical solutions with a provable performance bound, the group is also concerned about their practical implementation in real life. More details about her research group can be found at www.cise.ufl.edu/AppliedOptimization.

One of the most challenging aspects in studying this type of network is that it is complex, dynamic, and evolving. It is often greatly affected by uncertain factors which arise in many forms, including natural or man-made interferences. This can result in dynamic changes in network characteristics, thus leading to a set of fundamental issues: How does one characterize and forecast the interdependent response of network components in evolving networks with limited data? What are the quantitative benefits of using adaptive solutions via re-computing them from scratch for each change? How does one design an approximation algorithm which can be adaptive to changes with a guarantee on the performance bound? This new approximation technique is essential, yet difficult, for adaptive recovery strategies and once found, it can provide a mathematical framework for several existing problems in dynamic networks such as routing protocols, network design and management.

Thai's research group has introduced several new problems and concepts based on their studies, such as Size Constraint Group Testing, Minimum Weight Routing Backbone, and beta-Disruptor. They have recently introduced a novel framework and efficient algorithms against several types of DoS attacks on both the Internet and Wireless Sensor Networks. They have successfully brought the alpha-beta tree into wireless networks serving as a performance analysis framework for energy consumption

*Continued on next page*

DR. THAI'S RESEARCH GROUP HAS INTRODUCED SEVERAL NEW PROBLEMS AND CONCEPTS BASED ON THEIR STUDIES

• • •

THEY HAVE RECENTLY INTRODUCED A NOVEL FRAMEWORK AND EFFICIENT ALGORITHMS AGAINST SEVERAL TYPES OF DOS ATTACKS

• • •

# FACULTY NEWS

**MANUEL BERMUDEZ** / *associate professor* / was appointed as the Technical Committee Chair for the 2010 Latin American and Caribbean Consortium of Engineering Institutions Conference, taking place in Arequipa, Perú, May 31 - June 4, 2010.

**SHIGANG CHEN** / *associate professor* / received a $523,818 grant from the National Science Foundation for his project tilted "CPS:Small: Transforming a City's Transportation Infrastructure through an Embedded Pervasive Communication."

**AHMED HELMY** / *associate professor* / was a PC track chair and chair for the poster and Ph.D. forum sessions at the IEEE International Conference on Network Protocols (ICNP), in Princeton, NJ, Oct 13-16, 2009. Helmy was invited to give a tutorial at the fifth IEEE / ACM International Wireless Communications and Mobile Computing Conference in Leipzig, Germany, from June 21- 24, 2009. He will also act as general chair for the sixth IEEE / ACM International Wireless Communications and Mobile Computing Conference in Caen, France on June 28- July 2, 2010. He will also give a keynote speech at the conference. Helmy will be acting as TPC Vice-Chair for the seventh IEEE International Conference on Mobile Ad-hoc and Sensor Systems.

**BENJAMIN LOK** / *associate professor* / recently joined the steering committee for the IEEE Virtual Reality Conference and was appointed as program chair of the IEEE Virtual Reality 2010 Conference in Waltham, Massachusetts on March 20-24, 2010 and the ACM Virtual Reality Software and Technology Conference 2009 in Kyoto, Japan on Nov. 18 - 20, 2009. Lok was appointed area chair of the ACM / IEEE International Symposium of Mixed and Augmented Reality 2009 in Orlando Fl, from Oct. 19-22, 2009. Lok also became journal associate editor for the "International Journal of Human-Computer Studies and Simulation: Transactions of the Society for Modeling and Simulation International."

**PRABHAT MISHRA** / *assistant professor* / was awarded a $100,000 grant by the National Science Foundation for his research titled "Novel Techniques for Lossless Data Compression and Efficient Decompression in Heterogeneous Embedded Systems." He was also awarded $120,000 (40 percent as co-PI) grant by the National Science Foundation for his research titled "A Multi-Element and Multi-Objective Optimization Approach for Allocating tasks to Multi-Core Processors." Mishra was recently issued U.S. patent 7533294 titled "Functional Coverage driven Test Generation for Validation of Pipelined Processors".

**JORG PETERS** / *professor* / was appointed chair of the conference Shape Modeling International in Beijing June 26-28, 2009

**SARTAJ SAHNI** / *chair & distinguished professor* / was co-chair of the IEEE International Symposium on Computers and Communications, the International Conference on Information Systems, Technology, and Management, and the International Conference on Contemporary Computing. He was named an Honorary Professor of Asia University, Taiwan.

**MARKUS SCHNEIDER** / *associate professor* / received a $911,924 grant from NASA with colleagues from the Jet Propulsion Laboratory at NASA for his project titled "Moving Objects Database Technology for Weather Event Analysis and Tracking." He was also awarded a $354,292 grant from the National Science Foundation for his project "SPAL3D: Design and Implementation of a Type System for Three-Dimensional Spatial Data in Databases".

**MEERA SITHARAM** / *associate professor* / gave the keynote talk at International Conference on Advanced Computing ICAC 09 in Tiruchirapalli, India in August 2009.

**YE XIA** / *associate professor* / received a $375,000 grant from the National Science Foundation for "NeTS:Small: Fundamental Methods and Heuristics for Advanced, Network-Centric Content Distribution."

*Continued from front page*

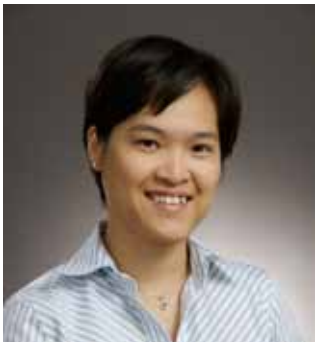

and transmission latency. They have shown several hardness and approximation algorithms for constructing a virtual backbone in the networks, both directional and unidirectional links, provided the first fault-tolerant scheme with a provable performance bound for transmitting data in sensor networks, and proposed the best interference-free broadcast model with a new coloring technique in 3D. Recently, they have provided the first mathematical framework to adaptively update the network community in online social networks and other complex networks, such as biological networks. They have also provided several approximation algorithms to identify the set of disruptors that play a key role in maintaining network connectivity, thus it can serve as a fundamental framework for the design of network topology, network vulnerability and reliability. □

My T. Thai received her Ph.D. degree from the University of Minnesota and joined the CISE department of the University of Florida in 2006. She is also serving as an associate editor for Optimization Letters and Journal of Combinatorial Optimization.

# GATORS @ BLOOMBERG

Bloomberg is about information: accessing it, reporting it, analyzing it and distributing it — faster and more accurately than anyone else. The ever-expanding array of electronic trading, data, news, publishing, innovation and TV and radio broadcast services create transparency and deliver innovation allowing businesses and professionals to transform knowledge into success.

Bloomberg is fast paced, high energy and entrepreneurial. We work hard. We value fresh thinking and initiative. We don't have titles or individual offices. We work shoulder to shoulder in a vibrant, open office environment encouraging interaction, communication, collaboration and teamwork. Glass meeting rooms are a visual reminder that the business is rooted in transparency.

The focus is on building a career rather than climbing a ladder. There is very little hierarchy. Exceptional performers are recognized and rewarded regardless of experience or tenure. The opportunities are limitless for those who have the entrepreneurial spirit, creativity and motivation to contribute to Bloomberg's success.

Since getting its first Gator in 2007 from the CISE Department, the Swamp at Bloomberg has grown wildly and gets larger with each trip to Gainesville. Bloomberg employs more than 20 Gators in the Research and Development Department alone. Most of these are graduates from the CISE department. A recent trip to Gainesville (Sep - Oct 2009), was an opportunity to include focus on other departments like electrical and computer engineering, and industrial and systems engineering. And waiting for these Baby Gators to become a part of the Bloomberg Swamp is really exciting.

Within Bloomberg, Gators have spread out into different teams and work to create next generation fixed Income, equity, trading system, news, messaging and many more applications.

When it is not snowing there are treks to nearby mountains and lots of enjoyment living the New York life. And working hard is natural, as is knowing how to party harder.

In our spare time, Gator football games are always watched and, of course, we root for the Gators. Yes you will see the Gator Bloomberg Team at the Gotham Gator gatherings, too. Go Tebow! Go get another Heisman!

**IN OUR SPARE TIME WE STILL WATCH THE GATOR FOOTBALL GAMES AND ROOT FOR GATORS EVERY WEEK. YES YOU WILL SEE US AT THE GOTHAM GATORS TOO. GO TEBOW! GO GET ANOTHER HEISMAN!**

We would like to thank the CISE Department for everything they taught us and helping us build our current and future careers.

"The one thing I totally believe in is 'The University of Florida is in Gainesville. The Gator Nation is everywhere. I try my bit to build the Gator Nation here at Bloomberg;" said Sumeet Maru. □

# PATTERNS, LANGUAGES, AND TOOLS FOR PARALLEL AND CONCURRENT PROGRAMMING

My group's research aims to make concurrent and parallel programming easier: We want programmers to be more productive, we want produced code to have fewer bugs, and we want the software to be more efficient to make better use of the available resources.

**PROFESSOR BEVERLY SANDERS**

## PATTERNS FOR PARALLEL COMPUTING

The goal of parallel computing is to speed up a long computation by allowing multiple processors to work on it simultaneously, and/or to allow larger problems to be solved by utilizing the memory of a large number of processors. It is the job of the programmer to find a way to decompose a problem into tasks that can be performed simultaneously in parallel on multiple processors, and then construct a program that does this correctly and makes effective use of the hardware.

## TO TAKE A FAMILIAR EXAMPLE, CONSIDER THE PROBLEM OF MATRIX MULTIPLICATION:

$$C_{i,j} = \sum_{k=0}^{N-1} T_{i,k} A_{k,j}$$

There are several ways that matrix multiplication could be decomposed. One way might be to think about computing each element $C_{I,J}$ in parallel, then each task would need a row of T and a column of A. This obvious solution seems fine at first glance and is easy to implement in a shared memory computer, but it turns out the performance is not good. The problem is that every multiplication operation requires reading a new value from memory and each element will need to be read many times by different processors. Since memory access is slow compared to floating point operations, memory access becomes a bottleneck. A better solution is to decompose the three matrices into submatrices, or blocks, and let the parallel tasks update the blocks of C. This allows data in caches to be reused, and the decomposition works especially well when the size of the blocks matches the size of the caches on the machine. The algorithm to compute the blocks has the same structure as the algorithm above with blocks replacing individual numbers and ordinary multiplication of two numbers replaced by a matrix multiplication of the blocks. If the matrices are too large to fit into the memory of a single processor, then breaking them up into blocks also turns out to be a good way to decompose them so the underlying communication system is used effectively.

Matrix multiplication is such an important problem that good solutions for its parallel implementation are well known. Nevertheless, there remain many issues. How can known solutions to commonly arising problems be communicated effectively to novice parallel programmers? How can solutions that only make sense in certain contexts be effectively communicated? Are there useful approaches for finding solutions to novel problems? How much of the work can we get the compiler and/or runtime system to do for us?

One approach to the communication problem is to describe solutions using patterns. Patterns describe good solutions to frequently occurring problems including the context where the solution might be applicable, the tradeoffs involved, a clear description of the solution, examples, and known uses. Drawing inspiration from the success of design patterns for object-oriented programming, we have attempted to codify a core of knowledge of good parallel programming practices into a set of patterns for parallel programming. Taking the work a step further, the patterns have been organized into a system called a pattern language. The Pattern Language for Parallel Programming is divided into four groups of patterns taking a programmer through the design process of a developing a parallel program, from finding the concurrency in a problem, to choosing an algorithm structure that matches the problem at hand, to then finding the way to organize the code, and to the low level implementation mechanisms. The pattern language was described in the book "Patterns for Parallel Programming" with Timothy G. Mattson and Berna L. Massingill, published by Addison-Wesley.

## A BLOCK-ORIENTED PARALLEL PROGRAMMING LANGUAGE

Once patterns are identified, they can be used directly as a resource by programmers. They can also be used as a starting point for tools and special purpose programming languages. A major effort in our groups is collaboration with computational chemists working on parallel software (see www.qtp.ufl.edu/ACES) to perform electronic structure calculations. These computations, which give information about the behavior of molecules that would be expensive or impossible to obtain experimentally, utilize extremely large multidimensional arrays. These arrays are usually too large to fit in the memory of a single processor, and may even need to be stored on disk. Like matrix multiplication, in this domain, decomposing the data into blocks is extremely important. To support this, a domain specific programming language has been developed where scientists can express their algorithms directly in terms of blocks. The runtime system for the language is specialized to handle large multidimensional arrays and takes care of managing the distributed data. The programmer needs only ask for a particular block of an array — the runtime system locates the block and brings it to the location where it is needed. While relatively transparent ways of handling distributed data have been found in other systems, the block-oriented programming

language is unique, and in combination with the distributed array management makes it easy for the runtime system to efficiently execute the program and overlap communication and computation. While a processor is computing with one block, the next block that will be needed is being fetched so that it will already be available when the processor is ready to start work on it. Because the data management, communication, and scheduling is handled effectively by the runtime system, the work of the scientists developing the applications programs, which are quite complex and require complicated tradeoffs between computation, memory usage, and communication, are significantly simplified. The person developing and tuning the runtime system needs to be an expert parallel programmer; the person developing the scientific application does not. Time to develop well-tuned programs has been reduced from several months (using MPI) to a couple of weeks. Porting programs to new supercomputer architecture only requires porting the runtime system; all of the application programs can run unchanged.

As programmers are gaining experience with the language, we are identifying the patterns that lead to the most efficient programs. Use of these patterns will be supported by an eclipse-based IDE that offers various analyses and refactoring operations that will improve the efficiency of a parallel program. A refactoring is an automated correctness preserving transformation of a program to improve its structure or, more typically, its performance. Because refactorings are automated, they can be performed quickly and are guaranteed to be correct. Refactorings can also be easily undone if for some reason the anticipated performance benefits don't materialize. This approach is in contrast to systems where compilers work as a black box. Much work has gone into general purpose parallel programming languages — these require a huge development effort and to date, none have fulfilled expectations. By focusing on a smaller task and doing it well, and proving more information and support to the programmer, better results can be obtained.

## MULTI-CORE SYSTEMS WITH RELAXED MEMORY MODELS

Most PCs for sale today include at least two cores and any program with a GUI is a multithreaded program, whether the programmer realizes this or not. Unfortunately, just as concurrent programming is becoming ubiquitous; it is also becoming more difficult. Advances in computer architecture and compiler technology that have led to the amazing increases in effective processor speed over the past decade make life difficult for the programmer of concurrent programs.

When a programmer writes two statements in a program, he or she usually assumes they will be executed in the order that they were written. For example, one might write

```
result = computation();
done = true;
```

where the value of the variable *done* is intended to indicate that the computation has finished and the result is ready. However, this may not be the case in reality. If this were sequential program where the computation does not access the variable *done* (or generate output), then the order of the two statements could be reversed without changing the visible behavior of the program. However, most likely in a multithreaded program, the done variable is intended as a flag to indicate the computation is finished, and reordering the operations could result in a different thread accessing result before the computation is finished. In general two threads are not required to see the same writes in the same order unless the programmer has inserted sufficient synchronization. On a single processor machine, even though the reordering is allowed, it most likely wouldn't actually happen. The situation is different on a multicore machine. Such errors may indeed manifest themselves, although only intermittently, and thus are very hard to reproduce. There is a significant amount of existing code that seems to be fine on a single processor, but may fail due to a memory model related error when executed on a multicore machine. Exactly what form the necessary synchronization must take varies according to the programming language and is specified by the language's memory model. In Java, assuming that another thread only tried to read results after it has read done and found its value to be true, it would suffice to mark the declaration of done with the volatile keyword, thus signaling to the compiler this variable is likely to be read by another thread. Unfortunately, the subtleties of the Java Memory Model are difficult to understand and it isn't always easy to determine if a program is correctly synchronized.

We are working on a type system that would allow the compiler to find incorrectly synchronized program. We have also developed a tool called Java RaceFinder (JRF) that analyzes Java programs and finds situations with inadequate synchronization that may result in memory model related errors. The tool is an extension of a tool called Java PathFinder developed at NASA to analyze Java programs using a technique called model checking. In our experiments with the JRF, we have found memory model related errors (data races) in programs from a concurrent programming textbook, in open source concurrent utilities, and in the Java Grande Forum benchmark suite. A recent addition to JRF allows it to analyze the information that was acquired during analysis to discover a data race and suggest corrections to the programmer. □

# BUILDING & DYNAMIC *customized* META-QUERIERS

PROFESSOR RANDY CHOW

The Internet has been a great source of information for daily life. However, the explosion of information has made it difficult to efficiently retrieve useful information, especially from on-line databases searchable only through HTML query interfaces. Studies have shown building meta-queriers is the only practical solution to searching effectively the contents from multiple on-line databases. A meta-querier refers to a system that provides a global query interface to multiple online databases through their own local query interfaces (see Figure 1). With a meta-querier, users only need to input a query in the global query form, which is respectively translated to multiple local queries with the same semantics (called query rewriting), and then the combined local query results are returned (called result rewriting). For example, kayak.com is a very popular meta-querier for searching airfares.

In general, meta-queriers are designed for specific user requirements to integrate a set of data sources. The differences in user interests and preferences require different functionalities and source selection. For example, zoomtra.com is a meta-querier for searching airfares. Although kayak.com is more popular in general, zoomtra.com provides cheaper air tickets to India most of the time. Thus, users who plan to travel to India by air prefer zoomtra.com than kayak.com. The difference between the two travel sites is mainly due to their different selection of data sources. Thus, it is desirable users have the capability of building their own meta-queriers with the freedom to select their preferred data sources (i.e., customizability).

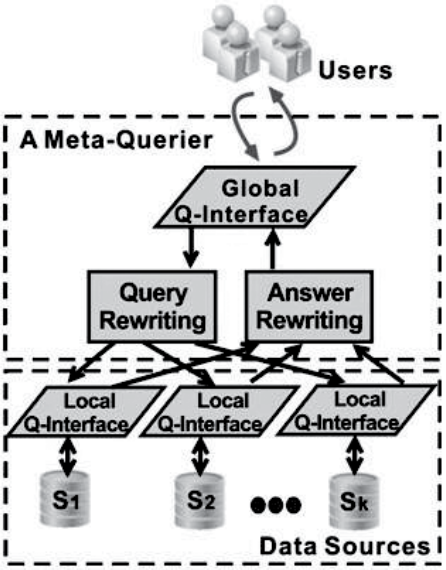

Figure 1: The Structure of a Meta-Querier.

Furthermore, user-specified source selection is not static but dynamic, and the query interfaces of data sources also are evolving. Therefore, meta-queriers should be allowed to dynamically evolve to adapt to the changing user needs and source interfaces (i.e., dynamic re-configurability).

In the academic community, much effort has been devoted to building a single meta-querier, but the existing approaches are impractical when the number of the systems becomes large due to increasing data sources and varying user needs. To address the construction of a large number of dynamic and customized meta-queriers, our group revisits the traditional research problems with different requirements: a) how to construct and maintain global interfaces from a set of local interfaces of user-specified data sources; b) how to discover mappings from global interfaces to local interfaces; c) how to store and manage the existing global interfaces and mappings.

Facing these challenges, three strategies are incorporated into our solution approach:

Mass collaboration: It is impractical for a small set of experts to build such a large number of systems. Instead, this workload of a small set of expert builders should be distributed to a considerable number of cooperative members in an application community, which can include both users and domain experts. In this setting, an intelligent platform is desirable for efficient communication among community members.

Reuse-oriented construction: Systems should be built holistically, rather than individually. Its main benefit is to decrease human workloads by avoiding repetitive tasks and reusing shared components. Therefore, a mechanism is required for organizing and managing the unstructured shared components for better reuse.

Mass customization: Grouping similar user needs can decrease the number of customized meta-queriers that need to be built. The corresponding models and algorithms are necessary for representing and clustering user needs and preferences.

Following the above three strategies, our group proposes a complete solution for holistically constructing and maintaining a large number of customized meta-queriers with the assistance of a domain-specific community. The core of our solution is a central knowledge base (called as M-Ontology) shared by multiple meta-queriers, each of which maintains its own global interface and mapping table. When global interfaces or mappings need to be created or updated, these meta-queriers can utilize the services provided by the schema merger and matcher respectively (see Figure 2). Schema merger and matcher are implemented by reusing the information stored in M-Ontology to construct global interfaces and to discover mappings. Interfaces and mappings validated by the community members are automatically recycled to construct a new version of the M-Ontology for the later reuse. M-Ontology organizes these unstructured shared components according to their semantics for easy understanding and management by the community members. Community members can further enrich the content of M-Ontology via sharing their opinions and information. To alleviate error-propagation and improve system robustness, intervention strategies and version management are also integrated into M-Ontology.

A prototype and simulation of the system have been implemented and analyzed to prove the concept and effectiveness. The system is built on top of an open-source system, Alignment Server. This system provides some basic services such as mapping storage and ontology mapping. In the analysis, we use more than 140 existing web query interfaces from three domains: Books (47), Movies (49), Records (45). These interfaces are represented in HTML and transformed to OWL as input to the M-Ontology. The simulation shows promising results in feasibility (high hit rate in mapping reuse) and effectiveness (high precision and recall in mapping discovery) of the system. □

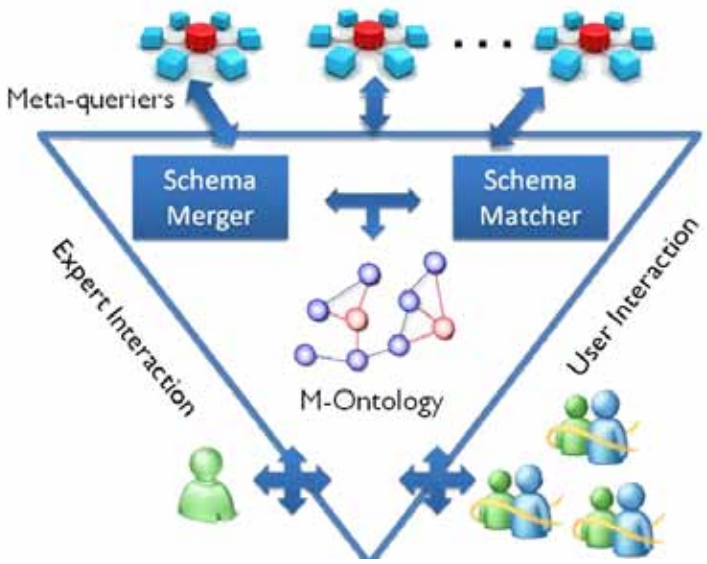

Figure 2: A Community-based and Ontology-based Construction of Meta-queriers.

# CISE CAREER DEVELOPMENT WORKSHOP


The 5th semi-annual Career Development Workshop was held September 28, 2009. It was another success. About 500 students attended to meet 50 representatives from 11 companies. More than 20 student volunteers from ASCIE and ACM assisted with the event which was held at the practice court of the O'Connell Center. This is an increase from the CDW in January when about 400 students and 35 representatives from nine companies attended. This year, about 75 of our students benefited considerably from a presentation given by UF's Career Resource Center about writing resumes and cover letters, and how to best impress a potential employer. Some of the attending companies, such as Harris, Lockheed Martin, Disney, the Florida Interactive Entertainment Academy, and Infinite Energy have been long-time participants. Others, like American Express, Bloomberg, the Nielsen Company, and Raytheon returned. Students enthusiastically welcomed organizations new to the CDW this time: Amazon, and the Naval Surface Warfare Center. □

## SUMMER 2009 GRADUATE AWARDS

**CHEN, MINGSONG**/ *Ph.D. student*/ was selected to attend NSF-SRC-SIGDA-DAC Design Automation Summer School, 2009.

**QIN, XIAOKE**/*Ph.D. student*/ received the Young Student Support Program award to attend Design Automation Conference, 2009.

## SUMMER 2009 STUDENT TRAVEL AWARDS

**BANDYOPADHYAY, SHIBDAS**: "Sorting On A Cell Broadband Engine SPU," IEEE Symposium on Computers and Communications, Sousse, Tunisia, July 05-08, 2009: www.comsoc.org/iscc/2009

**GURUMOORTHY, KARTHIK**: "A Schrodinger Equation for the Fast Computation of Approximate Euclidean Distance Functions," International Conference on Scale Space and Variational Methods in Computer Vision, Voss, Norway, June 01-05, 2009: www.math.uio.no/conference/ssvm2009

**HEO, GYEONGYONG**: "Fuzzy SVM for Noisy Data: A Robust Membership Calculation Method," IEEE International Conference on Fuzzy Systems, Jeju Island, South Korea, Aug. 20-24, 2009: www.fuzz-ieee2009.org

**LI, BO**: "A Refined Performance Characterization of Longest-Queue-First Policy in Wireless Networks," ACM International Symposium on Mobile Ad Hoc Networking and Computing, New Orleans, LA, May 18 - 21, 2009: www.sigmobile.org/mobihoc/2009

**LI, YAN**: "In-Advance Path Reservation for File Transfers in e-Science Applications," IEEE Symposium on Computers and Communications, Sousse, Tunisia, July 05-08, 2009: www.comsoc.org/iscc/2009

**MISHRA, TANIA**: "PETCAM — A Power Efficient TCAM for Forwarding Tables," IEEE Symposium on Computers and Communications, Sousse, Tunisia, July 05-08, 2009www.comsoc.org/iscc/2009

**RAJWADE, AJIT**: "Image Filtering Driven by Level Curves," International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition, Bonn, Germany, Aug. 24-27, 2009: emmcvpr09.iai.uni-bonn.de

**TIWARI, RAVI**: "On approximation algorithms for Interference-Aware Broadcast Scheduling in 2D and 3D Wireless Sensor Networks," International Conference on Wireless Algorithms, Systems, and Applications, Boston, MA, Aug. 16-18, 2009: www.wasaconf.org/wasa2009/index.html

**WANG, WEIXUN**: "Dynamic Reconfiguration of Two-Level Caches in Soft Real-Time Embedded Systems," IEEE Symposium on VLSI, Tampa, FL, May. 13-15, 2009. www.eng.ucy.ac.cy/theocharides/isvlsi09/index.htm

# STAFF SPOTLIGHT

# DAN EICHER

IT Expert.

Eicher moved to Florida from Indiana in 2002. Eicher joined the CISE Department in May 2007 as the senior systems administrator. During his time here Eicher has worked with others in the systems staff to upgrade the mail servers, replace old networking gear, add terabytes of storage, move from Solaris 8 to Solaris 10, upgrade from SUSE 10 Linux to Ubuntu Linux and migrate from in-house backup solution to TSM.

According to Eicher, he enjoys working with Unix, Linux and other open source software. He is teaching himself python and working on migrating legacy code in the Department from perl, ruby and shell scripts to python. Eicher's proudest accomplishment at the CISE Department has been increasing the quality and quantity of systems documentation. This has made it much easier to maintain systems and to increase the rate in which new employees are brought up to speed. Eicher is responsible for the network and system administration of the Department's servers and infrastructure. Eicher is a native of Indiana and graduated from Purdue University in 1993 with a B.S. in computer technology.

# NEW CISE STAFF

Jon Frum, Grants Specialist. Jon joined the CISE staff in August as a grants specialist. He is responsible for working with faculty members to assemble grant proposals as well as compiling budgets and managing the Department grant database. Jon has a B.S. in management from National-Louis University, McLean, Virginia.

John Kramer, IT specialist. Kramer joined the CISE staff in May as an IT Specialist. His primary responsibilities are the maintenance and development of the CISE Web site, Web applications, and the Department's server resources. Though raised in Gainesville, John spent several years in Atlanta, where he earned a B.S. in computer science at Georgia Tech.

Jon Frum

John Kramer

# INDUSTRIAL ADIVISORY BOARD MEETING

October 15, 2009

IAB members attending were: Chris Allen (Tower Hill Insurance Group, Inc.), Doug Pace (Bayshore Solutions, Inc.), Jack Needham (Harris Corp.), Nick Shanks, (Infinite Energy), Michael Parrish and Johnny Morgan (Lockheed Martin Corp.), Rhonda Holt (Turner Broadcasting), Gene Matter (Intel), Tony Barr (Barr Systems, Inc.), and Doug Opfer (Exxon-Mobil Corp.).

Meeting highlights included a College of Engineering overview presented by Dean Cammy Abernathy, a Department update presented by Prof. Sartaj Sahni, a status report from Prof. Manuel Bermudez on the Department's ABET accreditation activities, a technical presentation on simulation and modeling by Prof. Paul Fishwick, and student organization updates presented by members of both ASCIE and the Student ACM Chapter. Dean Abernathy was pleased to note that UF engineering students are currently ranked 2nd in the nation in winning student competitions and that UF Engineering is celebrating its 100th anniversary this year.

The spring IAB meeting was tentatively scheduled for March 25, 2010. □

UF UNIVERSITY of FLORIDA

DEPARTMENT OF COMPUTER &
INFORMATION SCIENCE & ENGINEERING
E301 CSE Building
P.O. Box 116120
Gainesville, FL 32611-6120



# YES! I WANT TO SUPPORT CISE

Amount (please circle):
$1,000 $500 $250 $100 $50 Other $ _____

Please choose a payment method:
Credit Card:

Type: VISA MasterCard Discover

Number: _____

Exp. Date: _____ / _____

Check: Please make checks payable to The University of Florida Foundation and note that your gift is for the CISE Department

_____
Name

_____
Address

_____
City State Zip

_____
Phone Email

Please return to:
University of Florida Foundation
PO Box 14425
Gainesville, FL 32604

Don't forget your company's matching gift form!

**DEPARTMENT OF COMPUTER & INFORMATION SCIENCE & ENGINEERING**
E301 CSE Building
P.O. Box 116120
Gainesville, FL 32611-6120

EDITED BY
**DR. SEEMA BANDYOPADHYAY**
**ANNA HERNANDEZ**

Produced by the College of Engineering
Office of Communications & Marketing