INSIDE CS

Department of Computer Science

FALL 2008

Notes from the Chair

by

Prof. Larry Davis

Midway into the fall semester of 2008, it is clear that this is going to be one challenging year for the University, with severe budget cuts looming both this year and next. Nevertheless, the department is committed to pushing forward along several educational and research fronts. The entire campus is involved in long range strategic planning and the department is participating by reviewing our programs and making some changes in focus and structure. As that plan is developed, it will be posted on our web page.

Prof. G.W. (Pete) Stewart retired this past summer. Pete was with the department for over 30 years, making a series of remarkable contributions to scientific computing. Pete was (is) our only National Academy of Engineering member. He's continuing his research program, but, to his relief, no longer has to attend department meetings. An article about Pete follows in this newsletter.

The department has new associate chairs for both undergraduate and graduate education. Jeff Hollingsworth is the new Undergraduate Associate Chair, replacing Jim Purtilo who is now in the Dean's office working with all departments in the College, not just CS. Howard Elman is the new Graduate Associate Chair, replacing Samir Khuller. I do want to thank Jim and Samir for their work and devotion to our education programs.

The department initiated an annual fall picnic a few years ago to give new graduate students and faculty a chance to meet other department members. Until this year, the weather had cooperated; but this year we had rain on September 26, and had to move the picnic into the lobby of the classroom building. This, at least, kept the food dry but prevented holding our traditional outdoor activities.

A few years ago, to save costs, we made this newsletter online only – but we retained a decidedly 20th century print mentality. That has changed! This newsletter, for the first time, has video content (and it is about time), provided by Amitabh Varshney and describing his lab and research. The newsletter also contains a nice overview of the research being done by our systems group (no video here, unfortunately).

I just returned from a trip to China with Adam Porter and Bobby Bhattarcharjee. We all visited South China Technical University (in Guangzhou), and then Adam and Bobby went on to Chengdu to visit Sichuan University while I went to Hangzhou to visit Zhejiang University. We're hoping that the trip has established some relationships that will lead to some of their best students applying to our graduate program.

I am pleased to announce that the department's Semi-Annual Award for the period January through June 2008 was given to Kathleen Barton, and there is an article in the newsletter about Kathleen and her very interesting hobbies.

Finally, I hope that you enjoy a new segment of the newsletter which focuses on recent UG alumni and what they are doing.

G.W. (Pete) Stewart Retires

Pete Stewart, Distinguished University Professor, retired June 2008 as a full-time Professor; however, he has continued to remain involved in research with the department. The title of Distinguished University Professor Emeritus was conferred by Pres. Mote. A reception was held on September 19 at the University's Golf Course with over 60 people in attendance. An undergraduate research award in his name will be made each year for the next five years with funding provided by the Department and the University of Maryland Institute for Advanced Computer Systems (UMIACS). Pete has a distinguished career with over 34 years being spent in the department. Following is an article run when he was selected by the University as a Distinguished University Professor.

Pete Stewart has been selected as a Distinguished University Professor. This is the first person in our Department to be awarded this honor by campus. We thought you would enjoy reading a brief summary of his work and the impact it has had on the general area of numerical analysis through his many seminal contributions to matrix computations, an area that is fundamental to all branches of science and engineering. Pete's contributions have been recognized by three prestigious honors. He received the F. L. Bauer Prize, awarded by the Technical University of Munich, in 1998. In 2004, he was elected to the National Academy of Engineering, cited "for development of numerical algorithms and

software widely used in engineering computation." Also in 2004, he was named a "Distinguished Editor" of the Journal Linear Algebra and Its Applications.

During his distinguished career, Prof. Stewart has made significant contributions to the following areas:

Mathematical Software. Pete was one of the LINPACK authors, creating the first portable library for solving linear systems of equations and computing matrix decompositions. Its success can be partially judged by the fact that the User's Guide became the all time best seller for SIAM, and by the number of later packages modeled after Linpack.

Matrix Algorithms. Pete produced a steady stream of new matrix algorithms, several of which changed the way computations are done. For example, together with Bartels, Stewart introduced the first stable direct method for solving the Lyapunov equation, $AX - XB = C$, and with Moler, he derived the first stable algorithm for the generalized eigenvalue problem $Ax = \lambda Bx$. Another major contribution is the introduction of a method for detecting near singularity, a method that became the LINPACK condition estimator. These algorithms are now widely used in control theory, signal processing, and other areas.

Error Analysis. Pete's work in rounding error analysis of numerical analysis has been seminal, showing how to extend results to problems for which it was not believed possible. For example, his analysis of downdating broke new ground and is widely cited, and his analysis of subspace iteration, an important method for sparse eigenvalue problems, is the definitive one.

Perturbation Theory. Pete's analysis of invariant subspaces of Hermitian matrices and its generalization to subspaces associated with generalized eigenvalue problems and singular value decompositions is definitive. He was the first to investigate the stability of matrix decompositions, a line of research that has subsequently flourished.

Prof. Stewart's publications are prolific, having published 8 books and 120 papers. In addition to his major research contributions, Prof. Stewart has published a highly successful textbook on matrix computations, an authoritative book on perturbation theory with J-G Sun, and two sets of class notes which have been praised widely. He also has been concerned with the roots of his subject, and has published articles surveying aspects of its history in the nineteenth and early twentieth centuries. He has published a translation from Latin and German of Gauss' later works on least squares, and translations of historically important papers as technical reports. He is currently undertaking a five volume survey of matrix algorithms; the first two volumes are already considered the definitive references on matrix decompositions and eigensystems. All these contributions bear the distinct mark of a truly outstanding scholar.

Computers Help Analyze the Bacteria That Live in Our Bodies

In September, Prof. Mihai Pop received a new R01 basic research grant from NIH . The grant titled "Assembly and Analysis Software for Exploring the Human Microbiome", was funded as part of a new NIH initiative - the Human Microbiome Project **http://nihroadmap.nih.gov/hmp/** - whose mission is to support the analysis of the complex microbial communities inhabiting the human body.

It is a little known fact that our bodies are inhabited by millions to trillions of microbes. In fact, our bodies contain an order of magnitude more microbial cells than human cells, though by weight microbes represent only a small fraction of our bodies. These microscopic hitchhikers play an important, yet poorly understood, role in our health. As an example, the microbes inhabiting the human intestinal tract are essential to our ability to extract nutrition from the food we eat: they help degrade complex starches into simple sugars and provide us with essential vitamins that our bodies cannot produce. Furthermore, there is evidence that our symbiotic microbes can help prevent disease by keeping harmful bacteria at bay.

Studying the microbial communities associated with the human body is difficult because many of these bacteria cannot be easily grown in a lab using standard approaches. New advances in DNA sequencing technologies have provided biologists with a new analysis tool - the ability to characterize the DNA of entire microbial communities without the need to isolate individual organisms - leading to the birth of a new scientific field called metagenomics. Using high-throughput sequencing technologies, scientists can now take a first glimpse at previously unknown organisms and can begin to analyze the complex relationships between the microbes populating our bodies and our health.

Metagenomic studies pose significant computational challenges. These studies will generate large amounts of data (hundreds of gigabytes to terabytes) requiring efficient computational analysis tools. Furthermore, most bioinformatics software tools developed to date are specifically targeted at the analysis of single organisms and are, therefore, ill suited for the analysis of complex mixtures of organisms. New algorithms will have to be developed that take into account the heterogeneous nature of metagenomic data. Such challenges will be tackled by the research proposed in this grant.

Under the guidance of professors Pop and Salzberg, a group of students and engineers at the Center for Bioinformatics and Computational Biology (CBCB) (**http://www.cbcb.umd.edu**) will initially address two fundamental computational tasks: assembly and gene finding within metagenomic data-sets. Assembly is the task of reconstructing the DNA sequence of organisms within a community from the multitude of short DNA sequences generated by sequencing instruments. Gene finding refers to the computational identification of genes within the reconstructed sequences.

Throughout the project, the investigators will continue current collaborations and develop new interactions with biologists actively generating metagenomic data in order to ensure the direct application of the newly developed tools in a practical setting. Like

other tools developed at the CBCB the software developed during this project will be released under an open-source license to the scientific community.

In Focus: Computer Systems Group

From left: Alan Sussman, Neil Spring, Jeff Hollingsworth and Ashok Agrawala. Not pictured: Bill Arbaugh, Bobby Bhattacharjee and Pete Keleher.

Computer Systems provides the foundation upon which all other software applications rely. The goal of systems research is to develop the key abstractions and services that enable software to be efficiently and portably run on hardware. Areas of interest to the systems group include operating systems, computer networks, parallel and distributed computation, and computer security.

The systems group tackles problems from both theoretical and experimental approaches. To support our experimental work, the group maintains several laboratories in the Computer Science department and in UMIACS. The Laboratory for Parallel and Distributed Systems includes a collection of parallel computers and clusters to support systems research. Current equipment includes a 24 processor SPARC SMP, 8 processor IBM Power 4 system, and a 128 processor Myrinet-connected Linux cluster. The Distributed Systems Software Laboratory contains a flexible networking environment for students to configure networking switches for experimental research. In addition, the laboratory includes about 20 machines to support experiments.

The history of the systems group at Maryland dates back more than 30 years. One early member of the group, Yohan Chu, wrote the book "Computer Organization" which was the first major book on the subject. This book was used extensively at many universities and colleges. David Mills was an early researcher in computer networks. He set up an ARPANet IMP (predecessor of the current Internet) in his basement using a PDP 11/45 (an early mini-computer). At the time, this was the only full ARPANet node not located at a University or a Government facility. In the late 1970s, Chuck Rieger and Mark Weiser built ZMOB, an early parallel computer based on commodity microprocessors. The system consisted of 128 Z-80 processors.

Students and postdocs from the systems group have gone on to faculty and industry positions around the world. Recent graduate students in faculty positions include: Gagan Agrawal (Ohio State University), Suman Banerjee (University of Wisconsin), Ugur Cetintemel (Brown University), Ibrahim Matta (Boston University), Bongki Moon (University of Arizona), Ron Larsen (Dean of College of Information Science, University of Pittsburgh), and Sang Son (University of Virginia). Many of our former students have careers at major research labs including AT&T Labs (Vijay Gopalakrishnan , Seungjoon Lee), Google (Ruggero Morselli), and IBM T.J.Watson (Henrique Andrade, I-Hsin Chung, Andrzej Kochut, Kyung Ryu). The group also has a rich history of post-doctoral researchers who have successful careers. For example, Anurag Acharya and Guy Edjlali are now at Google.

Current faculty in the systems group include:

Bill Arbaugh's specialty is information security and privacy. In the past, information systems were large, expensive, barely interconnected, and non-mobile. Today, information systems are small, lightweight, highly connected via wireless technology and mobile. Tomorrow, they will be even smaller and more mobile. The constant evolution in the design and operation of information systems presents new and increasingly complex challenges for computer security. One technology he developed is Copilot, a security and management monitor that is capable of detecting potential intruders in a high assurance manner and ensures that the intruders' activities do not compromise the main system. He is also the CEO of a start-up company, Komoku, Inc.

Ashok Agrawala, an AAAS Fellow, researches the basic nature of information and it's implications to the design and implementation of computer systems. He developed an Information Dynamics Framework, which distinguishes between information and its representation, recognizing that computers only deal with the representations. In 2004, he won the University of Maryland's "Invention of the Year" for Horus Technology. Horus, a novel location determination technology developed with Moustafa Amin Youssef, uses unique algorithms to efficiently process the signal information which is used to determine position.

Bobby Bhattacharjee's research interests are in the design and implementation of wide-area networking, distributed systems, and security protocols. His current focus is on the design of decentralized secure systems for multi-party applications and large scale data distribution, especially in the context of peer-to-peer and overlay systems. His group has build systems that demonstrate protocols for scalable media streaming, randomized resilience, anonymous communication, bulk data delivery, multicast rekeying, distributed directory service, unstructured lookup, secure lookup, and predicate-based search. In current work, he is working on applying techniques from game theory and mechanism design to problems in wireless networking and peer-to-peer systems.

Jeff Hollingsworth's research is in the areas of tools for high performance computing, program instrumentation, and programmer productivity. His work on high performance computing includes tools to measure the performance of parallel programs and Active Harmony, a system to allow runtime automatic tuning of parameters and algorithms for applications and middleware. His work on program instrumentation includes the Dyninst tool suite. Dyninst provides multi-platform instrumentation of binaries via both offline binary editing and online program modification. His

work on programmer productivity seeks to understand how parallel programmers spend their time, and how tools and programming environments can be improved to increase productivity of parallel programmers.

Pete Keleher's research is primarily in the field of distributed systems. His work spans shared memory protocols, adaptive grid schedulers, peer-to-peer systems (including distributed searching and ranking algorithms), and gossip-based consensus algorithms. Currently, much of his work is in the context of storage systems, specifically a wide-area file system called MoteFS. MoteFS is a novel system structured around the concepts of lightweight snapshots and principal-free capabilities. The combination allows extremely efficient and fine-grained control over access to the data storage.

Udaya Shankar's research interests are in the design and analysis of distributed systems and network protocols, in both correctness and performance aspects. His correctness work deals with compositional methods for specification, verification, and testing of concurrent (including distributed) systems, focusing on realistic problems, especially at the transport and routing layers. He's currently working on a "programmer-friendly" compositional framework called SeSF that can be used for system design as well be incorporated in existing concurrent programming languages (Java, C#). His performance work deals with analytical, simulation and experimental evaluation of queuing models of networking systems, under both steady-state and transient conditions. He's currently working on timestep stochastic simulation (TSS), developing a method to compute sample paths of general queuing networks with state-dependent delayed feedback (e.g., TCP/IP/WLAN networks) with the accuracy of packet-level simulation but at a cost several orders cheaper.

Neil Spring's current research focuses on the design of network protocols that are self diagnostic: that embed features that expose and explain faults directly to users and administrators. Neil uses implementation as a primary means of evaluating his research ideas: he has constructed software to measure ISP network topologies efficiently and accurately, software to support arbitrary but safe network measurement, and software that couples instrumentation probing with data transfer in the nearly-ubiquitous TCP. To provide useful conclusions about network structure and behavior requires the aggregation of various sources of information: routing protocol information from BGP, information embedded in host names in DNS, information returned by routers in packet identifiers, source addresses, and record route entries, etc. Each source of information can, by virtue of varied implementation and configuration decisions, provide false data; a current challenge is building the reasoning logic to find the best explanations for measured Internet data.

Alan Sussman's main research area is in software tools for high performance parallel and distributed computing, which is now widely known as Grid computing. Within that broad area, the InterComm project is investigating interoperability of parallel (and sequential) programs, in particular how that can be applied to complex coupled physical simulations. That interest ties in closely to related interests in software component technologies, in particular how they can be applied to high-end supercomputing applications. Another major research interest involves various types of runtime and compiler support for high performance data intensive applications, and their relation to high performance database systems. A third major research area is in applying peer-to-peer computing techniques to high-end computing problems, particularly focusing on utilizing desktop computers effectively to perform large-scale computations. All this work is strongly motivated by high-end applications, with a current emphasis on space science, astronomy and earth science applications, enabled by collaborations with scientists working in those areas.

The Systems group receives support from the Department of Defense, Department of Energy, NASA, and the National Science Foundation. Additional support is provided by industrial partners including DoCoMo, Fujitsu, IBM, Microsoft, Samsung, and Sun Microsystems.

Kathleen Barton Wins Semi-Annual Staff Award

Kathleen Barton is the coordinator of the department's state, DRIF, and faculty start-up accounts. She has been with Computer Science since August 2005 having been promoted in various past positions. Her nomination cited her excellent work while managing faculty accounts, salary obligations, and processing purchases/reimbursements. Kathy carries a heavy load and has done so with patience and professionalism while remaining helpful and pleasant to all.

Kathy and her husband, Dale, share a home with their 4 children (ages 16 to 21 years of age) and a menagerie of animals. Dale is a Master Sergeant in the U.S. Army who repairs musical instruments for the Army's field band. He has been in the army for 27 years, first as a musician and then gravitating to his current position. When both have time, and are not scooping 'poop' (lots of animals!), they are absorbed in the repair business…home repair (they replaced a roof recently) and instrument repair. Kathy is learning to repair wooden instruments and she especially likes working on clarinets and flutes. Dale also builds guitars and teaches musicians how to repair their own instruments so that if a problem develops while on tour, these musicians can make a quick fix until they arrive home and can turn their instruments over to the expert.

We guess it's not a surprise that Kathy has so much patience given her many interactions with all forms of human and animal pedigrees! Kathy says the only time she gets to sit down is when she comes to the office. We are thankful that she has chosen to rest her (fill in your own word here) in our business office.

Professors Visit China

In late October Larry Davis traveled with Bobby Bhattarchjee and Adam Porter to Guangzhou to meet faculty and students at South China Technical University and lecture about their research at Maryland. The goal was to generate interest from their best students in attending graduate school at Maryland, and forging ties with their faculty.

The lectures were given at a new "University City" that houses campuses for almost all of the colleges and universities in Guangzhou. The visit was arranged by Michael Ma, a Professor of Etymology at Maryland and a strong proponent of closer US-China educational and research ties. After spending three days in Guangzhou, Bobby and Adam went on to Chengdu where they gave another series of lectures at Sichuan University, a university that the department has now worked with for two years in a faculty visit program. Larry and Mike went on to Hangzhou to meet the honors students at Zhejiang University. The department is planning a summer internship program to bring some of the best undergraduate students identified during the trip to the department for 10-12 weeks next summer.

And...ACTION!
The Flexiview Project

Amitabh Varshney was the principal investigator of a recently sponsored DARPA seedling project at Maryland, called Flexiview, that investigated how real video and 3D graphical models of urban environments could be fused. The other senior faculty on the project were Rama Chellappa (UMIACS and Electrical and Computer Engineering) and Larry Davis. The project's technical goals were to detect and track people in surveillance video, and measure their body movements (walking, bending down, etc.) so that virtual humans that would move and behave like the real people in the surveillance video could be injected into the graphical urban model. Viewers could then control virtual camera parameters to re-create the actions in the real world from new, arbitrary perspectives. The demonstration video explains their technical approach and illustrates the system in action.

http://www.cs.umd.edu/newsletters/video/2008f/Flexiview.shtml

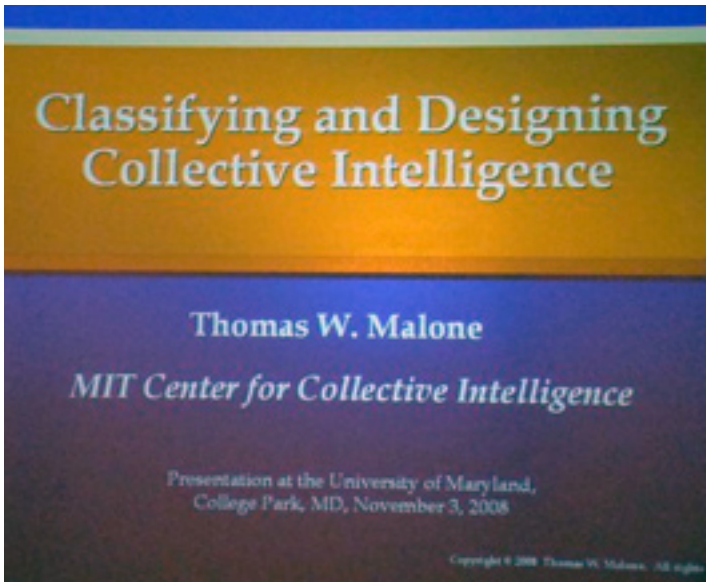

Distinguished Speaker Series

Prof. Tom Malone, MIT, gave a Distinguished Lecture on November 3. His talk can be viewed at **http://www.cs.umd.edu/talks/**

Ben Shneiderman and Larry Davis talk with Tom Malone

In The News...

Effective January 1, 2009, Prof. Ashok Agrawala will become a Life member of the IEEE, a distinguished status reserved for those with a long association with the IEEE. Life members contribute greatly to the fortunes of the IEEE through their service on committees and boards, involvement with local activities, and their contributions to the technical vitality of the IEEE.

Bonnie Dorr, and her students Matthew Snover and Nitin Madnani (in collaboration with Rich Schwartz at BBN Technologies) participated in the first ever NIST Metric MATR workshop to evaluate and compare automatic machine translation evaluation metrics. Their submission, TERp (Translation Edit Rate plus), was noted for its ability to automatically predict the quality of a translation. TERp was one of the top performing metrics at the workshop, and had the highest Pearson correlation coefficient, with human judgments in 9 of the 45 test conditions---more than any other metric. In addition, in 33 of the 45 test conditions, TERp was statistically indistinguishable from the top metric---again more than any other metric. Overall, TERp was consistently one of the best performing metrics in the workshop.

Amitabh Varshney has been elected as the Chair of the IEEE Visualization and Graphics Technical Committee (VGTC) for the 2008-2011 term. VGTC provides technical leadership and organization for technical activities in the areas of visualization, computer graphics, virtual and augmented reality and interaction.

Ben Bederson was quoted on September 18 in an article on America.gov, about the collaboration with children on the design of the International Children's Digital Library. The July 08 edition of the Review of Policy Research reviewed Bederson's recently published book "Voting Technology: The Not-So-Simple Act of Casting a Ballot."

Richard Upton, 1984 Ph.D. Computer Science (advisor Satish Tripathi), is Executive Vice President, Healthcare Informatics at BBN Technologies, MA, and responsible for developing new business areas. Prior to joining BBN, Upton was responsible for all business development and strategy for a BAE Systems North America line of business as Vice President, Business Development, National Security Solutions.

Neil Jhaveri, 2007 B.S. Computer Science and Finance, has been named a 2009 Siebel Scholar. The Siebel Scholars program was established by the Siebel Foundation in 2000 to recognize the most talented students at the world's leading graduate schools of business and computer science. Presently Jhaveri is a graduate student at Harvard University, School of Engineering and Applied Sciences.

V.S. Subrahmanian, Computer Science and UMIACS, was quoted in the Daily Star (Lebanon), October 22, in an article on the development of a software tool that can forecast the behavior of groups branded as "terrorist."

Congratulations to Matt McCutchen, Mitchell Katz and Alan Jackoway for placing first in the Mid-Atlantic ACM programming contest.

Qiang Yang, a 1989 PhD graduate of our department, has been elevated to IEEE Fellow.

Stanley Dunn, 1983 M.S., 1985 Ph.D. Computer Science (advisor Larry Davis) was named Vice Provost and Dean of the Graduate School at Rensselaer Polytechnic Institute in Troy, NY, effective August 1, 2008.

For more information on department activities, please visit our web site at ***http://www.cs.umd.edu/***

To subscribe or unsubscribe to the INSIDE CS newsletter, go to ***http://mailman.cs.umd.edu/mailman/listinfo/csnewsletterlist***