Spring 2008 Newsletter

Computer Science @ Virginia Tech

Special points of interest:

- Department Head's Message
- Investment in Excellence and CGI Scholarships
- Faculty Spotlight
- Of Student Interest
- Ali Butt's CAREER Award

If you'd rather receive this publication via email, or if you have an interesting story to tell us and other VT alumni, please send an email to **newsletter@cs.vt.edu**

Spring successes

Welcome to the Spring edition of our newsletter. We are happy to announce two new undergraduate scholarships created through the generosity of CGI and the combined support of other members of our industrial affiliates group. These scholarships recognize the contribution of our alumni to their companies and the potential of our future alumni.

In this issue you will also read about Dr. Ali Butt's NSF CAREER award for his research on advanced storage systems. We are very proud of Ali's accomplishment and know that this is just the first of many successes. Dr. Butt joins eight other faculty who have received NSF's most prestigious early career award.

You will also read about the strong presence of Virginia Tech at Supercomputing 2007, where Dr. Wu Feng and Jeremy Archuleta collaborated with Argonne National Labs and North Carolina State University to win the Storage Challenge Award for their ParaMEDIC framework.

Our department continues its tradition of high impact computing. Two significant examples are in neuroinformatics where General Motors Research is sponsoring Dr. Naren Ramakrishnan's highly successful research in data mining, and in bioinformatics where T.M. Murali collaborated in the study of human pathogens.

I hope you will share our feeling of accomplishment and hope to see you at www.cs.vt.edu for our downloadable, eco-friendly version of this newsletter and other news from your department.

Sincerely,

Dennis Kafura, Department Head

Investment in Excellence and CGI Scholarships

The Department of Computer Science is proud to announce the creation of a general scholarship fund for our undergraduate majors. The initial endowment of $50,000 comes directly from the 22 member companies of the **Computer Science Resources Consortium (CSRC).** The "Investment in Excellence" Computer Science undergraduate scholarship fund is established in recognition of the long-standing partnership between CS@VT and the CSRC. The support of CSRC companies for the programs, faculty and students in the Department of Computer Science is gratefully acknowledged. For over 25 years representatives of CSRC companies have provided valuable input to the department on academics and student programs, have served as role models and given insightful talks to CS student groups, and have supported a range of community-building activities within the department. In addition, many CS alumni work at CSRC companies, and there are several ongoing research collaborations between CS faculty and CSRC companies as well, making the relationships that much stronger. The criteria for the Investment in Excellence scholarships can be found on our scholarship information page. The first awards from this new scholarship will be made for the 2008/9 academic year. **CGI Group Inc.,** among the leading independent information technology and business process services firms in the world and a member of the CS Department's industrial affiliates program, the CSRC , has also created an endowed scholarship for undergraduate CS majors. CGI employees were deeply moved by the tragedy of April 16, 2007, and created this scholarship in memory of everyone impacted by the events of that day. Additionally, CGI wants to encourage students to attend Virginia Tech by helping the College of Engineering recruit and retain highly talented students interested in the field of Computer Science. The first CGI Computer Science Scholarship will be awarded for the 2008/9 academic year.

Spring 2008 Newsletter — Department of Computer Science

Virginia Polytechnic Institute and State University

Christina Daniilidi, M.Sc., Editor and Designer / Dennis Kafura, Department Head

Faculty spotlight

Virginia Tech had a strong presence at **SC/07, the International Conference for High-Performance Computing, Networking, Storage and Analysis**, held November 10-16 in Reno, NV. On Tuesday, November 13, **Ali Butt and Henry Monti** gave a talk entitled "Data Offloading Service for HPC Centers" at the ACM Petascale Data Storage Workshop at SC|07, and **Keith Bisset** spoke about "HPC Modeling of Epidemics" at the Virginia Tech booth. On Wednesday, November 14, **Pavan Balaji and Wu Feng**, as part of the SC|07 Technical Program, gave a talk entitled "Analyzing the Impact of Supporting Out-of-order Communication on In-order Performance with iWARP." **Dimitrios Nikolopoulos** also gave a talk entitled "Polymorphic Parallelization on Accelerator-Based Multi-core Architectures" as part of a Birds-of-a-Feather Session on "Unleashing the Power of the Cell Broadband Engine Processor for HPC." On Thursday 15, **Wu Feng (left) and Kirk Cameron (right)** presented the much-awaited "Green500 List" (photo). In addition, **Pavan Balaji** of Argonne National Laboratory**, Wu Feng and Jeremy Archuleta** of Virginia Tech, **and Heshan Lin** of North Carolina State University, won the Storage Challenge Award for "ParaMEDIC: Parallel Metadata Environment for Distributed I/O and Computing" (ParaMEDIC.) The collaborators leveraged mpiBLAST at Virginia Tech (VT) and MPICH2 at Argonne National Laboratory (ANL) to create the ParaMEDIC framework, which accelerated mpiBLAST by 25-fold over the distributed I/O and computing environment between ANL and Virginia Tech.

The Green500 List, which debuted in November 2007 at the Supercomputing 2007 (SC|07) conference to provide a ranking of the most energy-efficient supercomputers in the world and serve as a complementary view to the TOP500, has, once again, made the news. A cover feature article in *IEEE Computer* in the December 2007 issue is entitled The Green500 List: Encouraging Sustainable Supercomputing. In addition, a cover feature story in the January/February 2008 issue of *IT Professional* talks about how Green Supercomputing Comes of Age. This coming of age is supported by a recent article on the Green Computing movement, which points to the Green Destiny supercomputer as the start of the movement in supercomputing. A feature story by Matthew Dublin entitled "Green HPC has arrived" appeared at **Genome Technology** and includes quotes by Drs. Kirk Cameron and Wu Feng, lengthy descriptions of their energy-saving HPC research projects, including Green Destiny and Tempest, as well as the well-known Green500 list.

General Motors (GM) officials visited the computer science (CS) department at the Virginia Tech's Corporate Research Center on Friday, October 12, to launch the Laboratory for Neuroinformatics, where **Naren Ramakrishnan**, associate professor of computer science, and his students will collaborate with **K.P. Unnikrishnan**, research scientist at the General Motors R&D Center, to create new algorithms that will process the massive amounts of data neuroscientists are now able to collect from the brain. Ramakrishnan and Unnikrishnan will be the co-directors of the Laboratory. Photo, from left to right: *Don Leo*, associate dean of engineering for research and graduate studies at Virginia Tech, General Motors scientist *K.P. Unnikrishnan*, *Jeffrey D. Tew*, GM Technical Fellow and Group Manager at the GM Research and Development Center, *Naren Ramakrishnan*, associate professor of computer science (CS) at Virginia Tech, Virginia Tech Dean of Engineering *Richard Benson*, and CS Department Head *Dennis Kafura* have gathered to inaugurate the project at the CS department's new building in the Corporate Research Center. **Naren Ramakrishnan** also served as the program chair for **ICDM'07, the IEEE International Conference on Data Mining**, held Oct. 28-31 in Omaha, Nebraska. His co-chair was Osmar Zaiane of the University of Alberta, Canada. The conference location enabled partnership by representatives from the banking, insurance, and telecommunication industries.

Manuel Pérez-Quiñones, associate professor of Computer Science (CS), and **Steve Harrison**, CS research faculty, both at Virginia Tech, as well as **Lillian Boots Cassel**, professor of CS at Villanova University have recently earned a National Science Foundation (NSF) - Computer and Information Science and Engineering (CISE) - Pathways to Revitalized Undergraduate Computing Education (CPATH) – Community Building (CB) award of $80,000 for their work focused in advancing computing education in an interdisciplinary fashion. Manuel A. Perez-Quinones also participated as a panelist in the NSF-funded Academic Workshops for Underrepresented Faculty and Senior Graduate Students, scheduled for 11/30/07 – 12/2/07 in College Station, Texas, home of Texas A&M.

A proposal by Drs. **Ing-Ray Chen** (left) and **Chang-Tien Lu** (right), entitled "Intelligent Visualization and Analysis of Traffic Data on the Sun Platform for Transportation Government Agencies", has been selected for an AEG equipment award with a value close to $40,000. According to the Sun Microsystems website, Sun grants equipment to eligible organizations who have developed creative projects that address Sun's investment priorities and create partnerships for success.

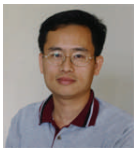

Researchers at the **Virginia Bioinformatics Institute (VBI)** and the **Department of Computer Science at Virginia Tech** have provided the first global analysis of human proteins interacting with viral proteins and proteins in other pathogens. The scientists examined publicly available experimental data for 190 different pathogens that comprise 10 477 interactions between human and pathogen proteins. This approach provides a highly detailed network map of human proteins interfacing with proteins in different pathogens. The network of interactions was published recently in the journal *PLoS Pathogens* and reveals possible key intervention points for the future development of therapeutics against infectious diseases. VBI researchers **Matt Dyer** and **Bruno Sobral**, and Department of Computer Science assistant professor and researcher **T. M. Murali** contributed to the paper, "The Landscape of Human Proteins Interacting with Viruses and other Pathogens." The paper will be featured in the February 15, 2008 edition of *PLoS Pathogens* **4**(2): e32. doi:10.1371/journal.ppat.0040032. The work was supported by Department of Defense grant #DAAD 13–02-C-0018 and National Institute of Allergy and Infectious Diseases grant HHSN26620040035C awarded to Bruno Sobral.

Wu Feng of the Departments of Computer Science (CS) and Electrical and Computer Engineering (ECE) at Virginia Tech, and **Gerald Sabin** and **V. "Nagu" Nagarajan** of RNET Technologies recently received a National Science Foundation Small Business Technology Transfer (STTR) grant of $150,000 for the first phase of a research project entitled "Network Offloading for Genome Sequence Searching using the SmartNIC." **Scott McCrickard**, associate professor of computer science (CS) at Virginia Tech, and **Todd Stevens**, development team leader, Meridium, Inc., the leader in asset performance management software and consulting solutions, also received a National Science Foundation Small Business Technology Transfer grant of $150,000 for the first phase of a software research and development project called "Integrating scenario-based usability engineering and agile software development practices."

Of student interest

Amrita Pati, a Ph.D. candidate in the CS department, has won the third prize in ACM's Student Research Competition at the Grace Hopper Conference this year. She will now compete in ACM-SRC's grand finals scheduled to be held in mid 2008. Amrita works with Dr. Lenwood Heath. Her research interests are in the areas of Algorithms, Computational Biology and Bioinformatics, Graph Theory, and Biological Data Integration.

Twelve women from Virginia Tech attended the **Grace Hopper Celebration of Women in Computing**, which was held from October 17th to 20th in Orlando, FL. The conference was designed "to showcase and model the success of technical women, to develop the next generation of technical leaders and to influence and change the culture of technology," said Telle Whitney, President of the Anita Borg Institute.
Two students were selected to present during the events. Amrita Pati presented a technical poster on *Genomic Signatures From DNA Word Graphs* and Laurian Vega presented at the session *20 Years of Empowering Women in Computing: Systers Past, Present, and Future*. Other students attending were Stacy Branham, Meg Kurdziolek, Tejinder Judge, Sarah Peck.

The **Association of Women in Computing** was proud to award Leadership Scholarships to three deserving applicants at its first meeting of Spring semester last Tuesday, Jan 22. The **Leadership Scholarship** was the brainchild of Meg Kurdziolek and was designed to promote organization activities and create leadership opportunities. AWC members were encouraged to come up with an idea for an event that either brings members together, helps support women in computing, or raises funds for the organization. They submitted their ideas to be juried by senior members, and winners were to receive a $100 cash prize and the opportunity to host their event with the help of fellow members. AWC was proud to be able to award all applicants for their quality ideas; congratulations to the winners **Mara Silva, Kristin Whetstone, and Lise Vincent!**

Matthew Curtis-Maury, a Computer Science graduate student advised by Dr. Dimitrios Nikolopoulos, gave an invited seminar in the Computer Systems Lab (CSL) seminar series at Cornell University on November 20, 2007.

His talk was entitled "Improving the Efficiency of Parallel Applications on Multi-threaded and Multi-core Systems through Prediction-based Adaptation."

"Power consumption is becoming a performance limiting factor in the development of high-end computing systems. Processors containing multiple/many cores are being hailed as a solution to the power crisis, but support for effectively managing parallelism -- in terms of both performance and power consumption -- on these architectures is far from mature. We present a performance prediction model for identifying energy-efficient operating points of concurrency at program phase boundaries in multithreaded scientific applications. We do so via statistical analysis of hardware event counters," said Curtis-Maury.

"I very much enjoyed my visit to Cornell. It was a wonderful opportunity both to increase general awareness of my research as well as to interact with peer researchers at a top university and see what kinds of projects they are working on," said Curtis-Maury.

Emil Constantinescu, Ph.D. student advised by Dr. Adrian Sandu, gave an invited seminar at NASA Goddard Space Flight Center's Office of Global Modeling and Assimilation, on December 18, 2007. The seminar was entitled "Ensemble Kalman filter data assimilation for atmospheric chemistry and transport models."

As Constantinescu explained, data assimilation is a process through which observations or measurements are incorporated into a computational simulation with the purpose of adjusting the results to better represent the reality. The typical example is found in weather forecasting; here long term predictions are unreliable due to the chaotic nature of the atmosphere. Real observations are needed to adjust the simulation variables in order to have a better representation of the atmosphere and thus improve the forecast skill.

"My visit to NASA was a great experience in many aspects. I had the opportunity to present a very important component of my research as well as interact with top scientists at NASA. I was also able to get a glimpse at the state-of-the-art projects that they are working on. Our hope is that the data assimilation techniques developed by our group at the Computational Science Laboratory may become a component used in operational models for climate simulations," Constantinescu concluded.

The **Scholars of the Future** (SoF) project provides research opportunities (through VTURCS) and mentoring to sophomore and junior computer science students, with a focus on under-represented participants. To encourage participation from underrepresented groups, fellowships are provided to the students selected to participate in the program. The students selected for the 2007-2008 Academic Year include, from left to right: Morgan Harris, Lauren Smith, Sergio Bernales, Sarah Dyer, Gregg Tabot, Emily Bollinger, Jaime Garcia, Crystal Kay Weil, Idris Eltahir, Tarish Smith. These students are working on various projects for the following CS professors: Drs. Stephen Edwards, Lenwood S. Heath, Chris North, Manuel A. Pérez-Quiñones, Clifford A. Shaffer, and Layne T. Watson. The Scholars of the Future project is funded by a grant from NSF to encourage underrepresented students to pursue graduate studies in computer science. The project is called Scholars of the Future (SoF) and it is on its third year. It is a collaboration between Auburn University, University of Colorado, and Virginia Tech. The Virginia Tech group is lead by Dr. Pérez-Quiñones, Dr. Pyla in CS and Whitney Edmister (CEED).

On Monday, February 11, 2008, the Computer Science Resources Consortium spring luncheon was held at The Inn at Virginia Tech, followed by CS Career Night later that evening in the Bowman Room of the Jamerson Athletic Center. Undergraduate students involved in VTURCS discussed projects ranging from Cave Visualization in a Virtual Environment to Gigapixel Image Synchronization. Later in the day, students attended CS Career Night in the Bowman Room.

The CS Career Night event, in its seventeenth year, gave CS majors an opportunity to visit with our industrial affiliates to discuss internship, co-operative education, and full-time positions. Seventeen companies were represented at this year's Career Night event, including three of the CSRC's newest members, Bloomberg, Cisco, and MicroStrategy.

The CS Department December Commencement Reception was held Thursday, December 13, in the Torgersen Atrium. CS graduates, their families and CS faculty and staff enjoyed a chance to meet together before the university commencement ceremonies on December 14.

Computer Science @ Virginia Tech

Ali R. Butt's CAREER Award

Ali R. Butt, assistant professor at the Computer Science (CS) department at Virginia Tech, has recently been awarded a **$400,000 National Science Foundation (NSF) Faculty Early Career Development (CAREER) award** for his research titled "A Scalable Hierarchical Framework for High-Performance Data Storage." "The goal of our research is to address the increasing performance gap between computing power and storage technology, especially for high performance computing (HPC) environments," said Butt. Modern scientific applications, such as analyzing information from large-scale distributed sensors, climate monitoring, and forecasting environmental impacts, require powerful computing resources and entail managing an ever-growing amount of data. "In terms of HPC compute power, we are seeing systems with tens-of-thousands or more processors that reach terabyte speeds, and soon will have the capability of processing a petabyte instructions per second. Unfortunately, this is not matched by a corresponding improvement in the input/output (I/O) systems," Butt said.

Ali Butt's CAREER research develops a framework for bridging the said performance gap and supporting efficient and reliable data management for HPC. "The CAREER award will allow me to engage quality graduate students and enable us to innovate, design, develop, and deploy tools and systems for improving the I/O performance of modern HPC setups." The target HPC environments present unique research challenges, namely, maintaining I/O performance with increasing storage capacity, low-cost administration of a large number of resources, high-volume long-distance data transfers, and adapting to the varying I/O demands of applications.

Butt's research addresses these challenges in storage management by employing a Scalable Hierarchical Framework for HPC data storage. The framework provides high-performance reliable storage within HPC cluster sites via hierarchical organization of storage resources, decentralized interactions between sites to support high-speed, high-volume data exchange and strategic data placement, and system-wide I/O optimizations. The overall goal is a data storage framework attuned to the needs of modern HPC applications, which mitigates the underlying performance gap between compute resources and the I/O system. "Our research adopts a holistic approach where all system components interact to give an efficient data management system for HPC," Butt said. "The CAREER will help us realize our research goals, and also serve as a symbol of trust from our peers in the research direction we have adopted. These are exciting times."

Department of Computer Science

Virginia Tech

2050 Torgersen Hall (0106)

Blacksburg, VA 24061