University of California, San Diego

Technologies and Foundations for Robust and Secure Networked Systems

Spring 2009 Newsletter

Center for Networked Systems

Advances in Networked Computing at CNS Winter 2009 Research Review

The members of CNS held their biannual research review January 14 and 15 in UC San Diego's Computer Science and Engineering (CSE) building to outline progress-to-date on their collective projects across all aspects of networked systems. Fifty faculty, research scientists and graduate students were joined by more than 40 industry participants from CNS member companies to discuss some of the top challenges facing their respective organizations and to provide feedback on the ongoing work at the center.

The review spanned a variety of topics, from conserving system memory through a technique known as "identical page sharing," to detecting malicious Web sites by analyzing the various characteristics associated with their URLs. CNS graduate students played a significant role in the review, both as presenters and as participants in a poster session and reception that exhibited their research. "The CNS research presentations showcased both the breadth and depth of our work in networked systems," remarked Amin Vahdat, director of CNS and a professor in UC San Diego's Computer Science and Engineering (CSE) department. "Our faculty, staff, and graduate students are working toward a vision of computing and storage delivered efficiently, reliably, securely and with low power. All presentations were made by our graduate students, giving them invaluable presentation experience and showcasing their talents for our member companies." Also speaking at the review were representatives from the computing industry, including Google Technical Manager Bob Felderman; Motorola Senior Fellow and Chief Architect Hamid Ahmadi; and Cisco Distinguished Engineer Flavio Bonomi.

The event's keynote speaker, Amazon.com Chief Technology Officer and Vice President Werner Vogels, presented a talk titled "Ahead in the Cloud: The Power of Infrastructure as a Service" to a standing-room-only audience. continued on page 3

CNS Initiates New Lecture Series

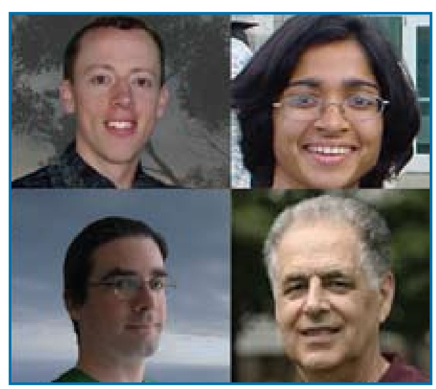

Starting in 2009, CNS now hosts guests from industry and academia to speak to faculty and students about cutting-edge topics of interest in systems and networking. Furthermore, beginning in March, most talks also are being webcast and video archived for later viewing on the CNS web site (speakers pictured clockwise from upper right: Ranjita Bhagwan, Sandy Fraser, John Howell, Adam Bechtel).

For the latest information on our upcoming lectures and other events, see http://cns.ucsd.edu/upcoming.shtml. For archives of abstracts, slides and webcasts of past lectures, go to http://cns.ucsd.edu/lecturearchive.shtml.

The CNS Lecture Series is free and open to the public. If you would like to receive emails advertising our talks or have any questions, please contact us at cns@ucsd.edu.

continued on page 8

CNS Members

AT&T, HP invent, Google, Motorola, Qualcomm, NetApp, Sun microsystems, Cisco

In This Issue

CNS Research Receives Further Support with Gifts from Google

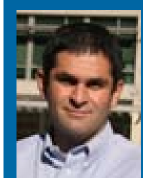

As part of its growing efforts to support excellent research in academia, CNS industry member company Google awarded two gifts to CNS researchers in early 2009.

The first gift went to Professor Amin Vahdat (pictured) in support of his CNS research project, “System Support for Cross-Data Center Replication.”

The second gift was awarded to CSE Associate Professor Stefan Savage in support of the CNS research project he is conducting with Associate Professor Geoffrey M. Voelker, “Scam Analysis and Defense Via Botnet Infiltration.”

CNS Awarded UC Discovery Opportunity Award

The purpose of the Industry-University Cooperative Research Program (IUCRP) is to leverage academic research and development investment on behalf of California-based companies with a view to accelerating technology transfer and promoting research that meets California’s needs. In recognition of the importance of the industrial and academic collaborations occurring at CNS, IUCRP awarded to CNS in 2009 a University of California Discovery Opportunity Award in support of the CNS Winter 2009 Research Review.

For more information about IUCRP and its UC Discovery Grant program, go to **http://ucdiscoverygrant.org/.**

CSE Professor Alex Snoeren Receives Prestigious Sloan Research Fellowship

The Alfred P. Sloan Foundation recognized CNS faculty member Alex C. Snoeren with a prestigious Sloan Research Fellowship in February 2009. The distinction is one of 118 given annually to early-career U.S. and Canadian “scientists and scholars of outstanding promise … in recognition of distinguished performance and a unique potential to make substantial contributions to their field.” Snoeren is a member of CSE’s systems group and received the award for his ground-breaking work on Decongestion Control, Secure and Policy-Compliant Source Routing, and Cloud Control with Distributed Rate Limiting.

Professor Snoeren joined the CSE faculty in 2003, where he is a member of the Systems and Networking research group. He received a Ph.D. in Computer Science from MIT the same year, and his earlier degrees from Georgia Tech. Snoeren is a recipient of an NSF CAREER Award (2004), the MIT EECS George M. Sprowls Doctoral Dissertation Award (Honorable Mention, 2003), and the Best Student Paper award at the ACM SIGCOMM conference (2001). As part of the UCSD Collaborative Center for Internet Epidemiology and Defenses, Snoeren helped develop Potemkin, the world’s largest high-interaction honey farm, to study the spread of worms and viruses. His research interests include operating systems, distributed computing, security as well as mobile and wide-area networking.

CNS Researchers Receive OSDI Best Paper Award for New Ways of Saving Memory

CSE graduate students Diwaker Gupta (pictured at left), Sangmin Lee and Michael Vrable, along with a number of CSE faculty (Stefan Savage, Alex Snoeren, George Varghese, Geoffrey Voelker and Amin Vahdat) received the inaugural Jay Lepreau Best Paper Award for their work on “Difference Engine: Harnessing Memory Redundancy in Virtual Machines” presented in December 2008 at the 8th Symposium on Operating Systems Design and Implementation (OSDI).

The researchers’ work solved a number of issues involving the loss of main memory in the use of virtual machine monitors (VMMs). Because VMMs can decrease the capital outlay and management overhead of hosting centers, they are a popular platform for Internet hosting and cloud-based computing services. However, main memory is not amenable to the multiplexing of hardware resources among virtual machines. To address this issue, the researchers built an extension of the Xen virtual machine monitor called Difference Engine. Difference Engine leverages sub-page level sharing (through page patching) and in-core memory compression and demonstrates substantial savings not only between virtual machines running similar applications and operating systems, but even across VMs running disparate workloads so that it outperforms other available VMMs.

Audio of the talk available here:

http://www.usenix.org/media/events/osdi08/tech/mp3s/gupta.mp3

PDF copy of paper available here:

http://www.usenix.org/events/osdi08/tech/full_papers/gupta/gupta.pdf

Recent CNS Graduates

Kirill Levchenko graduated in September 2008 with a Ph.D. in Computer Science. His advisors were Stefan Savage and Mohan Paturi. Mr. Levchenko is engaged in post-doctoral research at UC San Diego in the department of Computer Science and Engineering.

Erik Buchanan graduated in April 2009 with an M.S. in Computer Science. His advisor was Stefan Savage. Mr. Buchanan will be working as a Security Researcher in the Security Labs of Websense.

Jim Hong graduated in April 2009 with an M.S. in Computer Science. His advisor was Amin Vahdat. He is currently searching for a long-term position.

Ryan Roemer graduated in April 2009 with an M.S. in Computer Science. His advisors were Geoffrey M. Voelker and Stefan Savage. Mr. Roemer is taking a job as a Software Development Engineer in Microsoft's Azure Storage Group.

Jamie Bradley Steck graduated in March 2009 with an M.S. in Computer Science. Her advisor was Tajana Rosing. Ensign Steck, will be serving onboard the USS Bonhomme Richard (LHD-6) on her first sea tour as a division officer, rank Ensign (O-1), in the ship's engineering department.

Taurin Tan-atichat graduated in December 2008 with an M.S. in Computer Science. His advisor was Joseph Pasquale. Mr. Tan-Atichat works for CNS member company Google as a Software Engineer in the Platforms Group.

Yahoo! Inc. Awards Gift in Support of Computer Security Research

CNS faculty member and CSE Associate Professor Stefan Savage received a generous gift in December 2008 from Yahoo! Inc. The gift is intended to support Professor Savage's cutting-edge work in the area of computer security.

CNS Winter 2009 Research Review - continued from page 1

Cloud computing was a recurring theme at the review, and a number of presentations addressed the challenges inherent in this emerging style of computing. The general concept of cloud computing incorporates "software as a service," whereby a resource like Google Docs is provided over the Internet to users who need not have any expertise or control over the technology. Presenter Kevin Webb, a CSE graduate student, discussed techniques for "Seeding Cloud-Based Services" by addressing the need for distributed resource control and processing unstructured data (the associated research was conducted by Vahdat, CSE researcher Ken Yocum and CSE Associate Professor Alex Snoeren). In a presentation titled "Bluesky: System Support for Transparent Cloud Computing," CSE grad student James Anderson discussed research conducted by Associate Professor Geoffrey M. Voelker and Associate Professor Stefan Savage (both of CSE), which looks at ways to enhance networked systems built in a "transparent cloud" model. Transparent cloud computing is akin to using an out-sourced virtual data center, where organizations can decide how much computing power they need and what hardware/software systems they require.

The advent of cloud computing has placed a consequent emphasis on applications, and several of the CNS presentations focused on improving the ways applications are developed and shared among machines. A presentation by graduate student Cynthia Taylor, for example, investigated "Proximal Resource Architectures for Thinner Client Computing," which was essentially a primer on how to get applications like video to run faster and more smoothly on "thin-clients" such as iPhones and other tiny machines. Taylor advised using a cross-platform system called Virtual Network Computing (VNC) to speed up the "send-receive" update loop between clients and servers. By using the VNC with a smart proxy for added performance, the researchers were able to achieve a rate of 13 data updates per second, as opposed to 2.5 updates per second with no smart proxy. "We can improve VNC performance by having the smart proxy mediate the update rate despite high amounts of latency," referring to server-client lag time, said Taylor. "And we can do this in a way that doesn't change the existing source code. Our future work will focus on significantly increasing performance and adding functionality to the smart proxy, like face-detection technology for video applications."

The explosion of Internet-based "killer apps" and the simultaneous consolidation of datacenters have also posed a problem in terms of machine memory limits – a topic that was discussed in a presentation by CSE grad student Michael Vrable on "Extending Virtual Cluster Management and Resource Utilization." The research, which was conducted by CSE professors Savage, Snoeren, Varghese, Voelker and Vahdat, sought to minimize the hurdles involved in sharing memory among virtual machines. Noting that "memory is a key bottleneck to running virtual machines," Vrable explained the team's proposed solution: a system called "Difference Engine" that couples "identical page sharing" with methods of patching and compressing memory pages that are similar, but not identical.

Keeping the Internet safe and free of scams and malicious Web sites is also a priority for CNS, and two presentations focused on ways to identify and analyze the many ways spammers and hackers try to target users. A lecture titled "Beyond Blacklists: Learning to Detect Malicious Web Sites from Suspicious URLs," provided a number of lexical features and IP-based characteristics for detecting malicious URLs, spam phishing and other exploits. "There are various characteristics associated with these sites," said CSE grad student/presenter Justin Ma. "The question is: How do you relate these properties of the URLs to the maliciousness of the Web sites?"

To find out, Voelker and Savage drew malicious URLs from those submitted to "phish tanks" by online users, and compared them with benign URLS from

certain online directories that had been previously vetted for validity. Using a probabilistic linear model called "logistic regression" as a classifier, they reduced a set of 30,000 URL features down to 4,000 features for model analysis. They discovered that certain "red flags" indicate malicious intent, including 1) suspicious ownership of the site, 2) where the site is hosted geographically, 3) the registration date of the site, 4) what kind of connection the server is using, and 5) the presence of certain URL extensions. Ultimately, the researchers would like to create a URL reputation service that will allow users to query URLs via a database to determine their validity.

With more and more users not only downloading Web content but also creating it, security is a primary concern — but so is profit. Motorola's Hamid Ahmadi outlined his company's approach to expanding its market sphere of influence in a presentation called "Multimedia and Information Without Limit: A Motorola Applied Research and Technology Center." Ahmadi said Motorola's main focus at the moment is to use data in a more analytical way to enable new services. "The goal we have is to potentially create new business and new business opportunities," he said, noting that the core components of Motorola's current approach are broadband access, intelligent edge networking, next-generation video experiences, platforms and infrastructure. "But not all the research we do is internal — we have more and more joint projects with UCSD and other big universities." Ahmadi also pointed out the current challenges in the market, which include the gaps between those who create Internet content (YouTube, for example), those who aggregate or collect it (e.g. iTunes) and those, like Motorola, that distribute it through their devices.

"In the future, the promise will be 'any content on any device, any time you want it,'" predicted Ahmadi. "For example, right now, television programming is linear and not on-demand. Next-generation TV will be based on your schedule, and will be mobile." He even predicted that TV will become social, with Motorola-designed set-top boxes making it possible for users to create videos of themselves in their living rooms and immediately broadcast the footage to others' television sets in real time.

Cisco Systems' Flavio Bonomi made similarly mind-boggling predictions in his presentation, "Towards a More Application Aware Networking." By understanding the requirements of various applications and by allocating computing, storage and networking resources effectively across distributed systems, Bonomi expects that the Internet will evolve from its current arena of computers and hand-held devices to become the "Internet of things." It will support the interconnection of a huge number of smaller devices, including a proliferation of sensors, and also will connect "the elephant in the room," he said, "which is the car."

"Right now, the car is a black hole," in terms of computing and networking, added Bonomi. "We could do much more if cars were more connected: video conferencing, the mobile office ... the entire platform for computing and storage could be in the car, and all the cars on the road could be one great big distributed system."

"Some people say the electric car is just a router on wheels," he joked.

Stefan Savage noted that researchers at CNS and beyond are eager to develop such technologies, but are often stymied by a lack of appropriate platforms and testbeds. Vahdat suggested that companies like Motorola and Cisco set up testbeds at universities so that researchers could "live the experience" and serve as guinea pigs – and collaborators – as the new technologies are developed. Such collaborations are on the radar for CNS, which works with its member companies on all active research projects. In the past, CNS researchers worked with AT&T to deploy more efficient routing protocols running in the core of AT&T's global network. Currently, the unique perspectives of CNS member companies heavily influence the center's work on future data center networks and on spam analysis, and a number of joint publications and patents have come from ongoing collaborations between Sun Microsystems and CNS interns.

Looking to the future, Vahdat predicts that computing is at another transition point in its history. "With near-ubiquitous wireless network coverage and continued exponential increase in available computing power, network speed and storage capacity, we are approaching a time when all of our data, from applications, documents, video, music, photos, movies, to medical records, will be available to us instantly at any time," he noted. "Achieving such a vision will require a re-architecting of the global network from the core all the way to its edge, considering security, privacy, wireless and cellular protocols as first-class requirements. The infrastructure will be delivered by emerging mega data centers, consisting of upwards of 100,000 compute nodes with petabits/sec of network bandwidth and exabytes of storage. As we saw in the research review, CNS researchers – in collaboration with their industry partners – are at the forefront of this activity."

Good Code, Bad Computations

If you want to make sure your computer or server is not tricked into undertaking malicious or undesirable behavior, it's not enough to keep bad code out of the system.

Two CNS graduate students from the CSE department, Erik Buchanan and Ryan Roemer (pictured at right), published work at ACM's Conference on Communications and Computer Security (CCS) 2008 showing that the process of building bad programs from good code using "return-oriented programming" can be automated and that this vulnerability applies to RISC computer architectures and not just the x86 architecture (which includes the vast majority of personal computers).

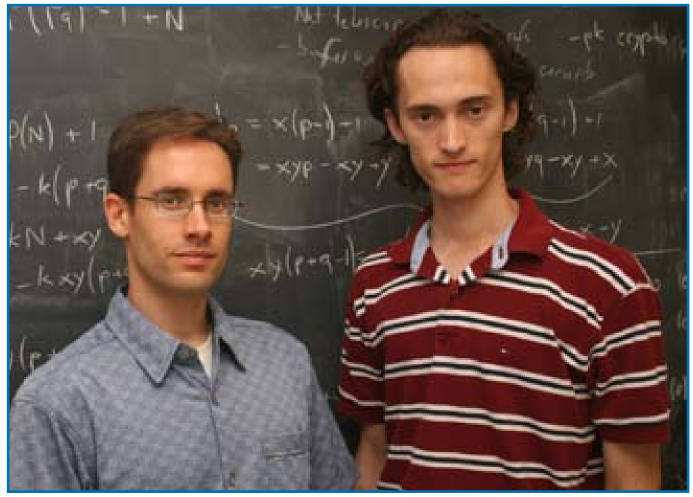

Last year, CNS faculty member and CSE Professor Hovav Shacham formally described how return-oriented programming could be used to force computers with the x86 architecture to behave maliciously without introducing any bad code into the system. However, the attack required painstaking construction by hand and appeared to rely on a unique quirk of the x86 design.

"Most computer security defenses are based on the notion that preventing the introduction of malicious code is sufficient to protect a computer. This assumption is at the core of trusted computing, anti-virus software, and various defenses like Intel and AMD's no-execute protections. There is a subtle fallacy in the logic, however: simply keeping out bad code is not sufficient to keep out bad computation," said CSE Professor Stefan Savage, another author on the CCS 2008 paper.

Return-oriented programming exploits start out like more familiar attacks on computers. The attacker takes advantage of a programming error in the target system to overwrite the runtime stack and divert program execution away from the path intended by the system's designers. But instead of injecting outside code—the approach used in traditional malicious exploits—return-oriented programming enables attackers to create any kind of nasty computation or program by using just the existing code. For example, a user's Web browser could be subverted to record passwords typed by the user or to send spam e-mail to all address book contacts, using only the code that makes up the browser itself.

"You can create any kind of malicious program you can imagine—Turing complete functionality," said Shacham.

"There is value in showing just how big of a potential problem return-oriented programming may turn out to be," added grad student Erik Buchanan.

The term "return-oriented programming" describes how the "good" instructions that can be strung together in order to build malicious programs need to end with a return command. The graduate students showed that the process of building these malicious programs from good code can be largely automated by grouping sets of instructions into "gadgets" and then abstracting much of the tedious work behind a programming language and compiler.

Imagine taking a 700-page book, picking words and phrases at random, and then assembling a 50-page story that has nothing to do with the original book. Return-oriented programming allows you to do something similar. Here the 700-page book is the code that makes up the system being attacked—for example, the standard C-language library libc—and the story is the malicious program the attacker wishes to have executed.

"The threat posed by return-oriented programming, across all architectures and systems, has negative implications for an entire class of security mechanisms: those that seek to prevent malicious computation by preventing the execution of malicious code," the authors wrote in their CCS 2008 paper.

For instance, Intel and AMD have implemented security functionality into their chips (NX/XD) that prevents code from being executed from certain memory regions. Operating systems in turn use these features to prevent input data from being executed as code (e.g., Microsoft's Data Execution Prevention feature introduced in Windows XP SP2). The new research from UC San Diego, however, highlights an entire class of exploits that would not be stopped by these security measures since no malicious code is actually executed. Instead, the stack is "hijacked" and forced to run good code in bad ways.

"We have demonstrated that return-oriented exploits are practical to write, as the complexity of gadget combination is abstracted behind a programming language and compiler. Finally, we argue that this approach provides a simple bypass for the vast majority of exploitation mitigations in use today," the computer scientists wrote in their CCS paper.

The authors outlined a series of approaches to combat return-oriented programming. Eliminating vulnerabilities permitting control flow manipulation remains a high priority—as it has for 20 years. Other possibilities: hardware and software support for further constraining control flow and addressing the power of the return-oriented approach itself.

"Finally, if the approaches fail," explained the authors, "we may be forced to abandon the convenient model that code is statically either good or bad, and instead focus on dynamically distinguishing whether a particular execution stream exhibits good or bad behavior."

“Six Degrees of Kevin Bacon” Game Provides Clue to Efficiency of Complex Networks

As the global population continues to grow, our social connections to one another remain relatively small, as if we’re all protagonists in the Kevin Bacon game inspired by “Six Degrees of Separation,” a Broadway play and Hollywood feature that were popular in the 1990s. In fact, classic studies show that if we were to route a letter to an unknown person using only friends or acquaintances who we thought might know the intended recipient, it would take five or six intermediary acquaintances before the letter reaches its intended destination.

The underlying success of this phenomenon, called the “small-world paradigm,” and discovered in the 1960s by sociologist Stanley Milgram, recently provided a source of inspiration for CNS researchers studying the Internet as a global complex network.

Photo by Sam Hodgson/voiceofsandiego.org

The result, a study undertaken by Kimberly Claffy and Dmitri Krioukov, (pictured at right), and Marián Boguñá, was published in Nature Physics on November 16, 2008. Their work reveals a previously unknown mathematical model called “hidden metric space.” This new model may explain the “small-world phenomenon” and its relationship to both man-made and natural networks such as human language, as well as to gene regulation or neural networks.

For these researchers, the concept of an underlying “hidden space” may also be relevant to their professional interests: how to remove mounting bottlenecks within the Internet that threaten the smooth passage of digital information around the globe.

“Internet experts are worried that the existing Internet routing architecture may not sustain even another decade,” said Krioukov, the study’s principal investigator with the Cooperative Association for Internet Data Analysis (CAIDA), based in the CNS-affiliated San Diego Supercomputer Center at UC San Diego. “Routing in the existing Internet has already reached its scalability limits. Black holes are appearing everywhere.”

“Discovery of such a metric space hidden beneath the Internet could point toward architectural innovations that would remove this bottleneck,” added Claffy, director of CAIDA and adjunct professor of computer science in UCSD’s Jacobs School of Engineering. “Although quite prevalent in the natural world, the idea of routing using only local rather than global knowledge of network connectivity represents a revolutionary change in how to think about engineering communications networks. And yet the current Internet architecture is fundamentally limited by the overwhelming amount of routing knowledge that must be continuously transmitted through what is now critical global infrastructure.”

According to the researchers, natural networks appear to transmit signals or messages with a high degree of efficiency, even though no single node – whether it’s an individual person in a social network or a single neuron in a neural network – can visualize the global structure of the entire network.

How is this possible? By constructing a mathematical model of geometry underlying the topology of these networks, the researchers discovered that many complex networks shared a similar characteristic – their global topological structure (or shape) maximizes their communication efficiency.

Take, for example, the “small-world phenomenon” described earlier. In this case, the only information people possessed to make their routing decisions was a set of descriptive attributes of the destined recipient – his or her home base and occupations. People then determined who among their contacts was “socially closest” to the target. For aficionados of the Kevin Bacon game, the goal was to connect any actor in Hollywood to Bacon through the films he made.

“The success of Milgram’s experiment indicates that social distances among individuals – although they may be difficult to define mathematically – have a role in shaping the network, and may also be essential for efficient navigation,” said Claffy. Added Krioukov: “When you know the network topology, you merely know the basic layout of a network. But when you discover its underlying geometry, or hidden space, you may know how this complex network really functions.”

So, what accounts for this inherent communication efficiency of complex networks? The study suggests the existence of an underlying geometric framework that contains all the nodes of the network, shapes its topology and guides routing decisions: a “hidden metric space.” Distances in this space are akin to social distances in the “small-world phenomenon.” They measure similarity between people. The more similar the two persons, the closer they are in the “social space,” and the more likely they are friends, connected in the acquaintance network. To route a message, a person forwards it to the friend socially closest to the message destination, as illustrated .

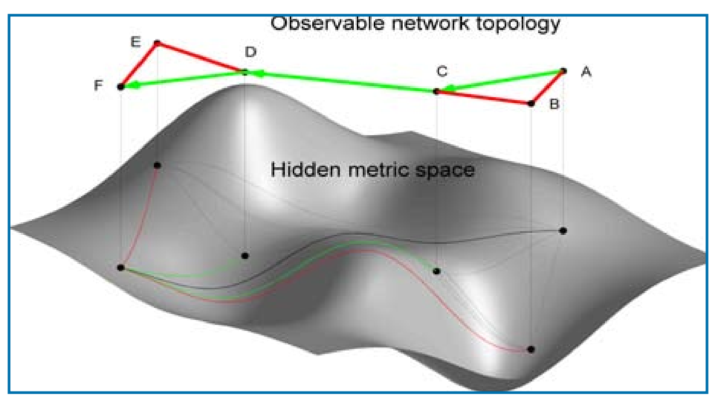

"Such routing allows networks to efficiently find intended communication targets even though they do not have a global view of the system," said Claffy.

The primary motivation for this study, according to Krioukov, was the constantly increasing size and dynamics of the Internet, leading to increasing incidences of routing bottlenecks. Discovery of the Internet's "hidden metric space" would allow messages to be forwarded to destinations based on local measurements of similarities between nodes, their positions in the "hidden space," rather than on their positions in the network, which requires global measurements of its structure.

Krioukov also suggests that reconstruction of hidden metric spaces underlying a variety of real complex networks may have other practical applications. For example, hidden spaces in social or communications networks could yield new, efficient strategies for searching for specific individuals or content. The metric spaces hidden under some biological networks also could lead to powerful tools for studying the structure of information or signal flows in these networks.

"This could be applied to cancer research, for example, in which studies rely heavily on gene regulation," said Krioukov. "Suppose you were able to find the hidden space here. One could then figure out what drives gene regulation networks and what drives them to failure. This would be an important contribution to the field."

Center for Integrated Access Networks

The National Science Foundation has awarded a five-year, $18.5 million grant to establish the Center for Integrated Access Networks (CIAN). The engineering research center, or ERC, is a 10-university partnership, directed by the University of Arizona (UA).

The new center's deputy director is affiliated with CNS: Jacobs School of Engineering electrical and computer engineering (ECE) professor Shaya Fainman, the school's Cymer Professor in Advanced Optical Technologies. Fainman is involved in the design and realization of ultrafast and miniature optical systems.

"Our goal over the next decade is to devise and adapt chip-scale optoelectronic integration technologies for an advanced optical access network capable of delivering data at 10 gigabits, or 10 billion operations per second, to single users anywhere, at anytime and at lower cost," said Fainman. "By contrast, the current data transfer rate is about 10 megabits, or 10 million operations per second."

In addition to Fainman, other UC San Diego faculty involved in CIAN include CNS director Amin Vahdat as well as ECE professors George Papen, Stojan Radic and Joseph Ford.

CIAN will focus on removing one of the last bottlenecks in the Internet by developing optoelectronic technologies for high-bandwidth, low-cost, widespread access networks. The universities will collaborate to create an advanced optical access network capable of delivering data more than a thousand times faster to users at lower cost than they now pay to connect to information data bases and communication networks.

"As the world increasingly relies on communications networks, we anticipate that CIAN will contribute the understanding and innovations needed to extend the reach and expand the capabilities of these networks," said Lynn Preston, NSF deputy division director and leaders of the ERC program. "We expect this area of research to interest many pre-college students in the program and in engineering, and we look forward to CIAN graduates becoming leaders and innovators in the creation of future communications systems."

Partner universities of UA and UC San Diego in CIAN include Caltech, Stanford, USC, UCLA, UC Berkeley, Columbia, Norfolk State and Tuskegee.

The new center offers corporate sponsors a voice in various functions of CIAN, including guiding the selection of research projects (not unlike the model used by the Center for Networked Systems), enabling technology transfer and participating in student education. Early industry members include Deutsche Telekom, Sun Microsystems, the U.S. unit of Japan's Nitto Denko Technical, and IBM. A total of 26 companies submitted letters of support for CIAN's proposal submission to the NSF.

In a unique approach, CIAN vertically integrates research from developing nanostructured photonic devices to demonstrating advanced network services.

The center will educate students from diverse backgrounds by piloting novel, multi-level "super courses" and student recruitment and retention programs. These will include partnerships with minority-serving institutions such as Pima Community College, Norfolk State University and Tuskegee University, as well as Native American tribes in Arizona. Outreach programs will include middle schools, high schools, undergraduate and graduate programs.

For more information, visit the CIAN Web site at www.cian-erc.org.

Winter and Spring 2009 CNS Lectures - continued from page 1

February 11, 2009
"Securing the Network"
Sandy Fraser, former head of AT&T Research and now head of Fraser Research, discussed his current work on new Internet design. Specifically, he proposed an alternate to IP – what he calls "the global ethernet."

February 20, 2009
"Leveraging Legacy Code to Deploy Desktop Applications on the Web"
John Howell, a Researcher at Microsoft, works on the intersection of security and scalability in distributed systems. He spoke on Xax, a browser plugin model that enables developers to leverage existing tools, libraries and entire programs to deliver feature-rich applications on the Web.

February 23, 2009
"Scaling Internet Datacenters"
Adam Bechtel, Director of the Global Network Architecture group at Yahoo!, spoke about how the company's data center networks faced scaling challenges as their infrastructure evolved from a single server to massive datacenters measured in the tens of megawatts.

March 10, 2009
"NetPrints: Diagnosing Home Network Misconfigurations Using Shared Knowledge"
Ranjita Bhagwan, a Researcher from Microsoft Research in Bangalore, spoke on NetPrints, a system that leverages shared knowledge in a population of users to diagnose and resolve misconfigurations.

March 20, 2009
"Failure Immunity: Teaching Systems to Defend Against Failures"
George Candea, a Researcher from the Swiss Federal Institute of Technology in Lausanne (EPFL), spoke on techniques his group is developing that enable programs to automatically gain failure immunity without assistance from programmers or users. In particular he focused on Dimmunix, a tool that gives deadlock immunity to both Java and C/C++ programs.

Upcoming Events

April 17, 2009
CNS Lectures: Balaji Prabhakar
Location: Calit2 Auditorium, Atkinson Hall, UC San Diego
Time: 2:00pm-3:00pm

Abstract: TBA
Bio: Balaji Prabhakar is an Associate Professor of Electrical Engineering and Computer Science at Stanford University. Balaji is interested in network algorithms, in scaleable methods for network performance monitoring and simulation, in wireless (imaging) sensor networks, stochastic network theory and information theory. He has designed algorithms for switching, routing, bandwidth partitioning, load balancing, and web caching. Balaji has been a Terman Fellow at Stanford University and a Fellow of the Alfred P. Sloan Foundation. He has received the NSF CAREER award and the Rollo Davidson Prize awarded to young scientists for their contributions to probability and its applications.

July 15-16, 2009 [All Day]
CNS Summer 2009 Research Review
Location: Room 1202, CSE Building, UC San Diego
Keynote Speaker: Greg Papadopoulos, Chief Technology Officer and Executive Vice President of Research and Development, Sun Microsystems.

The CNS Summer 2009 Research Review will feature talks by our industry members about their latest research challenges and innovative successes. Faculty and graduate students will present summaries of recently completed and ongoing CNS projects and will propose a new slate of projects to the CNS Advisory Board members. Graduate students will also present their current research at a poster session. The event provides numerous opportunities for informal interactions and collaborations among graduate students, faculty and industry research executives.

Attendance at the Summer 2009 Research Review is limited to industry sponsors and invited guests. If you are interested in attending, or have questions, please contact Kathy Krane at kkrane@ucsd.edu. More information will also be available at our Research Review Web page, **http://cns.ucsd.edu/2009Summerreview.shtml.**

Mission and Objectives of CNS

The mission of CNS is to develop key technologies and frameworks for networked systems. By combining our research talents and strengths in partnership with industrial leaders, CNS achieves critical mass and relevant focus, accelerating research progress and creating key technologies, frameworks and systems understanding for robust, secure networked systems and innovative new applications. CNS also works to educate the next generation of top students with a perspective on industry-relevant research and to train students on how to continue their leadership throughout their careers. This is accomplished by bringing together leading faculty, students, and companies to investigate the most challenging, interesting and important problems in computer networks.

If you are interested in joining the Center, please contact Director Amin Vahdat at vahdat@cs.ucsd.edu.