# CS@CU

NEWSLETTER OF THE DEPARTMENT OF COMPUTER SCIENCE AT COLUMBIA UNIVERSITY VOL.8 NO.1 SPRING 2012



# Two Computer Science Professors **Win Sloan Fellowships**

**Xi Chen** and **Junfeng Yang** Named Alfred P. Sloan Research Fellows Two assistant professors of computer science, **Xi Chen** and **Junfeng Yang**, have been named Alfred P. Sloan Research Fellows for 2012. They are among 126 early-career researchers of outstanding promise who have each been awarded a \$50,000 fellowship for use in their research.

"I am very excited to receive a Sloan research fellowship," says Chen. "It is a great honor, and the funding will support both my ongoing and new research projects in algorithmic game theory and complexity theory. I am very grateful especially to my department and colleagues for their great support ever since I joined Columbia two years ago."

Chen is interested in applying concepts and methodologies from game theory and economics

to areas of computer science, to understand how selfish behavior and outside influences impact decision-making processes in the computing world. He studies algorithmic game theory, a new and rapidly growing field, and works on computational problems that arise from the game-theoretic study of the Internet, e-commerce, and other decentralized systems. His current research examines algorithmic issues related to some of the most fundamental models and solution concepts in game theory and economics. Chen, who earned his B.S. in physics/ math and Ph.D. in computer science from Tsinghua University (Beijing), was a postdoctoral researcher at the Institute for Advanced Study, Princeton University and University of Southern California. He won an NSF CAREER award in 2012.

### **Cover** Story (continued)

"It feels great that my research has been recognized by the Sloan Foundation and the fellowship committee," says Yang. "I owe so many thanks to my collaborators, my colleagues at Columbia, and my brilliant students for their constant and tremendous support over the years. I am really thrilled to continue pushing the research frontiers in software systems."

Yang's research centers on creating effective tools to improve the reliability and security of real software systems. While in graduate school, he created an automated and comprehensive approach to detect storage system errors and then went to work at Microsoft Research, where he extended his research to distributed systems on large networks. His work led to numerous patches for the Linux Operating System and a technology transfer at Microsoft. At Columbia Engineering, he codirects the Software Systems Lab (SSL), and is focusing on efficient and reliable multithreading, tools for the cloud, and operating systems support for reliability. Yang, who joined Columbia in 2008, won an NSF CAREER award in 2011 and an AFSOR YIP (Air Force Office of Scientific Research Young Investigator Research Program) award in 2012. He earned his B.S. from Tsinghua University and his M.S. and Ph.D. from Stanford University.

Awarded annually since 1955, Sloan Fellowships are given to early-career scientists and scholars in recognition of achievement and the potential to contribute substantially to their fields. Administered and funded by the Sloan Foundation, the Fellowships are awarded in close cooperation with the scientific community. Potential fellows must be nominated for recognition by their peers and are subsequently selected by an independent panel of senior scholars.

### **Student** Awards

### Ph.D. Student **Jeremy Andrus Wins Facebook Fellowship**

Computer

Science

student

Jeremy

Andrus

has been

named a

Facebook

Fellow.

Jeremy

is one of

Ph.D.



Jeremy Andrus

only 12 fellows tapped nationwide from a field of more than 300 applicants.

This exciting news for Jeremy comes on the heels of winning a best paper prize with his advisor, Associate Professor Jason Nieh. As a Facebook Fellow, Jeremy will receive a \$30,000 stipend to cover study expenses, \$5,000 for conference travel, and \$2,500 for a personal computer.

In addition, the Facebook Fellowship comes with an invitation to visit Facebook headquarters in Palo Alto, CA, later this year to meet one-on-one with engineers working on problems relevant to his research.

"It's an honor to be chosen, and I would hope the recognition can fuel more productive collaboration," he says.

Jeremy is developing an app that would allow two separate personas—say, one for work and one for home-on a single smartphone.

"I would like to refine and unify the mobile computing experience using a system-level approach," he says. "Of course, I'm also interested in fun, exciting, and relevant research."

In its announcement, Facebook posted the following description of Jeremy and his research: "When Jeremy sees many people carrying multiple smartphones, he also sees a challenge: How can we unify all those phones and their functionalities so we need to carry only one device? His goal is to revolutionize the mobile market with virtualization solutions that will allow us to use multiple virtual devices on a single physical device.

Jeremy will work to unify the mobile user experience by investigating efficient and secure ways to virtualize whole platforms like Android. His work will help facilitate the next generation of mobile virtual appliances, open up new opportunities for mobile cloudbased applications and services, and clear our cluttered pockets of excess devices."

Jeremy anticipates leaving academia after completing his Ph.D. and hopes to return to industry to work at a startup or an established organization "I really enjoy creating new technologies that people can use," he says.

# **Engineering Senior Jacob Andreas Wins Churchill Scholarship**

Jacob

won the

Churchill

Scholar-

the first

ship,

has



Jacob Andreas

student from the School to receive the prestigious award. Jacob, a computer science major, is one of 14 students nationwide to receive the Churchill scholarship, which gives exemplary students the opportunity to study for one year at the University of Cambridge with all tuition and

fees paid, including a living allowance and travel and research expenses. Jacob was surprised when

he got the call in early January that he had won the scholarship from the Winston Churchill Foundation. Having misread an earlier email announcing the date on which finalists would be notified, he actually assumed he was already out of the running. The winners were announced on February 6.

"So for the last month or so I've been smiling at inappropriate moments without being able to explain why," said Jacob, 21, from Piedmont, CA. "It's a relief to finally be able to talk about it!"

Jacob, who began programming when he was just in elementary

Columbia school, said he feels "honored Engineerand humbled" to be a part of ing senior this year's group. "Everyone has a really extraordinary record Andreas of research and study."

> But so does Jacob, who is busy this semester working on his senior thesis on a model of machine translation driven by meaning, and as a researcher in the Natural Language Processing Group, under the guidance of his professors Kathleen McKeown and Michael Collins.

"I tell people I've been interested in natural language processing since the fifth grade," said Jacob. The first computer program he ever wrote was a Mad Libsstyle fill-in-the-blank game.

"Language has always excited me," added Jacob, who studied Spanish in high school and Mandarin here. "Computational linguistics really gets to the heart of some of the most interesting problems in artificial intelligence and machine learning that exist today."

Jacob began working in McKeown's lab as a sophomore. He is interested in semantic parsing and generationtransforming text into languageindependent representations of meaning and back again. When he attends Cambridge this fall, he plans to build a better semantic parser under the supervision of Professor Stephen Clark.

"Jacob is one of those amazing students who knew since he arrived what he wanted to work on," said McKeown. "He has published two papers with my group and just had a third accepted. That's in addition to another paper on machine translation that came out of his course work. This is absolutely amazing for an undergraduate, Not only is Jacob creative and thoughtful, he's also a great communicator and a nice guy."

As if there is more that can be squeezed into his day, Jacob also teaches; he has taught as a T.A. in the computer science department's Emerging Scholars Program and for Computer Science Theory, as a tutor, and as a coach on the robotics team at the Engineering School. He held a software engineering internship at Google's New York offices last year, and in 2010, at Microsoft in Bellevue, WA, where he developed several Windows Phone 7 applications using Live Labs technologies

Thinking ahead to life after Cambridge, Jacob said he definitely wants to return to the U.S. for a Ph.D., and after that, intends to pursue a professorship. Indeed, his parents were very pleased to hear of the Churchill news; and perhaps even more so since they both attended Cambridge. Said Jacob, "Perhaps unsurprisingly, most of our discussion so far has taken the form of dating advice.

# **Roxana Geambasu** Joins CS Faculty



Assistant Professor Roxana Geambasu

Roxana Geambasu received a Ph.D. in computer science from the University of Washington in 2011 and a B.S. in Computer Science from the Polytechnic University of Bucharest in 2005. Her research interests span broad areas of systems research, including cloud and mobile computing, operating systems, file systems, and databases, with a focus on security and privacy. She works on identifying and addressing the critical challenges and untapped opportunities created

by today's emerging technologies, such as the Web, cloud computing, and powerful mobile devices. She has won two best paper awards at USENIX Security 2009 and EuroSys 2011 and was the recipient of the first Google Fellowship in Cloud Computing in 2009.

#### CUCS: Tell us about your field of research.

Broadly speaking, my research area is systems, with a focus on security and privacy issues. In my research, I design, implement, evaluate, and sometimes deploy systems that help address some of the critical challenges raised by today's emerging technologies, such as cloud computing and mobile devices. Despite enormous advantages, these technologies threaten users' ability to control the use of their data, its lifetime, accessibility, privacy, management properties, etc. My research focuses on restoring to users control they've ceded to the cloud and mobile devices. As part of my Ph.D., I designed and implemented several systems aimed at re-empowering users with various aspects of control over their data, an effort which

I'm now continuing in my current work. As one example, I designed Keypad, an auditing file system for mobile devices that permits users to track and control accesses on their mobile data, even after a device has been stolen. As another example, I designed Vanish, a global-scale distributed-trust system that allows users to cause all copies of desired Web data objects, online or offline, to simultaneously self-destruct at a specified time.

#### What are some of your ongoing research projects and what motivated you to work on them?

My current work extends my previous efforts on restoring control and transparency over users' data in the cloud and on mobile devices. As one example, I am now working

on an operating system design overhaul for mobile devices with security in mind. Today's mobile operating systems including Android and iOS are descendants of traditional desktop OSes, such as Linux and OSX, and as such lack the appropriate mechanisms to cope with a whole new range of mobile-specific threats, including the pervasive threat of theft. For example, current mobile OSes and applications mismanage sensitive data such as passwords, credit card numbers, emails, and othersby accumulating it in enormous quantities on the device's RAM or persistent storage. This places it in danger if the device is stolen or lost. To secure mobile data against theft, we are now designing and implementing CleanOS, an Android-based OS that maintains itself "clean" at any point in time, keeping sensitive data away from the theft-prone device as much as possible.

I chose to work on this problem because it is important. As more personal or corporate users adopt mobile devices and as the military deploys mobile devices to troops on the ground, we need to ensure that these devices are equipped with the necessary mechanisms to deal with the new threats they present.

#### What attracted you most to a faculty position at Columbia?

Three aspects attracted me to my faculty position at Columbia. First, the undergraduate and Ph.D. students are amazing. I met with some of the systems/ security students during my visits at Columbia and was very impressed by their sharp remarks, fresh research ideas, and commitment to great work. Second, the broad systems group at Columbia is simply amazing. It includes some of the most respected, productive, and fun researchers in operating systems, software reliability, hardware, and security in the country. Apart from the topnotch research they produce, many of these researchers are determined to build systems that are as close to reality as possible, a notion I deeply believe in (please see below). Third, NYC is a really fun place to live :).

#### Do you have a method or approach that you like to use in your research?

The innovative use of cross-field ideas, large-scale design and deployment, and collaborative synergy characterize my research approach. First, my research creatively applies ideas taken from multiple domains to solve complex problems. For example, I integrated ideas from cryptography and databases to address challenges in cloud and mobile computing. Second, my research focus on building, deploying, and measuring real systems at large scale helps to uncover important hidden problems and reveal new opportunities with broad system applicability. For example, my large-scale experience with building the Vanish self-destructing data system yielded insights into the design of a new type of extensible storage system relevant for cloud storage environments like Amazon S3. Finally, my research is invigorated by collaborations with others from diverse backgrounds: professors, industrial researchers and engineers, and students. In the past, I led projects that engaged and leveraged my collaborators' special skills to best meet our research objectives.

### **Feature** Article

# Can you **Trust your Printer?**



Computer Science Professor Salvatore J. Stolfo and his Ph.D. student, Ang Cui, say it is. They've discovered that Hewlett-Packard LaserJet printers, for one, have a major security flaw. Most likely other printers do, too, and this problem is not an easy fix.

"We've discovered a whole new class of flaws that could impact all of us-consumers, businesses, universities, government agencies," Stolfo said. "The primary flaw we found is that the HP 2055DN printer allows any firmware update without checking the authenticity of that firmwarethe firmware isn't digitally signed to determine whether it is legitimate HP-issued firmware or not. This means that the printer's operating system could be entirely replaced by anyone, including malicious attackers."

"While we tested just three models of HP printers, not all models, it's pretty clear that even home printers that aren't directly connected to the Internet are at risk—as long as your printer is hooked up to a computer, it could be used to launch attacks or as part of a botnet. These devices are completely open and ready to be exploited."

Stolfo and Cui were working on developing new advanced firmware update mechanisms. security technology to protect We did a guick scan of unproembedded systems widely tected printers available on the deployed on the Internet, such internet left open to Internet as routers, VOIP phones, and attack and found more than printer firmware, and how they 40,000 that could be infected might inject this new technology within minutes." -called Software Symbiotes-The researchers say that now into existing embedded systems. is the time to alert printer Several months ago, they manufacturers and to get them successfully injected Symbiotes to think more seriously about into Cisco routers to protect their existing security architecthem from hacker attacks.

Imagine your tax return with all kinds of personal information printing out not just on your printer but also on a criminal's in another country thousands of miles away. Is this really possible?

While studying how to embed Symbiotes into HP printers, they discovered the major flaw.

"It is fortunate that we discovered the flaw and alerted HP," Stolfo noted. "We provided the technical details we uncovered, and have offered a number of strategies for HP to develop specific solutions to mitigate the risks. We are looking forward to working with them."

Stolfo's group discovered the security flaw involving the firmware that runs most of these devices. HP LaserJet printers allow firmware upgrades, and each time a printer accepts a job, it checks to see if a software update is included. So Cui figured out how the firmware upgrade feature worked and noted that the printer accepted any firmware that was unsigned. This permits the printer to erase its operating software and install a booby-trapped version.

"This security vulnerability is so fundamental that it may impact tens of millions of printers," said Cui. "And it could affect other devices that use similarly flawed

tures. The primary goal of their research is to identify weaknesses and improve the security of a large body of deployed embedded systems, including printers, which represent only one kind of embedded system "It's conceivable that most of the printers that are currently deployed are vulnerable," Stolfo added. "Printers that permit unsigned software upgrades from print jobs are open to this kind of hack. The question is how can we push out fixes to ensure these devices are not successfully attacked by malicious adversaries? This is a fundamental industry-wide call to arms to protect embedded systems in much the same way our PCs and servers are automatically updated with patches and fixes. Alerting industry to focus on securing embedded systems will ultimately benefit all home and office users, and make the internet safer than it is today."

### Faculty News

# Professor Grinspun one of **"Brilliant 10"**

# Professor Nayar Elected to American Academy of Arts & Sciences



*Popular Science* has named **Eitan Grinspun** among its "Brilliant 10," the magazine's annual list of the top 10 researchers in the U.S.

Associate Professor Eitan Grinspun

Eitan Grinspun, an Associate Professor of Computer Science. studies the basic rules of motion and turns them into computer programs that are animating Hollywood movies and creating new tools for graphic designers. The programs are also used for medical research, and to study problems involving flexible strands, sheets, liquids, and even icicles. These could have applications for the laying down of transoceanic communications cables, design of nanoscale wiring for stretchable electronics, and faster more compact bottling of shampoo and other thick coiling liquids.

"I am honored that the work conducted by my group at Columbia Engineering has received this unique recognition," says Grinspun, who also won a Sloan Research Fellowship in 2010. "It reflects the good fortune that we have had in recruiting top graduate students and researchers to work together, and establishing close collaborations with the best of partners in both academia and industry, collaborations that have enriched our research experience many times over. The ten scientists honored by *Popular Science* are carrying out a diverse range of explorations, and to me the magazine's greatest service with this recognition is the emphasis it places on communicating both the promise and the relevance of far-reaching scientific research."

"We are very pleased that Professor Grinspun has been selected as one of *Popular* Science's 'Brilliant 10.'" says Feniosky Peña-Mora, dean of Columbia Engineering. "He is an outstanding example of Columbia Engineering's exceptional faculty and exemplifies our School's great tradition of excellence and leadership in both teaching and research. His work has had and will continue to have a broad impact on understanding the motion of materials and we are proud of his wonderful achievements."

Grinspun, who directs Columbia Engineering's Computer Graphics Group, part of the Columbia Vision and Graphics Center, has been collaborating with Disney ("Tangled") and Weta Digital ("Rise of the Planet of the Apes"), which use his technology to make animated objects move more realistically and to animate grand, complex scenes that would be very expensive to film using real objects. Adobe Systems Inc. included a new paintbrush tool based on his work as part of its most recent editions of Adobe Photoshop and Adobe Illustrator. "For all of these applications," he says, "you need to understand the motion of materials and how they behave. We're bringing to it our knowledge of computers."

Grinspun, who was born in Israel to Chilean parents and grew up in Ann Arbor, MI, and Toronto, partners with physicists and mathematicians to determine the best formulas to use as a starting point for his work. From there, his research team refines and customizes the formulas they use in their programs. His approach to computing the motion of materials is radically different from his peers and predecessors. He and his collaborators have pioneered the use of a new kind of mathematicsdiscrete differential geometry (DDG)—as a new mathematical "language" in which the behavior of physical materials can be described simply and succinctly and directly translated into fast

computer codes. By using DDG, Grinspun is able to produce simpler, faster algorithms that "get the physics right" with fewer computer cycles.

For over a decade, the editors of Popular Science have been seeking out promising young researchers at labs across the nation. This year's "Brilliant 10" represent the best of what science can achieve and demonstrate America's continuing cutting-edge research. From diagnosing diseases in developing countries to inventing a new kind of geometry, each of these brilliant scientists and researchers are under 40 and offer everyone reason to be optimistic about the future

"Our annual 'Brilliant 10' feature is a testament to the importance of scientific research and a salute to the dazzling young minds driving it," said Mark Jannot, editorial director, the Bonier Technology Group and Editor-in-Chief, Popular Science. "Each year, we solicit nominations from hundreds of eminent scientists and whittle the candidates down to the ones whose work really blows the tops of our heads off. Over the past 10 years, we've celebrated the achievements of 100 scientists who are changing the way we look at, and live in, our world, and I can't wait to see what the next decade brings."



of Computer Science

Shree Navar

**Shree K. Nayar** has been elected to an elite group of accomplished world leaders from academia, business, the humanities, and the arts.

Computer Science Professor

As the T.C. Chang Professor and chairman of Computer Science, co-director of Columbia Vision and Graphics Center and head of the Computer Vision Laboratory, **Shree Nayar** can now claim membership in the American Academy of Arts and Sciences (AAAS).

"I feel greatly honored to be included among such an august group of individuals from the sciences, the humanities, and the arts," Nayar says. "It is a particular thrill for me to be in the same entering class as Bob Dylan, of whom I have been a fan from my college days. I am hoping my chances of getting his autograph have just gone up."

Nayar is among 212 new members who joined one of the nation's most prestigious honorary societies and a leading center for independent policy research in 2011. Members contribute to Academy studies of science and technology policy, global security, social policy and American institutions, the humanities, and education.

Professor Alfred Aho, Nayar's Columbia Engineering colleague and a fellow AAAS member, was delighted to share the news of Nayar's election with the department. He says Nayar is internationally known for his pioneering research into novel digital cameras, physics based models for vision, and algorithms for scene understanding.

"Shree's work has had significant impact on the practice of digital imaging, computer graphics, and robotics," Aho says, "and he recently launched an effort to create an inexpensive educational camera kit, called BigShot, that can be distributed to children around the world to learn about science, art, and culture."

"Shree's interests make him a Renaissance man," said Columbia Engineering Dean Feniosky Peña-Mora. "It is fitting that he has been elected into such an esteemed academy whose members represent the full range of human intellectual endeavors. We are all proud of his accomplishment."

The AAAS was founded in 1780, to "cultivate every art and science which may tend to advance the interest, honor, dignity, and happiness of a free, independent, and virtuous people." Among the Academy's founders are Benjamin Franklin, Thomas Jefferson, and George Washington. Its current members include many of the nation's most prominent computer scientists, more than 250 Nobel and Pulitzer Prize winners, as well as Columbia University President Lee Bollinger.

The 2011 elected class includes astronomer Paul Butler, discoverer of over 330 planets; cancer researcher Clara Bloomfield, who proved that adult acute leukemia can be cured; David Page, whose genome sequencing work has advanced understanding of human reproduction; former Congressional Budget Office Director Robert Reischauer; University of Cambridge classicist Mary Beard; Nobel laureate writer Mario Vargas Llosa; actors Daniel Day-Lewis, Helen Mirren and Sam Waterston; singersongwriter Paul Simon; jazz icon Dave Brubeck; and documentary filmmaker Ken Burns.

#### Faculty News (continued)

# FCC Names Professor Schulzrinne **Chief Technology Officer**

"The communications technology revolution is key to our economy and broad opportunity. With the appointment of **Henning** a world-class technologist—we extend our commitment to technology excellence at the FCC and to strong engagement with outside technology experts." – FCC CHAIRMAN JULIUS GENACHOWSKI

Julian Clarence Levi Professor of Mathematical Methods **Henning Schulzrinne**  Columbia Engineering School Professor **Henning Schulzrinne** has been appointed as Chief Technology Officer by FCC Chairman Julius Genachowski.

"The FCC is engaged in helping bridge the digital divide, increase public safety, protect consumers, and help foster new innovative mobile networks," said Schulzrinne, who is the Julian Clarence Levi Professor of Mathematical Methods and Computer Science, and Professor of Electrical Engineering. "I look forward to participating in these and other activities. I also hope to better connect the engineering community and the FCC, so that we can jointly tackle those important challenges."

"I'm delighted that Henning will be joining us," said FCC Chairman Genachowski. "The communications technology revolution is key to our economy and broad opportunity. With the appointment of Henning a world-class technologist we extend our commitment to technology excellence at the FCC and to strong engagement with outside technology experts."

As Chief Technology Officer, Schulzrinne will guide the FCC's work on technology and engineering issues, together with the FCC's Office of Engineering and Technology. He will advise on matters across the agency to ensure that FCC policies are driving technological innovation, including serving as a resource to FCC Commissioners. Schulzrinne will also help the FCC engage with technology experts outside the agency and promote technical excellence among agency staff. He will be based in the FCC's Office of Strategic Planning and Policy Analysis.

"We are extremely pleased that Henning has been named to such a prominent post," said Feniosky Peña-Mora, Dean of The Fu Foundation School of Engineering at Columbia University. "Not only has he been a leader for many years in the field of computer science and has had a strong academic impact here at Columbia Engineering, he will now play a major role in helping to effect technology policy for both our nation and the world. He exemplifies Columbia Engineering's far-reaching faculty and we are honored to have him as a colleague."

Schulzrinne has been an Engineering Fellow at the FCC since 2010. He has published more than 250 journal and conference papers, and more than 77 Internet Requests for Comment (RFCs). He is widely known for the development of key protocols that enable voiceover-IP (VoIP) and other multimedia applications that are now Internet standards, including the Session Initiation Protocol (SIP). His research interests include Internet multimedia systems, applied network engineering, quality of service, and performance evaluation.

Schulzrinne received his undergraduate degree in economics and electrical engineering from the Darmstadt University of Technology, Germany; his M.S.E.E. degree as a Fulbright Scholar from the University of Cincinnati, Ohio; and his Ph.D. from the University of Massachusetts in Amherst. He was a member of technical staff at AT&T Bell Laboratories, Murray Hill, and an associate department head at GMD-Fokus (Berlin), before joining the computer science and electrical engineering departments in 1996 at The Fu Foundation School of Engineering at Columbia University, New York.



Henry and Gertrude Rothschild Professor of Computer Science Kathleen McKeown

Professor **Julia Hirschberg** 

Vikram S. Pandit Professor **Michael Collins** 

# Three CS Professors Named Founding Fellows of Association for Computational Linguistics

8 CS@CU SPRING 2012

Three Computer Science Department professors are among the 17 founding fellows of the Association for **Computational Linguistics** (ACL). They are Kathleen McKeown, Henry and Gertrude Rothschild Professor of Computer Science and Vice Dean of Research; Julia Hirschberg; and Michael Collins, Vikram S. Pandit Professor of Computer Science. The ACL Fellows Program was created last fall to recognize members whose contributions to the field have been most extraordinary.

"It's truly gratifying to be honored by your peers in this way," Hirschberg says.

Each fellow was nominated for their specific contributions, as noted by the ACL: McKeown, for her significant contributions to natural language generation and multi-document summarization; Hirschberg, for her significant contributions to intonation, discourse, text-to-speech systems, and labeling standards for speech corpora; and Collins, for his significant contributions to natural language parsing and discriminative training. McKeown notes that Columbia had the largest number of founding fellows of any university.

"This is an indication of the comparative strength of the natural language processing team at Columbia, which includes four strong researchers at the Center for Computational Learning Systems as well," she says.

The Association for Computational Linguistics is the international scientific and professional society for people working on problems involving natural language and computation. Its journal, Computational Linguistics, is considered the primary forum for research on computational linguistics and natural language processing.

#### Feature Article

# Professor Yang's **Peregrine System** Improves Software Reliability & Security

Assistant Professor Junfeng Yang

Anyone who uses multithreaded computer programs—and that's all of us, as these are the programs that power nearly all software applications including Office, Windows, MacOS, and Google Chrome Browser, and web services like Google Search, Microsoft Bing, and iCloud knows well the frustration of computer crashes, bugs, and other aggravating problems.

The most widely used method to harness the power we require from multicore processors, multithreaded programs can be difficult for programmers to get right and they often contain elusive bugs called races. Data races can cause very serious problems, like the software bug that set off the 2003 power blackout in the Northeast. Now there is a new system that will combat this problem.

Peregrine, a new software system developed by a team of researchers led by Assistant Professor of Computer Science Junfeng Yang, will improve the reliability and security of multithreaded programs, benefiting virtually every computer user across the globe. Peregrine can be used by software vendors like Microsoft and Apple and web service providers like Google and Facebook, to provide reliable services to computer users. This new research was published in the 23rd ACM Symposium on Operating Systems Principles, considered to be the most prestigious systems conference held each year, and presented by Yang's graduate student Heming Cui.

"Multithreaded programs are becoming more and more critical and pervasive," says Professor Yang. "But these programs are nondeterministic, so running them is like tossing a coin or rolling dice-sometimes we get correct results, and sometimes we get wrong results or the program crashes. Our main finding in developing Peregrine is that we can make threads deterministic in an efficient and stable way: Peregrine can compute a plan for allowing when and where a thread can "change lanes" and can then place barriers between the lanes, allowing threads to change lanes only at fixed locations, following a fixed order. This prevents the random collisions that can occur in a nondeterministic system.

"My research is focused on creating effective

tools to improve the reliability and security

of real software systems. I'm excited about

cyberspace a better place and benefit every government, business, and individual who

uses computers." - JUNFENGYANG

this area because it has the potential to make

"Once Peregrine computes a good plan without collisions for one group of threads," adds Yang, "it can reuse the plan on subsequent groups to avoid the cost of computing a new plan for each new group. This approach matches our natural tendency to follow familiar routes so we can avoid both potential hazards in unknown routes and efforts to find a new route."

Yang notes that in contrast to many earlier systems that address only resultant problems but not the root cause, Peregrine addresses nondeterminism—a system that is unpredictable as each input has multiple potential outcomes—and thus simultaneously addresses all the problems that are caused by nondeterminism.

Peregrine also deals with data races or bugs, unlike most previous efforts that do not provide such fine-grained control over the execution of a program. And it's very fast-many earlier systems may slow down the execution of a program by up to ten times. Peregrine is also a practical system that works with current hardware and programming languages—it does not require new hardware or new languages, all of which can take years to develop. It reuses execution plans, whereas some previous work makes a different plan for each group of threads: as Yang points out, "The more plans one makes, the more likely some plans have errors and will lead to collisions."

"Today's software systems are large, complex, and plaqued with errors, some of which have caused critical system failures and exploits," adds Yang. "My research is focused on creating effective tools to improve the reliability and security of real software systems. I'm excited about this area because it has the potential to make cyberspace a better place and benefit every government, business, and individual who uses computers." Yang's research was funded by the National Science Foundation, including an NSF CAREER award, the Defense Advanced Research Projects Agency (DARPA), the Air Force Research Laboratory (AFRL), and the Intelligence Advanced Research Projects Activity (IARPA).

#### Faculty Awards

# Professor Vapnik **Wins Pair of Major Honors**



Professor Vladimir Vanni Computer Science Professor **Vladimir Vapnik** has won two significant awards for his career contributions to his field.

#### Computer Science Professor Vladimir Vapnik won the 2012 Benjamin Franklin Medal in Computer and Cognitive Science from the Franklin

Institute. The award citation is "for his fundamental contributions to our understanding of machine learning, which allows computers to classify new data based on statistical models derived from earlier examples, and for his invention of widely used machine learning techniques." In recognition of this award Professor Vapnik was honored with a one-day symposium on "Statistical Machine Learning: Theory and Applications" held at the University of Pennsylvania in April.

Professor Vapnik is also the 2012 winner of the Frank Rosenblatt Award, a technical field award from the Institute of Electrical and Electronics Engineers (IEEE). The award is presented for outstanding contributions to the advancement of the design, practice, techniques, or theory in biologically and linguistically motivated computational paradigms, including but not limited to neural networks, connectionist systems, evolutionary computation, fuzzy systems, and hybrid intelligent systems in which these paradigms are contained.

Vapnik is widely recognized as one of the main developers of the Vapnik-Chervonenkis Theory, which is a major component of statistical learning theory, and as one of the originators of the Support Vector Machine, a highly successful and widely used approach in machine learning. Previous awards for Professor Vapnik include the 2005 Gabor Award, induction into the National Academy of Engineering in 2006, the 2008 Paris Kanellakis Award, and the 2010 Neural Networks Pioneer Award.

Vapnik joined the Columbia faculty in 2003, and also serves as a senior research scientist at Columbia Engineering's Center for Computational Learning Systems. He received his master's degree in mathematics from Uzbek State University, Samarkand, USSR and joined the Institute of Control Sciences, Moscow, where he worked from 1961 to 1990, becoming head of the Computer Science Research Department. He then joined AT&T Bell Laboratories in Holmdel, N.J., where he continues as a consultant. He is also a Fellow of NEC Laboratories in Princeton.

### **Faculty** Awards (continued)

# **Professor Hirschberg** Receives Pair of Speech **Research Awards**



Professor Julia Hirschberg has received two awards recognizing her outstanding research contributions in speech and audio processing.

Professor Julia Hirschberg

#### Professor Julia Hirschberg received the International **Speech Communication** Association (ISCA) Medal for **Scientific Achievement for**

**2011** at the annual Interspeech Conference, held in Florence. The award citation reads: "She has made outstanding contributions to text-to-speech synthesis, prosody research, and many other topics in spoken language processing."

Professor Hirschberg was also honored by the Institute of **Electrical and Electronics** Engineers with its James L. Flanagan Speech and Audio Processing Award. The IEEE has given the James L. Flanagan award annually since 2002 to an individual or team for outstanding contribution to the advancement of speech and/or audio signal processing. The group honored Hirschberg for her "pioneering contributions to speech synthesis and prosody research."

"I am particularly honored to receive this award since Jim Flanagan was my first center director when I arrived at Bell Labs," Hirschberg says.

The award citation calls Hirschberg an innovator in building viable computational models of human prosody for use in speech

synthesis. "Her goal," the citation continues, "has been to make human-computer interaction in spoken dialog systems more natural and effective. She was one of the architects of the Tone and Break Indices (ToBI) system for labeling human prosodic contours that is used in many text to speech systems and is widely used in prosody research. Initially used for intonational description of standard American English, ToBI has been extended to model other languages. Professor Hirschberg has also pioneered work in emotional and deceptive speech and on audio browsing and retrieval research, both in techniques to improve audio search and to make audio search engines more usable."

Hirschberg has been a fellow of the American Association for Artificial Intelligence since 1994 and a fellow of the International Speech Communication Association (ISCA) since 2008. She served as President of ISCA from 2005-2007, editor-in-chief of Computational Linguistics from 1993-2003, and co-editor-inchief of Speech Communication from 2003-2005. In 2009 she received a Columbia Engineering School Alumni Association (CESAA) Distinguished Faculty Teaching Award.

# Professor Belhumeur Receives Edward O. Wilson **Biodiversity Technology Pioneer Award**



Professor Peter Belhume

This free mobile app uses visual recognition software to help identify tree species from photographs of their leaves.

Leafsnap contains high-resolution images of leaves, flowers, fruit, petiole, seeds, and bark. It currently includes the trees of New York City and Washington, D.C., and will soon grow to include the trees of the entire continental United States.

Professor Peter Belhumeur's app has generated tremendous media attention, including from the New York Times, NPR, the major television networks and numerous print and online publications around the world.

Professor Belhumeur, who also directs Columbia's Laboratory for the Study of Visual Appearance, has worked on facial recognition since the mid-1990s. In a Columbia News story in May, he said he quickly saw that the same algorithms that can process the curve of an eyebrow or the angle of a cheekbone could be applied to the shape of a leaf.

The American Computer Museum has named Computer Science Professor Peter Belhumeur winner of the 2011 Edward O. Wilson Biodiversity Technology Pioneer Award for co-creating LeafSnap, an electronic field quide for iPhones and iPads that identifies trees.



"The idea of building classifiers that say, 'Is this person in the photo a man or a woman?' or 'Is that leaf a sugar maple or a silver maple?' uses a lot of the same sort of math and technology." he said.

The award, which honors individuals who have significantly contributed to the preservation of biodiversity on Earth, was presented in October by Dr. Edward O. Wilson, Professor

Emeritus, Harvard University, a biologist recognized worldwide as the Father of Biodiversity, in a public forum of 2,500 individuals in Bozeman. Montana. The co-recipients are Belhumeur's collaborators, John Kress, a research botanist and curator with the Smithsonian Museum of Natural History, and David Jacobs, computer science professor at the University of Maryland.

#### **Department** News & Awards

single-track

once every two

years, has given

award to "Cells:

A Virtual Mobile

Architecture" by

Jeremy Andrus,

Christoffer Dall,

and Alex Van't

ubiguitous, and

Hof, postdoc-

Ph.D. students

a best paper

Smartphone

systems



Jeremy Andrus



Christopher Dall



Dr. Oren Laaden



carry multiple phones to accommodate work, personal, and geographic mobility needs The authors created Cells,

a virtualization

Assoc. Professor Jason Nieh

architecture for enabling multiple virtual smartphones to run simultaneously on the same physical cellphone in an isolated, secure manner. *Cells* introduces a usage model of having one foreground virtual phone and multiple background virtual phones. This model enables a new device namespace mechanism and novel device proxies that integrate with lightweight operating system virtualization to multiplex phone hardware across multiple virtual phones while providing native hardware device performance. Cells virtual phone features include fully accelerated 3D graphics, complete power management features, and full telephony functionality with separately assignable telephone numbers and caller ID support. The researchers have implemented a prototype of *Cells* that supports multiple Android virtual phones on the same phone. Their performance

SOSP, a premier results demonstrate that *Cells* imposes only modest runtime and memory overhead, works conference held seamlessly across multiple hardware devices including Google Nexus 1 and Nexus S phones, and transparently runs Android applications at native speed without any modifications.

> Computer Science Ph.D. student Jeremy Andrus and his advisor, Associate Professor Jason Nieh, co-authored a paper that won the Best Paper Award at the Association for Computing Machinery's Technical Symposium on Computer Science Education. The paper, "Teaching Operating Systems Using Android", was selected from a pool of 289 papers submitted to the conference, which provides a forum for sharing new ideas for courses, laboratories, and other elements of teaching and pedagogy. The paper describes their experience in the fall of 2010 teaching and working as a teaching assistant in the Columbia University operating systems course (COMS W4118).



Aaron Bernstein has been awarded the Best Student Paper Award at SODA 2012, the ACM-SIAM

Symposium on Discrete Algorithms and the top algorithms conference, for his single-authored work on "Near Linear Time Epsilon-Approximation for Restricted Shortest Paths in Undirected Graphs." Aaron is the sole winner of this award at SODA 2012.



Associate Professor Luca Carloni was among 85 engineering faculty nationwide selected to take part

in the National Academy of Engineering's (NAE) 17th annual U.S. Frontiers of Engineering symposium. The symposium took place at Google's headguarters in Mountain View, CA, and examined additive manufacturing, engineering sustainable buildings, neuroprosthetics, and semantic processing. The selection marks Carloni as being among the nation's brightest young engineers between 30 and 45 years old who are performing exceptional engineering research and technical work in industry, academia, and government. Carloni is working to develop smaller, faster, and more energy efficient computer chips by reinventing chip architectures and the tools used to design them.

A demo presented by members of the Energy Harvesting Active Networked Tags (EnHANTs) Project received the Best Student Demo Award in the ACM Conference on Embedded Networked Sensor Systems (ACM SenSys 2011), which is the premier conference of the sensor networking community. The demo, titled "Organic Solar Cell-equipped Energy Harvesting Active Networked Tag (EnHANT) Prototypes" was developed by 10 students (Gerald Stanje, Paul Miller, Jianxun Zhu, Alexander Smith, Olivia Winn, Robert Margolies, Maria Gorlatova, John Sarik, Marcin Szczodrak, and Baradwaj Vigraham) from the groups of Associate Professor Luca Carloni (Computer Science), Professor Peter Kinget (Electrical Engineering), Associate Professor Ioannis Kymissis (Electrical Engineering), and Assistant Professor Gil Zussman.

The EnHANTs Project is an interdisciplinary project that focuses on developing small, flexible, and energetically self-reliant devices. These devices can be attached to objects that are traditionally not networked (e.g., books, furniture, walls, doors, toys, keys, clothing, and produce), thereby providing the infrastructure for various novel tracking applications. Examples of these applications include locating misplaced items, continuous monitoring of objects (e.g., items in a store and boxes in transit), and determining locations of disaster survivors.

In 2009, the project won first place in the Vodafone Americas Foundation Wireless Innovation Competition, and in 2011 it received the IEEE Communications Society Award for Outstanding Paper on New Communication Topics. The project has been supported by the National Science Foundation, the Department of Energy, the Department of Homeland Security, Google, and Vodafone.



with the Electrical Engineering Department, was named the **Richard Dicker** Professor of Telecommunications. Professor Chang also received the 2011 ACM SIGMM Technical Achieve-

Shih-Fu Chang,

a Joint Professor

ment Award. The award cites Chang's significant contributions that shape directions in many key areas of multimedia, including multimedia search, video summarization, compresseddomain manipulation, and trustworthy media.



Game Theory, Economics and Computer Science: Equilibria, Fixed Points, and Beyond".

Recently concepts and methodologies from game theory and economics have found numerous successful applications in the study of the Internet and e-commerce. The main goal of this proposal is to bridge the algorithmic gap between these three disciplines. Chen will work to develop efficient algorithms for some of the fundamental models and solution concepts and to understand the computational difficulties inherent within them, with the aim to inspire and enable the next-generation e-commerce systems. The proposed research will contribute

to a more solid algorithmic and complexity-theoretic foundation for the interdisciplinary field of Algorithmic Game Theory.





March. Their work, which demonstrates advantages of overlaying virtual graphics on both live and cached imagery in 3D Professor **Steve Feiner** layout tasks,

is focused on the design and development of 3D user interfaces. The award-winning poster is titled "Manipulating Virtual Objects in Hand-Held Augmented Reality using Stored Snapshots."

Their work is an example of augmented reality, in which camera imagery is overlaid with live 3D graphics. The poster presents a set of interaction techniques that allow a user to first take snapshots of a scene using a tablet computer, and then jump back and forth between the snapshots, to revisit them virtually for interaction. By storing for each snapshot a still image of the scene, along with the camera position and orientation determined by computer vision software, this approach allows the overlaid 3D graphics to be dynamic and interactive. This makes it possible for the user to move and rotate virtual 3D objects from the vantage points of different locations, without the overhead of physically traveling between those locations.

Technology Student Paper Award at IEEE **ISMAR** (IEEE International Symposium on Mixed and Augmented Reality), the premier conference in its field.

Ph.D. student

Henderson

won a Best

Science and

Steve

Presented in Basel, Switzerland, "Augmented Reality in the Psychomotor Phase of a Procedural Task" reports on a key part of Steve Henderson's spring 2011 dissertation, and was coauthored by Dr. Henderson and his advisor, Professor Steve Feiner. It presents the design and evaluation of a prototype augmented reality user interface designed to assist users in performing an aircraft maintenance assembly task. The prototype tracks the user and multiple physical task objects, and provides dynamic, prescriptive, overlaid instructions on a tracked, see-through, head-worn display in response to the user's ongoing activity. A user study showed that participants were able to complete aspects of the assembly task in which they physically manipulated task objects significantly faster and with significantly greater accuracy when using augmented reality than when using 3D-graphicsbased assistance presented on a stationary LCD panel

nation's most advanced networking consortium, has awarded Ph.D. student Kyung-Hwa Kim one

of two 2011 Internet2 Driving Exemplary Applications (IDEA) student awards for innovation in advanced network applications for collaborative research and education. Kyung-Hwa is a Ph.D. student in the Internet Real-Time Lab, headed by Professor Henning Schulzrinne.

Kyung-Hwa's project, DYSWIS, is a collaborative network

Internet2, the

fault diagnosis system, with a complete framework for fault detection, user collaboration and fault diagnosis for advanced networks. With the increase in application complexity, the need for network fault diagnosis for end-users has increased. However, existing failure diagnosis techniques fail to assist end-users in accessing applications and services. The key idea of DYSWIS is a collaboration of end-users to diagnose a network fault in real-time to collect diverse information from different parts of the networks and infer the cause of failure.



Assistant Professor Simha Sethumadhavan has received an NSF CAREER Award for an ongoing research project

on hardware security.

Hardware components can contain malicious, illegal modifications that can siphon sensitive information to transmit to adversaries or shutdown critical operations. Such modifications to the hardware—the root of trust in computing-can compromise trustworthiness of systems because all software runs on hardware. This research investigates techniques to build trustworthy hardware systems even with such untrustworthy, malicious hardware components. Sethumadhavan is the founding director of the Computer Architecture and Security Technologies Lab (CASTL).



**CCLS** Director **David Waltz** was the 2011 recipient of the AAAI Distinguished Service Award. The award was

established in 1999 to honor an individual for extraordinary and sustained service to the artificial intelligence community. The AAAI Awards Committee mentioned his extraordinary and long-term technical and

organizational leadership at both the community and individual levels, and the numerous Al scientists who have been directly touched in more personal ways by his insights, wisdom, and mentorship. The award was presented at AAAI-11 in San Francisco during the opening ceremony.



Assistant Professor **Junfeng Yang** has received the Young Investigator Research Program award

from the Air Force Office of Scientific Research and its Program for Information Operations and Security. The award honors and supports scientists and engineers who show exceptional ability and promise for conducting basic research.

Yang will investigate concurrency attacks and defenses. Nowadays multithreaded programs power almost all web services such as Facebook, Google, and iCloud and many applications such as Microsoft Office and Google Chrome Browser. Yet these programs are plagued with subtle but serious concurrency vulnerabilities such as race conditions. Just as vulnerabilities in sequential programs can lead to security exploits, concurrency vulnerabilities can also be exploited by attackers to steal private information such as credit card numbers, send spams, etc. Few existing defense techniques can deal with concurrency vulnerabilities, and in fact many of the traditional defense techniques are rendered unsafe by concurrency vulnerabilities. The goal of this project is to take a holistic approach to creating novel program analysis/protection techniques and a system called DASH to secure multithreaded programs and harden traditional defense techniques in a concurrent environment.

#### **Columbia University in the City of New York**

Department of Computer Science 1214 Amsterdam Avenue Mailcode: 0401 New York, NY 10027-7003

ADDRESS SERVICE REQUESTED

NON-PROFIT ORG. U.S. POSTAGE **PAID** NEW YORK, NY PERMIT 3593

#### Extra News

### Computer History Museum Records History of Professor Joseph Traub's Career



Edwin Howard Armstrong Professor of Computer Science Joseph Traub

The Computer History Museum has recorded an extensive oral history of Edwin Howard Armstrong Professor of Computer Science **Joseph Traub**'s career. Professor Traub was the founding chair of the Computer Science Department at Columbia, founded the Journal of Complexity, and was the first chair of the Computer Science and Telecommunications Board (CSTB) of the National

Academies. His research contributions include the creation of optimal iteration theory and, together with Professor Henryk Wozniakowski, the creation of the field of information-based complexity, which studies algorithms and complexity for continuous problems. In the wide-ranging interview, conducted by Prabhakar Raghavan, Professor Traub summarizes his life and career as being "very lucky."

Further information, including a transcript of the interview, can be found at http://www.computerhistory.org/ collections/accession/102745087