
cns
Center for Networked Systems

Director's Message

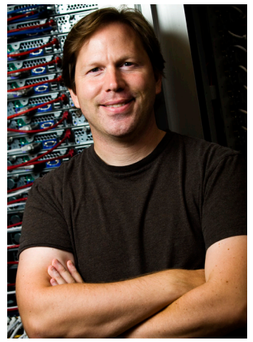

In recent years, networked systems and the countless technological applications they have made possible have achieved a prevalence in contemporary life that nears ubiquity. Networked systems and applications open up new avenues and means for global commerce: they make available vast libraries of interactive databases and archives; they manage civic, military, and corporate infrastructures; they record, organize, and analyze medical, biological, financial, and environmental research; and they make possible multiple platforms for communication and social networking. And as each new application of networked technology is developed, problems associated with issues of security, privacy, reliability, robustness, energy efficiency, and thermal management become more dramatic and pressing.

The Center for Networked Systems (CNS), established in 2004, now forms a hub for some of the leading research in networked systems and security. The reputation of the systems and networking group at UCSD has grown rapidly, so that it is now considered among the top places in the world for the exploration of some of the most important topics in our field. CNS brings researchers in the UCSD community together with counterparts in industry to develop key technologies and frameworks for networked systems as well as to foster a lively intellectual community that works to accelerate the advance of innovation in the field. Our areas of expertise and activity involve modeling network traffic and topology, network security (including botnets, spam, and the underground economy), data center architecture, energy efficiency, thermal and energy management in compute infrastructure and wireless devices, "cloud" computing, data center networking, wireless system performance and diagnosis, and database organization for unstructured ("XML") data.

The high regard for our researchers is not only a reflection of their demonstrable excellence, but also of their willingness to interact with and respond to the issues facing those sectors of the technology industry that seek to innovate in these research spaces. Our leadership position would not have happened if we were not informed by the "real world" problems our partners struggle to solve. Our ability to keep all participants involved in a general conversation has been one of the keys to the high level of intellectual exchange and productivity in our relationships and events. Additionally, we pride ourselves on being highly reactive to our members' individual needs and points of view, in large part because we have proven that input from industry provides invigorating perspectives on the opportunities for innovation in systems and networking.

We believe that we have created an institution that is in the vanguard of the communications networks revolution. With the combined strengths of our world-class students and faculty and their proven effectiveness in collaborations with our industry partners, UC San Diego's Center for Networked Systems will continue to be a leader in the field.

Stefan Savage
Interim Director, CNS

Is Your Cloud Leaking? Security Researchers Investigate

The recent technological developments that enable companies and institutions to outsource their computing needs to third-party providers is an economic boon in terms of savings from investments in infrastructure, personnel, space and utilities. However, every advance in networked technologies presents the potential for new security vulnerabilities, and cloud computing is no exception.

Computer Science and Engineering (CSE) professors Hovav Shacham and Stefan Savage, along with then graduate student (now Assistant Professor at University of Wisconsin, Madison) Thomas Ristenpart, set out to explore the possible vulnerabilities of third-party cloud infrastructures as they are being developed in the marketplace. While most critics of the cloud have talked about the dangers inherent to the trust relationship between the customer and the cloud services provider, the CNS investigators envisioned a less obvious threat. "Cloud infrastructures can also introduce non-obvious threats from other customers due to the subtleties of how physical resources can be transparently shared between virtual machines," explains Shacham. Because cloud providers maximize efficiency through the use of virtual machines that multiplex distinct customers on common physical hardware, there is a possibility that adversarial customers might share space on the same physical servers. Since any two residents sharing server resources are open to attack from their neighbors, this characteristic of cloud computing infrastructure creates an incentive for malicious actors to locate adversaries within the cloud, arrange to place themselves as customers within the same physical servers, and then to launch an attack from this privileged position.

While cloud providers could attempt to obfuscate the location of targets in their network, the investigators demonstrated how any possible network configuration remains open to the exploration of a sufficiently resourceful and determined cloud cartographer. Therefore, the researchers propose that providers of cloud services should make the risks transparent to users who might then, based upon their need for privacy, elect to bear the opportunity costs of using physical machines populated with only their virtual machines.

CNS-Industry Success Stories

CISCO SYSTEMS: KANGAROO

This a project between Professor George Varghese, Ph.D. student Christos Kozanitis and Cisco Systems researchers Sushil Singh and John Huber. The team worked to develop Kangaroo as a programmable parser that can adapt as new protocols are added and yet retains speed by "leaping" over multiple packet headers in a single step. Kangaroo mitigates the cost of leaps by providing an offline algorithm to find where the leaps should occur in the parse tree. The researchers co-authored a paper on Kangaroo that was published in INFOCOM 2010.

> "The collaboration on Kangaroo has allowed Cisco to explore some exciting new flexible parsing methodologies that could potentially allow future ASICs to maintain line rates even when new protocols need to be supported."
>
> *- John Huber, Technical Leader, Cisco Systems*

ERICSSON: Mobile Malware

Professor Geoffrey M. Voelker, Ph.D student Gjergji Zyba and Calit2 Principal Development Engineer Per Johansson have worked since 2007 on several projects involving mobile malware with Ericsson researchers Michael Liljenstam and András Méhes. This partnership has resulted in co-authored papers in such respected venues as ACM WORM 2007 and IEEE INFOCOM 2009.

> "In order to maintain the high security level in mobile broadband, Ericsson is proactive with respect to potential security threats from malware targeting smart phones and advanced mobile devices. The current attention directed to mobile malware accentuates the value of the collaboration with UCSD over the last few years, which has generated a better understanding of the potential threat scenarios and solution strategies to protect users and networks. Ericsson benefits further from being able to draw upon the great expertise at UCSD regarding Internet threats, and the collaboration contributes to Ericsson's position as a leading communications provider with security as an important focus area."
>
> *- Michael Liljenstam, Ericsson Researcher*

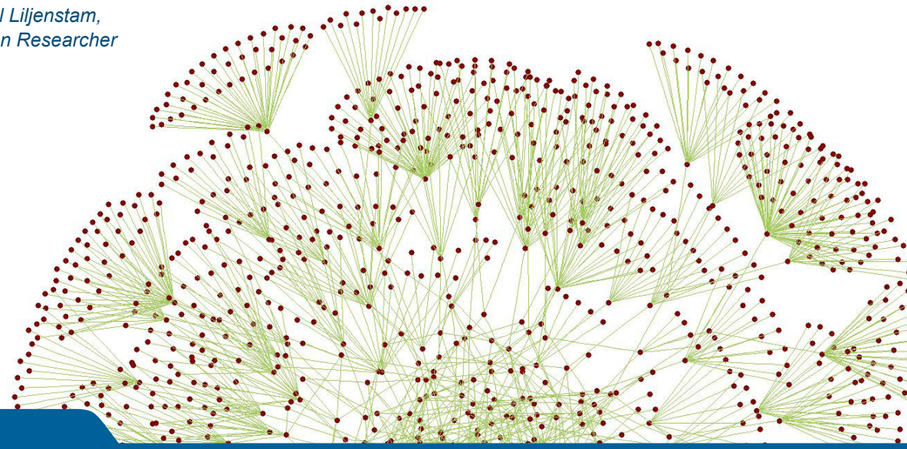

ORACLE: Thermal Management

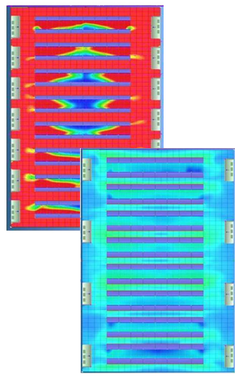

Dr. Tajana Rosing has collaborated steadily with Kenny Gross of Oracle on a range of thermal and cooling management issues for both multi-core systems and data centers. In one project, Dr. Rosing and Dr. Gross, working with Dr. Ayse Coskun (then a doctoral student, now an Assistant Professor at Boston University), developed proactive, dynamic thermal management techniques for multiprocessor system-on-chips that achieved thermal profiles superior to those of systems managed reactively while avoiding the majority of the performance costs normally associated with dynamic thermal management systems.

> "Our collaborative investigations with Prof. Rosing have led to new innovations and new intellectual property for Oracle and for UCSD that will ultimately benefit the broader computing industry by simultaneously increasing both energy efficiency and workload throughput performance for future data center computing assets."
>
> *- Kenny Gross, Distinguished Engineer, Oracle*

Multiple Threads and No Idle Processes

The core of CNS work is performed through two-year grants funded by our industry members. These grants are awarded annually to an average of six research teams. The research has, in turn, had a concrete impact on the development of products and operations at our member companies, in addition to generating doctoral dissertations, award-winning papers at top conferences, and federal funding awards.

In addition to the center grants, CNS researchers collaborate with industry members on special projects to analyze, solve, and test solutions to the most pressing current problems in systems and networking.

Partnering with Industry

CNS researchers work closely with their industry partners through direct collaboration, a regular lecture series, informal interactions and sponsored projects. The benefits of membership are considerable: members are exposed to the latest research months before new data is published, they make connections with key researchers at UC San Diego, and they are able to recruit directly from one of the most talented pools of students in the world. Each member company has a seat on the Center's Advisory Board, which gives them a voice in the research direction of the Center. CNS additionally extends an open invitation to its member companies to host industry researchers as Visiting Scholars or as speakers in the CNS Lecture Series.

Step I: *Conception:*
CNS researcher, in collaboration with an industry member or independently, formulates an idea for a CNS research grant proposal.

Step II: *Pre-proposal period:*
Advisory Board members vet proposals, provide feedback, and propose industry collaborations prior to the official proposal process.

Step III: *Proposal period:*
Proposals are presented at the CNS Summer Research Review; Advisory Board members vote on which will earn CNS support.

Step IV: *Award process:*
CNS Research Grant awardees are announced.

Step V: *Research period:*
Summaries of ongoing research progress and achievements are presented at subsequent Winter and Summer Research Reviews throughout the two year period of the grant.

Step VI: *Final presentation:*
At the end of the two-year period, completed research grants are presented in final reports at the Summer Research Review.

CNS Faculty

CNS draws its distinguished faculty from the Jacobs School of Engineering's Electrical and Computer Engineering (ECE) as well as Computer Science and Engineering (CSE) departments; the San Diego Supercomputer Center (SDSC); and the UCSD division of the California Institute for Telecommunications and Information Technology (Calit2).

Anthony Acampora, Professor, ECE
Research: Bridging gaps in America's digital infrastructure
Honors: IEEE Fellow; founding Director, Center for Wireless Communications (1995–'99)

Tara Javidi, Associate Professor, ECE
Research: Communication networks, stochastic resource allocation, and wireless communication
Honors: NSF CAREER Award (2004)

Joseph Pasquale, Professor, CSE
Research: Network/operating system software architectures to support large-scale, peer-to-peer computing on the Internet; extended client/server structures for mobile computing
Honors: J. Robert Beyster Chair in Engineering; NSF Presidential Young Investigator Award (1989)

Alex C. Snoeren, Associate Professor, CSE
Research: Operating systems, distributed computing, and mobile and wide-area networking; Internet manageability, performance and reliability
Honors: NSF CAREER Award (2004), Sloan Fellow (2009)

KC Claffy, Research Scientist, SDSC
Research: Internet workload/performance data collection, analysis and visualization; commercial ISP collaboration/cooperation and sharing of analysis resources
Honors: PI, Cooperative Association for Internet Data Analysis

Bill Lin, Professor, ECE
Research: Design of modern VLSI systems; balancing programmability, high performance and energy efficiency
Honors: Best Paper, IEEE Transactions on VLSI Systems; 2 patents

Ramesh Rao, Professor, ECE; and Director, UCSD Division, Calit2
Research: Architectures, protocols and performance analysis of wireless, wireline and photonic networks for integrated multimedia services
Honors: Qualcomm Endowed Chair in Telecommunications and Information Technologies

Amin Vahdat, Professor, CSE; and Director, CNS
Research: High-performance, robust computer systems, including data center architecture, virtual machines, operating systems, network protocols, and overlays
Honors: Science Applications International Corporation Professor in Engineering; NSF CAREER Award (2000); Sloan Fellowship (2003)

Rene Cruz, Professor, ECE
Research: Design and performance analysis of communication networks; high-speed wireless and optical systems, scheduling, routing, and network calculus
Honors: IEEE Fellow; co-founder, Mushroom Networks, Inc.

Keith Marzullo, Professor, CSE
Research: Fault-tolerant distributed computing in grid computing, mobile computing and clusters
Honors: Former Chair, CSE (2005-'10)

Stefan Savage, Professor, CSE
Research: Operating-system kernel design, disk array design, network measurement tools and overlay analysis
Honors: Director, NSF Collaborative Center for Internet Epidemiology and Defenses; Interim Director, CNS; Sloan Fellow; co-founder, Asta Networks and NetSift; ACM Fellow

George Varghese, Professor, CSE
Research: Network algorithms, using algorithmic and systems techniques to speed up key networking tasks; scheduling, timers and IP lookups
Honors: 14 patents; co-founder, NetSift, Inc., acquired by Cisco Systems (2005).

Alin Deutsch, Associate Professor, CSE
Research: XML and efficient/effective integration of heterogeneous data from multiple sources
Honors: U.S. Patent based on query optimization (2000)

Yannis Papakonstantinou, Professor, CSE
Research: Query processing, especially data integration and semistructured data (XML)
Honors: NSF CAREER Award (1998); founder, Enosys Software, Inc.

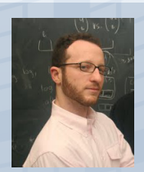

Hovav Shacham, Assistant Professor, CSE
Research: Applied cryptography, systems security, and technology policy; pairings to construct cryptographic systems
Honors: Participant in California Secretary of State's "Top-to-Bottom" review of voting machines certified for use in California (2007)

Geoffrey M. Voelker, Associate Professor, CSE
Research: Operating and distributed systems, networking, security and mobile computing
Honors: Ericsson Distinguished Scholar (2007); co-director, Reliable-Adaptive Multi-Path network project

Shaya Fainman, Professor, ECE
Research: Optical networks, systems and devices, especially ultra-fast information processing
Honors: Cymer Professor in Advanced Optical Technologies; IEEE Fellow; Fellow of the Optical Society of America; Deputy Director, NSF Center for Integrated Access Networks

George Papen, Professor, CSE
Research: Advanced photonic systems, optical communication systems, optical networking and atmospheric remote sensing.
Honors: 3 patents; vice-chair, CSE

Tajana Simunic Rosing, Assistant Professor, CSE
Research: Energy-efficient computing, particularly in embedded and wireless systems; optimization of power and topology design for systems-on-chip
Honors: Lead, Large Scale Systems, Multi-Scale Systems Center (MuSyC)

Yuanyuan (YY) Zhou, Professor, CSE
Research: Operating systems, computer architecture, storage systems and software reliability
Honors: Qualcomm Endowed Chair in Mobile Computing; NSF CAREER Award (2004); Sloan Fellowship (2007); co-founder, Pattern Insight

Botnet Judo:
using spam to fight itself

In 2004, bot-based spam distribution emerged as a new and extremely effective method for generating spam. A botnet is a network of software agents that run automatically on computer systems, usually without the knowledge or permission of the system's owner and typically with malicious intent. One of the most frequent uses for botnets is as a platform for launching large-scale spam campaigns.

To meet the challenge posed by bots, CNS graduate students Andreas Pitsillidis and Chris Kanich, working with post-doctoral researcher Kirill Levchenko, Professors Geoffrey M. Voelker and Stefan Savage, and with researchers from UC Berkeley, devised Judo, a system that can derive a near-perfect spam filter using a novel approach that turns the strength of botnets' spam-producing tactics against itself.

Currently, most bot-based spam campaigns are able to distribute massive amounts of spam with the goal being to deliver the same basic message over a vast network. However, the spammers must balance their need to maintain their basic "message" against their need to vary the overall wording of the message so that it avoids being caught by content filters. To do so, spammers developed template-based systems. These systems work like form letters in that they produce a basic text into which are inserted throughout substitution macros. This often makes the resulting e-mails unique enough to evade current content filtering programs.

Through an analysis of the nature of template systems, CNS researchers discovered that they can apply a template inference algorithm to a set of messages generated by the same bot and thereby reverse-engineer that spam generation template. Once the template can be identified, so can any spam generated by it.

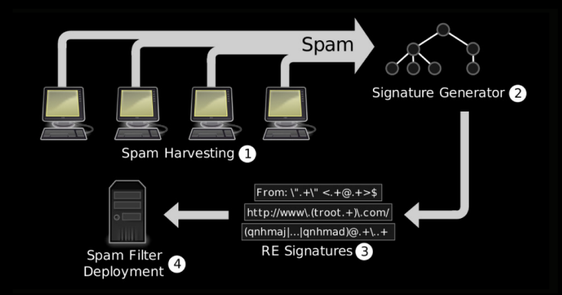

Automatic template inference makes it possible to deploy template signatures as soon as they appear "in the wild:" bots (1) running in a contained environment generate spam processed by the Judo system (2); signatures (3) are generated in real time and disseminated to mail filtering appliances (4).

"The principal requirements of a spam filtering system are that it should be both safe and effective," explains Levchenko, "meaning that it does not classify legitimate mail as spam, and it correctly recognizes the targeted class of spam." Another concern about the system is the time that elapses between the spam production of a new template and Judo's generation of a new signature for that template. In tests where Judo's performance was measured on actual spam generated by four bots, Judo was able to generate effective template signatures that filtered out spam and that produced almost no false positives. Additionally, execution time in almost all tests of the system remained under 10 seconds.

Analyzing the Gap Between Bugs and Patches

The discovery of security vulnerabilities in software and the subsequent distribution of security patches for those "bugs" is a commonplace of contemporary technology lifecycles. And while the way users and administrators respond to the discovery of software vulnerabilities has been documented, the rate and efficacy of such security fixes for cryptographic compromises is not so well-studied. CNS Professors Hovav Shacham and Stefan Savage took advantage of the discovery and repair of just such a severe vulnerability in the Debian Linux version of OpenSSL in order to observe and analyze the pattern of recovery from a cryptographic breach.

The goal of this work was to measure recovery from this type of vulnerability and to compare it to what is known about recovery from other vulnerabilities. The group found an extremely slow rate of fixing that -- unlike with conventional vulnerability patching (where fixing phases tend to be short) -- displayed a much flatter curve, with fixing extending six months after the announcement of the discovery of the bug. Their findings focused on the continued vulnerabilities in systems after the application of the fix, and on the unusually extended fixing phase of this particular bug. Though the results were troublesome, they allowed the researchers to identify predictive factors for the rate of upgrading with this sort of vulnerability.

NEON:
A System for Managing the Unmanageable

"Modern organizations face increasingly complex information management requirements," explains Qing Zhang, the lead graduate student for the team developing Neon, a system designed to manage security mandates for data. Despite numerous laws and policies that have been put in place to stop them, large-scale security breaches of information such as social security numbers or credit card numbers routinely make headlines, serving both as a source of embarrassment to the institution that mismanaged the private data and as a disruption to the lives and finances of those whose data has been compromised.

This is in part because control mechanisms put in place to implement policy considerations are applied most often to objects such as files that serve as containers for information and not to the information itself. Information is easily transformed into derived data -- that is, information that has been manipulated in some way through such processes as compressing a file, "cutting and pasting," attaching data to an email, or transmitting data over a network. Because it is both easy and convenient to manipulate data, while implementing management policies is difficult and prone to human error, information can quickly become exposed. According to Ms. Zhang, "Mandating restrictions on information use is futile unless one has a mechanism for enforcing their use both on individual hosts and between them."

The team of CNS researchers, consisting of graduate students Qing Zhang, John McCullough, Justin Ma, and Michael Vrable along with professors Amin Vahdat, Alex C. Snoeren, Geoffrey M. Voelker, and Stefan Savage, decided to meet this challenge by creating a new way to enforce information management laws and policies. Their solution takes a pragmatic approach that works with and not against the main threat to security: the way that people want to (and usually do) share information across networks.

The researchers developed Neon, a system that automatically and transparently tracks and enforces data management policies and so neutralizes the element of human error. Neon does this by integrating a mechanism within a virtual machine monitor (VMM) that "implements byte-level policy labels, called 'tints,' that are transparently propagated and combined as part of normal instruction execution," says Ms. Zhang. "Neon propagates tints across the network to and from storage, thus maintaining a binding between a policy and any derived data." Though still in its development stage, Neon has so far successfully maintained data "tints" even when tinted data has been compressed, compiled, port-forwarded, as well as copied and pasted.

Hedera: Dynamically Reducing Cost While Increasing Efficiency

In a technological development that could not have been predicted even a few years ago, more and more organizations are either building large-scale data centers or are using data center resources accessed through cloud-computing host providers. These data centers often aggregate bandwidth to thousands or tens of thousands of machines to create complicated services such as those found in social networking or e-commerce web sites or in powerful distributed computing frameworks like Hadoop, MapReduce, or Dryad. However, while the scale and pattern of utilization for data centers has changed, data center routing and forwarding protocols have yet to adapt to the growing demands of this new paradigm.

In a bid to meet the challenge of this emerging issue, Mohammad Al-Fares, a CNS graduate student researcher in a team with other CNS students and faculty, introduced Hedera, calling it "a scalable, dynamic flow scheduling system that adaptively schedules a multi-stage switching fabric to utilize aggregate network resources efficiently." Hedera is designed to address the main problems facing network designers of modern data centers.

A full implementation of Hedera has been completed on the Portland testbed. Using commodity switches and unmodified hosts, Al-Fares and his colleagues demonstrated that Hedera produces bisection bandwidth for a simulated host data center that is 96% optimal and 113% better than current static load-balancing methods. While Hedera's gains in performance are dependent upon network flow and are more apparent when the network is stressed by several large data transfers, the improvements significantly outperform current routing and forwarding protocol solutions. Moreover, notes Al-Fares, given "the large investment in network infrastructure associated with data centers (many millions of dollars), and the incremental cost of Hedera's deployment (for example, one or two servers), we show that dynamic flow scheduling has the potential to deliver substantial bandwidth gains with moderate additional cost."

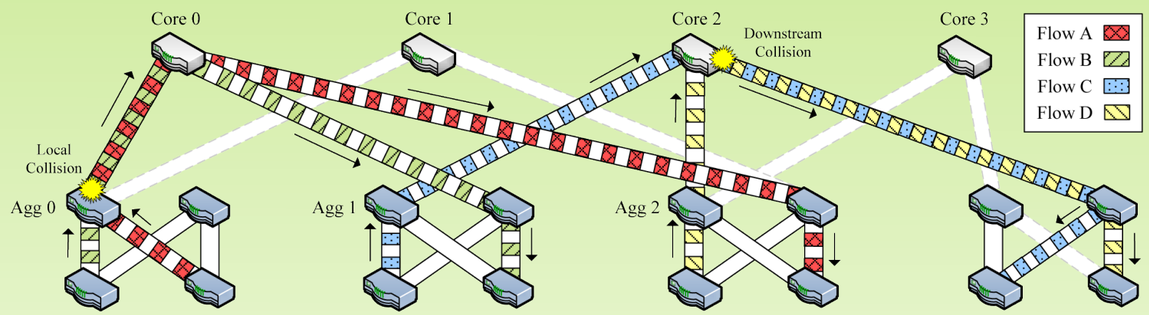

Examples of ECMP collisions resulting in reduced bisection bandwidth. Unused links omitted for clarity.

CNS Member Companies

Networking Systems, Networking People

Some of the brightest minds in systems and networking research can be found at CNS. Not only can companies collaborate on projects carried out within the Center, but they can also make early connections to tomorrow's next generation of star researchers while they are still in graduate school. This can be done informally through networking opportunities provided at CNS events, or more formally through the CNS Summer Internship Program. Since the Center's inception, dozens of graduate students have completed summer internships, many of whom have gone on to employment in industry, often at those companies where they interned.

Vikram Subramanya, an M.S. student who graduated in 2010, made contact with a high-level researcher from Google at a CNS Research Review. "After interacting for a while, I… expressed my interest in a Google internship. Readily he gave me his contact and later connected me to people in his group looking for interns. I got my interviews scheduled within a couple of days. A week later… I had an internship offer in the group exactly aligned with my research interest (data center networking)! All went well in the summer, and I have now accepted a full-time job in Mountain View, CA."

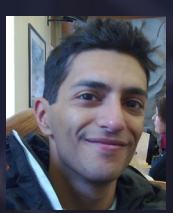

Ph.D candidate Mohammed Al-Fares interned at HP Labs under the supervision of Sujata Banerjee on a project that implemented and evaluated datac enter routing techniques. Al-Fares completed a fully operational prototype of SPAIN, a host-based, fault-tolerant L2 multipathing agent for large arbitrary data center networks, and had a paper based on this work accepted at the prestigious USENIX/ACM NSDI 2010 conference. He described working at HP Labs as "a tremendously valuable experience for me; one that allowed bridging the gap between an abstract research prototype and a robust, real-world implementation."

Email: cns@ucsd.edu
Tel: (858) 822-5964
Fax: (858) 534-7029
http://cns.ucsd.edu

Center for Networked Systems
University of California, San Diego
9500 Gilman Drive, #0404
La Jolla, Calif. 92093-0404