INSIDE CS

Department of Computer Science

SPRING 2009

Notes from the Chair

by

Prof. Larry Davis

An item hot off the press is that the Department's undergraduate programming team, coached by Amol Deshpande, went to Stockholm to compete with squads from all over the world in a programming contest known as the "Battle of the Brains." Our team placed 20th, and only three other U.S. universities did better or as well – MIT placed 7'th, CMU 12th and Stanford tied with us for 20th. The top ten teams were mostly from Russia and China. The students on the UM team were Alan Jackoway, Mitchell Katz and Matt McCutchen. This is a tremendous accomplishment, plus a good time in Europe.

I attended the campus's annual alumni Gala just the other night. At this event, each College on the campus names their alumni of the year, and then there are campus wide awards for international alumni and young alumni. Remarkably, four out of twenty one award winners were computer scientists. Prof. Chan Park, now an advisor to the President of Korea, was formally a graduate of the Engineering College (and was the international alumni of the year – three of the past four international alumni of the year have been CS alumni!). Prof. Park did most of his research at Maryland on computer vision with Prof. Azriel Rosenfeld, and then spent his entire academic career in computer science. Dave Bagget, who was a double major in Linguistics and CS, and a member of the Department's Alumni Hall of Fame, was the Alumni of the Year for the College of Arts and Humanities; two brothers who cofounded web.com – Haroon and Zekeria Moktarzada – were honored as outstanding young alumnus across the entire campus. Zekeria is a B.S. alum of our department, and thanked the late John Gannon for giving him an opportunity to study CS in spite of sub-optimal grades. Finally, Anh Duong was a double major in CS and Engineering, and now has a high level position with Homeland Security. She was named as the Engineering Alumni of the Year.

Next year will see some important enhancements to our undergraduate education program. First, we have funding to build an undergraduate education laboratory in multi-core computing; it will contain a variety of "conventional" multi-core machines and GPU's. We will be developing a set of 1 credit courses that not only our students, but other students in science and engineering, would take to learn basic programming skills for these machines, and we are developing a more advanced senior level course on multi-core computing for our majors. Our department has a substantial amount of research activity in GPU systems; Nvidia named UMD as a center of excellence in that area and funds a number of research projects on Scientific Computing (Prof. Duraiswami), Geographic Information Systems (Prof. Samet) and Graphics (Prof. Varshney). There is an article in the newsletter on Prof. Varshney's GPU activities.

Additional articles in this newsletter describe other interesting and noteworthy activities of department members. If you've done something that you'd like to share with other alumni, please drop me a note and let me know.

Large-scale Testing and Analysis of Highly-Configurable Software Systems

Modern software systems are designed to be highly configurable. From a business perspective this is certainly beneficial. It allows generic systems to serve many different customers by tailoring the system to the customer's individual run-time environment and usage scenarios. For developers however, such configurability makes understanding, testing and optimizing their systems much, much harder. That's because current software development tools and approaches, designed for yesterday's single-configuration systems, do not scale to the enormous numbers of configurations found in today's systems.

Computer Science Professor Adam Porter and his colleagues are working to create revolutionary new tools and approaches to remedy this problem. Dr. Porter has developed a new approach to software testing and analysis that he calls Distributed, Continuous Quality Assurance (DCQA). Inspired by volunteer computing projects such as SETI@Home and Folding@Home-which distribute pieces of enormous scientific computations to computers spread across the world- he is redesigning traditional QA analyses so they can be efficiently run across extensive grids of computing resources in a distributed and continuous manner. This approach greatly improves the quality and speed of QA processes, gives developers greatly expanded insight into system performance on diverse runtime environments and workloads, and allows efficient, coordinated and transparent execution of very large-scale QA processes.

To support this vision, Dr. Porter's research tackles diverse problems such as: formally modeling a system's software configuration space (the set of all valid system configurations), defining efficient

and effective algorithms for choosing which configurations to test or analyze, and creating infrastructure and tools for testing/analyzing selected configurations across large computing grids.

The resulting toolset and approaches have been used to perform sophisticated build, functional and performance testing on large software systems such as the ACE+TAO+CIAO (ATC) open source CORBA implementation and MySQL. The ATC project is a widely used software suite with over 500 configuration options. It currently runs in thousands of commercial, academic and military systems, including DIRECTV satellite TV, the Keck observatory, and onboard the USS Ronald Reagan aircraft carrier. MySQL is one of the world's most popular database management systems. It has well over 150 configuration options, runs on 20+ OS platforms and is used around the world by everyone from individual pc users to the world's largest corporations.

Dr. Porter's research has been funded and supported by the National Science Foundation, the Office of Naval Research (ONR), DARPA, Raytheon, Lockheed Martin, MySQL and other IT companies and software projects.

The Artificial Intelligence Group

Artificial Intelligence (AI) has a long history in our department, and currently supports a very dynamic program of research and education. Our educational curriculum provides a broad range of courses including introductory AI, automated planning, cognitive modeling, commonsense reasoning, evolutionary computation, game theory, machine learning, multi-agent systems, natural language processing, and neural computation. The AI group has consistently ranked high in external national assessments: for example, in the US News ranking of best graduate schools, our AI program is ranked 9th among all universities and 6th among public universities.

Many of our former students have gone on to very high levels of achievement. Examples include Vipin Kumar (PhD 1982), Fellow of the AAAS, ACM, and IEEE; Qiang Yang (PhD 1989), Fellow of the IEEE; Naresh Gupta (PhD 1993), Senior Vice President at Adobe; Lee Spector (PhD 1992), Fellow of the ISGEC; Gary Flake (PhD 1993), Microsoft Distinguished Engineer and well known author of The Computational Beauty of Nature; Narendra Ahuja (PhD 1979), Fellow of the IEEE, AAAI , SPIE, and ACM; and Granger Sutton (PhD 1992), whose software was used in the first-ever assembly of the complete whole genome of a free-living organism, the Haemophilus influenzae genome.

Early prominent members of the AI faculty included Laveen Kanal (heuristic search), Jack Minker (logic and databases), Charles Rieger (common sense reasoning and language), and Azriel Rosenfeld (computer vision). Following are brief profiles of our current faculty:

Bonnie Dorr does research on computational linguistics and co-directs the Computational Linguistics and Information Processing (CLIP) Laboratory. She is known for her research on linguistically-informed, semantically-inspired statistical models that are cross-linguistically applicable, yet practical to train and use. Her research group's systems for summarization, machine translation, and translation metrics have led to first place standings in several NIST evaluation forums, including the Document Understanding Conference, the Machine Translation Evaluation Forum, and the Metrics MATR competition. She is past-president of the Association for Computational Linguistics and recipient of the NSF Presidential Faculty Fellowship Award, the Maryland's Distinguished Young Scientist Award, the Alfred P. Sloan Research Award, and the NSF Young Investigator Award. She is the author of *Machine Translation: A View from the Lexicon.*

Lise Getoor does research in machine learning and reasoning under uncertainty, especially as applied to structured and semi-structured data. She is one of the founders of the statistical relational learning (SRL) research area and is well-known for her work in probabilistic relational models and link mining. Her group does research in machine learning applied to networks and graphs, with a recent emphasis on social networks and social media. In addition, they work on related topics such as data integration and visual analytics. She is recipient of an NSF Career Award in 2008, several best-paper awards, and was a Microsoft New Faculty Award finalist. More information can be found at http://www.cs.umd.edu/~getoor.

Dana Nau does research on automated planning and game theory. He is known for the discovery of "pathological" games (in which, counter-intuitively, looking ahead leads to worse decision-making), and for influential work on the theory and applications of automated planning. The algorithms he and his students have developed for AI planning, manufacturing planning, zero-sum games, and non-zero-sum games have won many awards and championships. His SHOP and SHOP2 planning systems have been downloaded more than 13,000 times and have been used in thousands of projects worldwide. He has more than 300 published papers and several best-paper awards, and is co-author of Automated Planning: Theory and Practice, the standard textbook on its topic. He is an elected Fellow of the Association for the Advancement of Artificial Intelligence.

Don Perlis studies commonsense reasoning—the ability to

muddle through without expertise, even in unfamiliar situations—and the related areas of cognitive modeling and philosophy of mind and language. One of his main ongoing projects is the use of time-situated metacognitive computation for flexible, domain-general, self-adjusting autonomous systems, centered on a "commonsense-core hypothesis": that human-level intelligence hinges on a particular processing technique employing a concise set of basic error-handling rules, and that this can be programmed into automated systems as well. As this work progresses, Perlis applies it to various test domains, currently focusing on human-computer natural-language dialog.

Jim Reggia does research in the area of nature-inspired computing, including neural networks, genetic algorithms/programming, and artificial life. His research group's current/recent work includes such topics as developing integrated neurocomputational systems as models of human language processing and short term memory, symbolically interpreting the representations learned by neural networks, combining cause-effect reasoning with evolutionary search as a creativity-support tool for design, the use of swarm intelligence for problem solving, optimization, and self-assembly, and the use of genetic programming to evolve modular neural networks and self-replicating "machines". He is an elected Fellow of the American College of Medical Informatics and a Senior Member of the International Neural Network Society.

V. S. Subrahmanian works at the intersection of AI and databases, identifying methods by which to perform intelligent reasoning on massive, diversely represented data. His work encompasses logic and probabilistic methods. His work includes algorithms to reason about graph and network data (such as RDF-based semantic web data, algorithms to learn probabilistic models of behaviors from historical information, algorithms to predict possible future scenarios from large histories, as well as methods to analyze opinion expressed in text and extract information scalably from massive text archives and frameworks to identify (probabilistically) activities occurring both in video and transaction data. He has applied these methods to a variety of problems including learning models of terror group behavior and using them to make forecasts. Prof. Subrahmanian is an elected Fellow of both the American Association for the Advancement of Science and the Association for the Advancement of Artificial Intelligence.

In addition to the above faculty, the AI group has close ties with our department's computer vision and spatial reasoning group, where faculty work in research areas that overlap with AI (Yiannis Aloimonas, Larry Davis, David Jacobs, Hanan Samet). Also involved with our AI program are faculty from other campus units, including John Horty (Philosophy), Sarit Kraus (Institute for Advanced Computer Studies, and Bar Ilan University's Computer Science Department), Ugur Kuter (Institute for Advanced Computer Studies), Satyandra Gupta (Mechanical Engineering), Bill Rand (Business School), Phil Resnik (Linguistics), and Amy Weinberg (Linguistics).

Ken Knudsen Wins Semi-Annual Staff Award

Ken Knudsen is the recipient of the Computer Science Department's Spring 2009 staff award. Ken has been the department's webmaster and web applications developer since he started at the University of Maryland in December 2005.

Ken has worked with the majority of department offices helping to revamp web pages and develop electronic systems that have made our work environments more productive. The selection committee commented in its recommendation that the award was based on "Ken's overall work performance, the ease with which he works with staff and faculty members, and his most recent work with the undergraduate office on several projects."

When Ken is not busy with department assignments, he enjoys the opportunity to sharpen his skills as a part-time CS graduate student. In his spare time, he enjoys playing golf, watching sports, making improvements to his condo, and spending time with his wife, Katrina. Katrina can also be spotted around campus; she is employed at the University of Maryland as the Graduate Studies Coordinator for the Sociology Department and is working to finish her M.A. degree in Higher Education. Ken and Katrina are both very happy to be employed at the University of Maryland which also makes commuting to and from work easier.

"Where Are They Now?" The CS Alumni Corner

Dr. Qing Xie, a former CS Ph.D. student (advisor Prof. Atif Memon), is currently employed by Accenture Technology Labs in Chicago, Illinois. In this article, she shares some of her remembrances with us.

"I was recently preparing for a presentation that jogged some fond memories. I was going over the presentation slides for my Ph.D. dissertation, preparing a guest lecture about my research on software testing for a class at DePaul University. Surprisingly, I found

that I could still remember the years-old Ph.D. materials. Just then, it also dawned on me that I am following the same problem solving strategies that I developed during my days in UMD and have applied them in my post-Ph.D. career ever since. This realization prompted me to write this article and share my experiences.

When I first joined the CS graduate program, I did not imagine that software engineering, especially software testing, would be my future research area. In fact, I used to think of software engineering research as not real research: it was just too practical, perhaps occasionally even trivial. If I had kept thinking that way, however, I would not be where I am today. Luckily, three experiences changed my mind. As a fresh graduate student at UMD, I worked at the Office of Information Technology (OIT) as a graduate assistant. My routine job was web hosting system administration. Occasionally, I developed or debugged applications for clients. I constantly received bug fix requests from clients because of the recent changes to their applications. It always took me a lot of time to become familiar with their application, understand their changes before fixing the bugs. At that time, I dreamed of having an automated testing tool that would help me (or the clients) finish these tasks quickly. One year later, I read a report from NIST, stating that software errors cost the U.S. economy $59.5 billion annually. Approximately 80% of the cost of software development can be attributed to developers identifying and correcting defects. One of the key solutions to improve software quality is to improve software testing process, which affected my opinion on software testing. After finishing my coursework, I approached Prof. Memon and asked him to be my research advisor. Prof. Memon shared with me his vision on software testing and explained to me his plan on the GUITAR project (a GUI Testing frAmewoRk). This project drew my attention immediately and I decided to get on board.

Graphical-user Interfaces (GUIs) are ubiquitous. End users interact with the GUI by performing events such as clicking on buttons, typing in text boxes, etc. GUITAR tries to address the challenges posed by GUI testing that traditional software testing techniques do not. The challenges include 1) determining the coverage criteria in terms of events, as a GUI test case is a sequence of events, 2) generating a small subset of test sequences that can detect faults effectively from the enormous input event space, 3) creating expected output that takes the GUI's features (e.g., its layout, its title and color) into account and compares with the actual output, 4) executing test sequences and verifying output after each event and after the entire sequence has executed, and 5) creating a testing process that addresses the agile development nature of GUIs, i.e., rapid prototyping and iterative development and testing. The time spent working on the GUITAR project is one of the highlights of my Ph.D. years (besides free pizzas and ice cream). We (Atif, Ishan, Adithya and I) developed the framework that addresses the above challenges successfully, experimented with real world applications, and published numerous papers. Moreover, I really enjoyed the procedure we used to abstract and tackle the problems, work in a team, and got trained during the process. This training and experience has been life changing for me.

In 2006, due to the success of GUITAR, I was invited to give a talk at Accenture Technology Labs. Months later, I joined Accenture as a researcher in the Systems Integration group right after my graduation. One of the reasons that I came to Accenture was its closeness to practice, which potentially could be a great commercialization path. As one of the largest IT service companies, Accenture has more than 180,000 employees worldwide, and nearly 10,000 among them are test engineers. Our research on systems integration focuses on helping the company and clients to achieve advances in software engineering through the software lifecycle, namely, requirements, design, development, testing and maintenance. My main focus is on software testing and maintenance. I will now describe two projects that were directly influenced by my Ph.D. work.

The first project in which I was actively involved is SMART - SysteM for Application Reference Testing for test cases generation and test harnesses for web services from their reference GUI-based applications. The context for starting this project was that our application renewal group (whose main responsibility is to remedy, re-platform and migrate legacy systems into a new system to meet the goal of IT flexibility) always found it difficult to regression test the functional-equivalent new version of a legacy application against the old one, especially when the two versions have totally different interfaces. Typically, the old version has a GUI front-end, while the new version is a web service which exposes its operations defined in Web Service Description Language (WSDL) file. My intuitive solution to this problem was to apply reference testing techniques developed in my Ph.D. research. Reference testing is a form of regression testing in that the new version is tested against the old version which serves as the baseline (reference). However, the biggest challenge here is the significantly different interfaces. My teammates and I found that the test data requested and responded by web services (target) can be exposed via GUI objects of the GUI application (reference), so we developed a generic and non-invasive mechanism that enables test personnel to extract test data from the reference legacy GUI applications without accessing the source code and the schemas of their proprietary data storages and use this data to generate test cases and test harnesses for target web services. SMART allows users to specify how they use reference GUI applications, and then replay different input data using a prerecorded operational path on these applications, retrieving data from different GUI elements en route. SMART generates unit test cases using these retrieved data to test target web services. Our experimental results have proved that this approach is efficient and effective. This approach has been presented at the International Conference on Software Maintenance (ICSM) 2007; and the tool has been demonstrated at the International Conference on Automated Software Engineering (ASE) 2007. In addition, the company has filed patents for this achievement.

The more we talked with test engineers, the more we knew about the reality of testing practice; we grew more eager to push for test automation. Currently, most testing tasks are performed manually for various reasons such as lack of right tool support, lack of experience and skills of testers on the tools, or simply the uncertainty of the return on investment of the tools. Even for the projects that achieved test automation in which test engineers write programs using scripting languages (e.g., JavaScript and VBScript), and these programs (test scripts) mimic users by performing actions on GUI objects of these applications using some underlying testing frameworks, we found these test scripts are rarely reused on a new release as releasing new versions of GUI applications with modified GUI objects breaks their corresponding test scripts. It typically takes hours to days to repair these test scripts so that they can be tested on successive releases. Test engineers often lack time

and necessary skills to understand and fix old scripts, especially if these scripts were created by other engineers. Existing approaches provide little help to address this pervasive problem. The annual cost of manual maintenance and evolution of test scripts is estimated to be at least $50 million just at Accenture alone. We have proposed an approach for Reducing Effort in Script-based Testing (REST) that automatically identifies changes between GUI objects and locates test script statements that reference these modified GUI objects. The input is GUIs of the successive releases of the same application and the test script for the prior release of the application. These GUIs are compared and modified GUI objects are located. Then, the test script is analyzed statically to invalidate statements that reference these modified GUI objects. This analysis produces warnings that enable test engineers to fix errors in test scripts in a way similar to how compilers issue warnings that enable programmers to fix programs. We have conducted a large scale case study with 45 professional programmers and testers to evaluate the effectiveness of our tool. The results of the evaluation show that users find more broken statements and report fewer false positives in test scripts with our tool than with competitive approaches. A paper illustrated our approach and results of the case study will be presented at the International Conference on Software Engineering (ICSE) 2009. Another paper explains in detail how we infer the types of references to GUI objects in test scripts has just won the best paper award at the International Conference on Software Testing, Verification and Validation (ICST) 2009. Six patents were filed in March 2008.

Both SMART and REST address practical problems that are faced by test engineers today. I have been exposed to many of these new problems only after working in industry; I feel that I may not have seen them had I stayed in academia. I am happy that I joined Accenture because I always wanted to be close to practice. It is very rewarding that that our solutions have been acknowledged by both researchers as well as practitioners. Lastly, I have identified large gaps between the state of the art and state of the practice, especially in the domain of software testing. In the future, I want to help close this gap; I want to bring the state of the practice to researchers to help them to better understand the domain and to attract improved solutions. I also want to bring the state of the art to practitioners to advance their solution domain (GUITAR is on the top of my list). To this end, I continue to be in touch with the GUITAR group, its newly developed solutions, and new directions.

Graphics Processing Units: UMD leading a new direction in High-Performance Visualization and Computing

The University of Maryland is 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 under a $1.1 million infrastructure grant from the National Science Foundation. This project is being led by Amitabh Varshney with participation of 14 other faculty at the University of Maryland. 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.

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 (Amitabh Varshney, Dianne O'Leary, Sergei Sukharev, David Fushman, and Dave Thirumalai) and sequence alignment (Michael Schatz and Cole Trapnell). We are also using the cluster for applications in real-time computer vision (Larry Davis and Rama Chellappa), real-time 3D virtual audio (Ramani Duraiswami), large-scale modeling of neural networks (Jim Reggia), proximity calculations (Hanan Samet), mold design for mechanical CAD (Satyandra Gupta), and computational physics (Bill Dorland and Manuel Tiglio). 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.

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. A number of grants have been awarded based on the cluster. They include DARPA-sponsored research on reconstruction of spatio-temporal 3D environments from multiple videos (Amitabh Varshney, Rama Chellappa, and Larry Davis), Army-sponsored research on computation and visualization of computational fluid dynamics datasets (Amitabh Varshney and Joseph JaJa), NOAA-sponsored research on Climate Modeling (Ramani Duraiswami, Amitabh Varshney, Raghu Murtugudde) and industrial sponsorships from Kitware and NVIDIA. Ongoing research includes using GPUs for simulating trap and nanocomponent interactions for automated nanoassembly (Satyandra Gupta, Mechanical Engineering and Amitabh Varshney). We are also working on visualization and simulation tools for climate analysis and assessing the effects of weather on urban environments with Eugenia Kalnay (AOSC), Joseph JaJa, and Amitabh Varshney. Another very exciting project that has just been funded by NSF is to create tools that expand the computing power freely available to all ATOL (Assembling the Tree of Life), and other phylogenetic researchers. This project is being led by Michael Cummings, with participation by Amitabh Varshney and Charlie Mitter and the central idea is to combine Grid and CUDA computing to take better advantage of a diversity of computing resources, particularly existing desktop processing capacity available through public-computing.

Figures illustrate some of the visualizations being done in the lab:

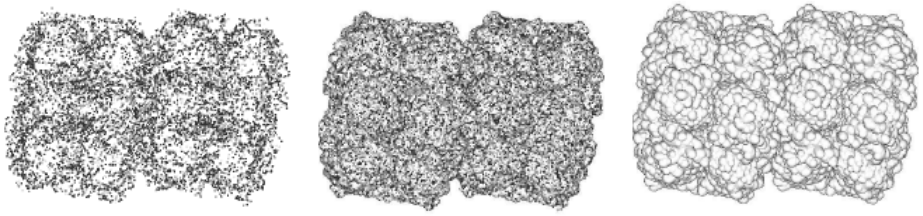

Figure 1: Stochastic Sampling of the molecular surface of protein 1HTO from the Protein Data Bank using GPUs: 10K samples (left), 100K samples (middle), 1000K samples (right) (factor of 8 speedup on GPUs) (Juba and Varshney 2008)

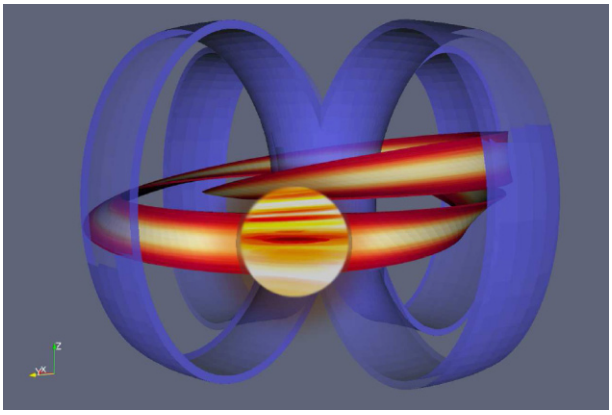

Figure 2: This figure shows the gyrokinetic simulation of the Cyclone benchmark case for ion temperature gradient-driven turbulence in a tokamak and illustrates the geometry of the flux tube domain and features a blown-up region that corresponds to a small ball around the center of the computational domain. We have been developing novel and efficient volume rendering systems using GPUs to visualize large scientific datasets such as these (Stantchev, Juba, Dorland, Varshney 2008).

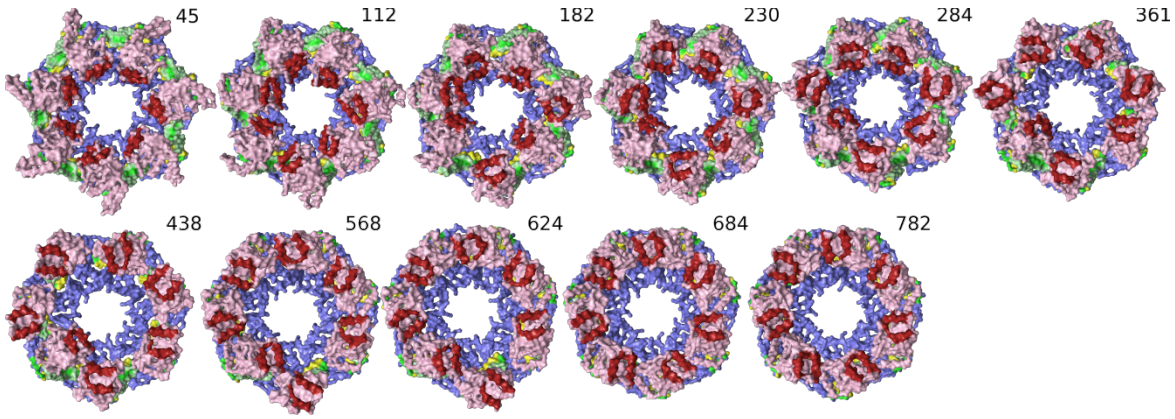

Figure 3: Recent advances in sophisticated computational techniques have facilitated simulation of incredibly-detailed time-varying trajectories and in the process have generated vast quantities of simulation data. In collaboration with scientists from the UMD BioPhysics program (Thirumalai, Hyeon, Sukharev,and Anishkin) researchers from CS (Patro, Kim, Ip, O'Leary, and Varshney) have developed techniques to efficiently summarize large molecular dynamics datasets. The figure shows the 11 most salient time steps as detected by our algorithm from a 834 time step simulation of the R → R" transition of the GroEL nanomachine.

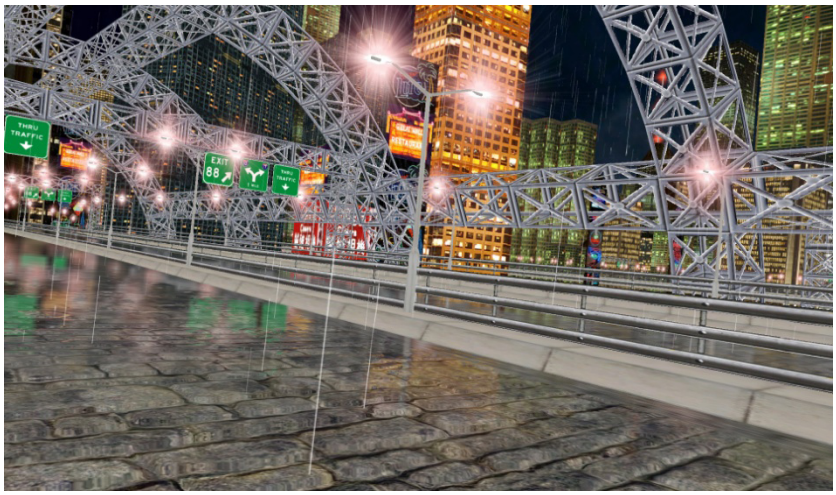

Figure 4: In collaboration with researchers from AOSC and ESSIC we have been developing a visual environment to assess and convey the impact of changes in climate and weather in urban environments using GPU-based codes (Bista, Patro, Varshney 2009)

In The News...

Profs. Howard Elman and Dianne O'Leary and Distinguished University Prof. Pete Stewart have been selected as Fellows of the Society for Industrial and Applied Mathematics (SIAM). Fellowship honors SIAM members who have made outstanding contributions to the fields served by SIAM, and they are among the distinguished members in the initial class of Fellows.

Shomir Wilson, a CS graduate student whose advisor is Prof. Don Perlis, has received an East Asia and Pacific Summer Institutes (EAPSI) award from NSF. He will spend eight weeks this summer in Australia at Macquarie University where he will work with Prof. Robert Dale in the Centre for Language Technology. The primary goals of EAPSI are to introduce students to science and engineering in the context of a different research setting and help initiate scientific relationships that will better enable future collaboration with foreign counterparts. Additional information can be found at: *http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=5284*

Mihai Pop is highlighted in an NSF's Cloud Computing press release and web page at :
http://www.nsf.gov/news/news_summ.jsp?cntn_id=114686&org=NSF&from=news

Steven Salzberg's work on the cow genome was highlighted in Science News – go to the following web site for further details: *http://www.sciencenews.org/view/generic/id/43190/description/Cattle_genome_sequenced* . Steven was also interviewed on WTOP on Sunday, April 26, about the swine flu in Mexico.

Jonathan Katz was one of 12 professors selected to be a member of the DARPA Computer Science Study Panel (CS2P) for 2009. This is a multi-year program, consisting of a funded educational experience to familiarize the participants with DoD practices, challenges and risks, and up to three years of funded research to explore and develop technologies that have the potential to transition innovative and revolutionary computer science and technology advances to the government.

Amitabh Varshney gave an invited CIBM Distinguished Lecture at the University of Wisconsin at Madison on "Visual Insights for Molecular Biology".

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***