Technologies and Foundations for Robust and Secure Networked Systems

WInter 2011 Newsletter

Center for Networked Systems

## HELIOS: Shedding Light on the Future of Data Center Architecture

Delivering scalable bandwidth among pods, each of which can hold between 250 and 1000 servers, has emerged as a new challenge in the design of data centers. One of the key issues in supporting efficient inter-pod communication is the issue of flexibility in the placement of computation and services. Existing data center architectures address this issue through the over-provisioning of bisection bandwidth so that the network can handle the worst case, not the common case, demand. This means that if network designers want to avoid producing localized bottlenecks that they must do so by accepting the greater expense, resource allocation inefficiency, and complexity of a non-blocking topology.

To solve this issue, a team of CNS graduate students, research scientists, and faculty decided to take a novel approach. They developed and evaluated HELIOS, a fully-functioning prototype of a hybrid electrical/optical switch architecture that combines the performance benefits of a fully-functioning packet-switched network for many workloads while it delivers significant reductions in the number of switching elements, cabling, cost and energy consumption relative to all other actual or recently proposed data center network architectures. Their work on HELIOS was presented in a paper this fall at SIGCOMM 2010, a premier conference in systems and networking.

The HELIOS prototype "consists of 24 servers," explained Nathan Farrington, a CSE Ph.D. student (pictured above) who is the lead author on the paper, "[as well as] one Glimmerglass 64-port optical circuit switch, three Fulcrum Monaco 24-port 10 GigE packet switches, and one Dell 48-port GigE packet switch for control traffic and out-of-band provisioning." In addition to the hardware components, HELIOS is comprised of three pieces of software: A Topology Manager(TM), a Circuit Switch Manager (CSM), and a Pod Switch Manager (PSM). Farrington continued, "The TM is the heart of HELIOS, monitoring dynamically shifting communication patterns, estimating the inter-pod traffic demands, and calculating new topologies/circuit configurations." Part of the TM's importance stems from the fact that it manages the control loop, and the length of the control loop is a key factor in the speed with which HELIOS responds to changing communication patterns.

*Continues on page 3*

## GentleCool Reduces Impact to Environment... and Bottom Line

High temperatures negatively impact the reliability and leakage power of modern data centers and enterprise servers, so effective cooling of servers is a matter of vital concern. The most common method to remove excess heat is forced convection, which some studies suggest can account for up to 51 percent of the overall server power budget, which also has a significant effect upon the bottom line. In an attempt to lower the costs associated with fan power, various methods to lower fan usage and speed have been explored. For example, most data centers currently leverage virtualization to improve server utilization and resource efficiency, yielding significant savings in energy usage.

The problem with the current system, explains Raid Ayoub (left), a CSE Ph.D. student and lead author on a paper exploring a more proactive workload scheduler for data centers, is that "modern job schedulers do not consider either thermal/cooling characteristics of the machines, or the inherent power characteristics of the workload running in the system." This means that the efficiency gained through the use of virtualization has the unfortunate side effect of increasing the chances of creating damaging thermal hotspots in the system.

*Continues on page 7*

## CNS Members

## In This Issue



## CNS Student Wins Scholarship from Yahoo!

In November, Yahoo! awarded a CNS student from the CSE department, Malveeka Tewari (pictured below), a scholarship for writing the winning entry in its Random Hacks of Kindness (RHoK) essay contest. The scholarship paid for Tewari's travel in early December to the third annual RHoK conference, a multi-site event that brings together information and communication technology personnel with development experts to craft innovative solutions to real-life problems. This year's event themes were climate change and disaster risk management.

The theme of the contest was "Why I Want to Change the World." In her entry, Tewari states: "I want to ensure the future generations a safe world with equal opportunity and consider this my social responsibility… I have always hoped that my education will help me in building tools that will have social outreach — tools that will benefit others across any boundaries."

The full essay as well as an interview with Tewari can be found at: *http://developer.yahoo.com/blogs/ydn/posts/2010/11/rhok-scholarship-changing-the-world/*.

## CNS Interim Director Named ACM Fellow

CNS Director Stefan Savage (right) is one of 41 computer scientists recently named a 2010 ACM Fellow. He was cited "for contributions to large scale systems and network security."

The ACM, or Association for Computing Machinery, is the world's largest educational and scientific computing society. The ACM Fellows Program, initiated in 1993, celebrates the exceptional contributions of the leading members in the computing field. Fellow is the ACM's most prestigious member grade and is bestowed upon only the top one percent of ACM members. Fellows work at the world's leading universities, corporations, and research labs and are recognized for their outstanding accomplishments in computing and information technology, and/or outstanding service to ACM and the larger computing community.

Professor Savage's career has focused on the security threats enabled by broad Internet connectivity; including worms, viruses, denial-of-service attacks, botnets and spam. He is known for advancing a quantitative approach towards computer security, including the development of empirical techniques to measure and analyze global-scale attacks and efforts to identify the economics driving modern attackers.

## Graduating Students

**Priti Aghera,** an MS student working under CSE professor Tajana Rosing, submitted her thesis in fall 2010. She has accepted a position as a Software Engineer at Broadcom.

**Denis Dondi**, a Postdoctoral Researcher, completed his work in Tajana Rosing's System Energy Efficiency Lab and left to become a member of the research faculty at Università di Modena e Reggio Emilia.

**Krishnam Raju Indukuri,** an ECE MS student, also graduated after working with Professor Rosing. He is currently employed as a Digital Design Intern at Qualcomm.

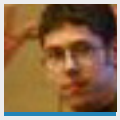

**Ryan Braud,** a PhD student working with CSE professor Amin Vahdat, successfully defended his dissertation in the fall quarter of 2010. He will be joining Thousand Eyes as a Software Engineer.

## HELIOS Hybrid Electrical/Optical Switch Architecture

*Continues from page 1...* In the past, the up-front business costs and risks associated with data center expansion have been considerable. Businesses needed to guess at their long-term data center growth needs and then construct climate-controlled spaces to house the number of machines demanded by their current operating needs while also planning for the integration of new switches and the resolution of performance bottlenecks as their operations expanded. But as an increasing number of industries have found uses for large-scale data centers, the demand for more efficient ways to build and operate them has created a market for new approaches. One successful solution has been the development of modular data centers, popularly referred to as pods.

In contrast to the original method for data center formation, pods are self-contained units with no external HVAC requirements that can be stored in any secure interior or exterior space, thereby reducing the majority of the brick and mortar costs associated with data center expansion and maintenance. Additionally, the pod vendor completes the bulk of the design and configuration of the data center prior to shipment and so obviates the company's need to recruit and employ in the long term their own specialized data center IT personnel. The result is a "pay-as-you-go" scalable data center that yields a reduction in time to installation from months to weeks with a radical decrease in costs in terms of technical expertise, energy, and capital expenses.

HELIOS achieves efficiency in part by allocating resources from an available pool of bandwidth on the basis of dynamically changing communication patterns instead of provisioning at all times for the worst case communication requirements. The research team evaluated HELIOS against a fully-provisioned electrical packet-switched network arranged in a traditional multi-root tree topology. The results of the evaluation were significant: when compared in a simulation against a traditional topology, HELIOS performed with a reduction in cost by a factor of three ($40M in savings) and a reduction in power consumption by a factor of 9 (800 kW). Though the research team details some concerns and optimizations which must be taken into consideration and utilized in order to realize these benefits, Farrington concluded, HELIOS "can deliver performance comparable to a non-blocking electrical switch with significantly less cost, energy, and complexity."

A copy of the 2010 SIGCOMM paper can be viewed at *http://www.nathanfarrington.com/papers/helios-sigcomm10.pdf*

## CNS Advisory Board Votes to Fund Seven New Projects

The CNS Advisory Board, comprised of representatives from CNS industry members, met this past August at the 2010 Summer Research Review and voted to fund seven new projects for the 2010-'12 support cycle. The projects span a wide spectrum of current hot topics in systems and networking research: routing in wireless networks; data center switch architecture; novel approaches to web security; and management of networked devices and applications.

This year the Advisory Board selected the following seven proposals as beneficiaries of the grant program:

- **Congestion Aware Routing for Mesh Wireless Networks**
  Tara Javidi and Alex C. Snoeren
- **High-fidelity Datacenter Emulation**
  Ken Yocum and Amin Vahdat
- **Sender Anonymity for Mobile Clients in Location-Based Services**
  Alin Deutsch
- **Context in the Loop: Context Aware Management of Collaborative Mobile Systems and Applications**
  Tajana Rosing and Piero Zappi
- **Networked Device Drivers**
  Cynthia Taylor, Joseph Pasquale and Amin Vahdat
- **Monetizing Stolen Financial Credentials and Web Resources**
  Chris Kanich, Kirill Levchenko, Stefan Savage, and Geoffrey M. Voelker
- **Efficient Data-Intensive Scalable Computing**
  George Porter and Amin Vahdat


CNS accepts proposals for the CNS Research Grants Program every June. These research grants generally run for two years and are awarded to support cutting-edge research in the area of systems and networking.

The grant program is the core of CNS's research activity, with the work of many past projects resulting in the further awarding of large federal grants, groundbreaking publications in top conferences, Ph.D. dissertations and best-paper awards. Researchers are invited to present their proposals at the summer Research Review, when the CNS Advisory Board meets to vote on which of the projects will receive funding.

*(Left) CNS Research Scientist George Porter is PI on the new efficient data-intensive scalable computing project.*

## Security Researchers Play CAPTCHA with Spammers

The World Wide Web was initially conceived of as a vast open resource accessible to all participants. However, the misuse and abuse by spammers of such resources as free web accounts or access to discussion forums prompted the need to erect barriers to exploitation. The difficultly lay in devising security methods that could maintain the availability of these web resources to legitimate users while they bar access to deceitful and fraudulent ones. About ten years ago, the CAPTCHA ("Completely Automated Public Turing test to tell Computers and Humans Apart") was developed to accomplish this task.

CAPTCHAs are "reverse Turing tests" -simple tests that distinguish humans from computers- that were developed to guard open web resources from mass, automated exploitation by malicious or abusive agents. CAPTCHA development involves a careful balancing of elements. The CAPTCHA needs to be distorted enough that it is rendered unreadable using current implementations of computer vision technology while not becoming so illegible that deciphering it becomes a deterrent to use for potential customers. That CAPTCHAs have proven resistant to automatic solving techniques is testified by their wide deployment as a security device for such diverse activities as purchasing tickets, registering email accounts, and posting comments on popular web sites. Indeed, the use of CAPTCHAs has reversed the normal role of spammers and defenders in the internet security arms race. Whereas defenders generally occupy the unenviable position of having to update constantly their security measures against a flood of innovations generated by the adversary, with CAPTCHAs, the adversaries who seek to employ automatic solving algorithms have had to work to stay abreast of the latest types of CAPTCHAs.

However, the economic value of the resources guarded by CAPTCHAs remains and so incentivizes the development of new methods by malicious actors to solve them. In recent years, easy and cheap access to the Internet has combined with the low cost of human labor in some parts of the globe to supply CAPTCHA-solving service providers with an effective solution: recruiting workers from the lowest-cost labor markets to solve

*Playing CAPTCHA: CNS Ph.D. student Marti Motoyama*

CAPTCHAs that are employed to guard high value resources. "Since CAPTCHAs are only intended to obstruct automated solvers," explains CSE Ph.D. student Marti Motoyama (pictured above), "their design point can be entirely sidestepped by outsourcing the task to human labor pools."

While the very existence of this global market for CAPTCHA solutions could be considered proof of the economic value of CAPTCHA-protected resources vis-à-vis the cost of the labor involved with solving passing the tests, no empirical study has yet to be conducted of this market. Additionally, the economic feasibility of a spammer business model which employs a CAPTCHA-solving provider has also not yet been considered. Motoyama is the lead author of a novel study that evaluates "the behavior and dynamics of CAPTCHA-solving service providers, their price performance, and the underlying labor markets driving this economy."

The first documented proof of the existence of a CAPTCHA-solving service was in 2006. Since then the business model was evolved and even standardized. "Today, there are many service providers that can solve large numbers of CAPTCHAs via on-demand services with retail prices as low as $1 per thousand," continues Motoyama. And since the current industry has the capacity to solve up to one million CAPTCHAs a day, this means that the retail price for this daily output is only $1000.

This stunning fact then begs the question: since CAPTCHAs can be solved so simply and cheaply, can they still be considered useful as security devices against mass exploitation? The answer, the study authors argue, depends on the value of the resources being protected. For lower value web properties such as less popular blogs, the cost of solving the CAPTCHA in order to gain access to the property remains prohibitive and so it can be argued that CAPTCHAs will continue to dissuade abuse. However, for an adversary with a sufficiently efficient business model such

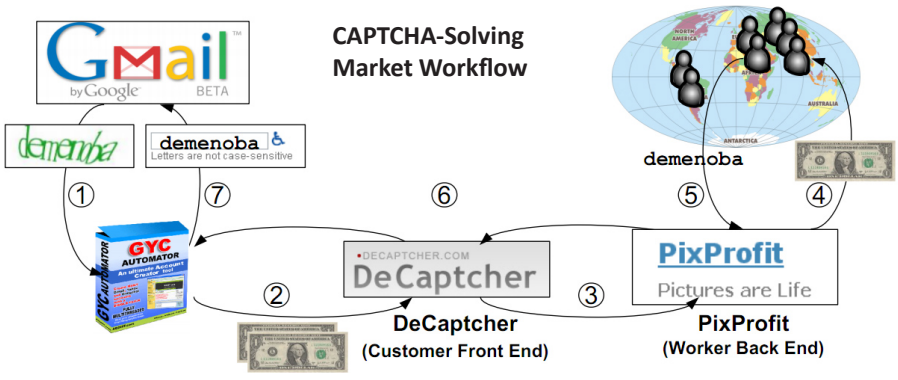

as a spammer with a successful spam campaign running from fraudulently acquired web accounts, purchasing CAPTCHA-solving services remains a fiscally sound strategy.

As a part of their study, Motoyama and his team traced the history of the wages of CAPTCHA-solving workers, which appears to have declined as solving efficiency of the workers increased and the source of workers moved to countries with increasingly lower labor costs. The researchers additionally conducted experiments to estimate the likely geographic location of most CAPTCHA solvers (currently in China, Bangladesh, India, and Japan) and to gauge the range of quality in relationship to service pricing that currently exists on the open market (their evidence indicates that money buys greater accuracy).

The work provides a new perspective to the evaluation of Internet security measures. "CAPTCHAs, while traditionally viewed as a technological impediment to an attacker, should more properly be regarded as an economic one," explain the researchers. "The profitability of any particular scam is a function of three factors: the cost of CAPTCHA-solving, the effectiveness of any secondary defenses (e.g. SMS validation, account shutdowns, additional CAPTCHA screens, etc.), and the efficiency of the attacker's business model." In the end, they argue, defenses don't need to be impregnable, they only need to be prohibitively expensive to solve. This means that as the labor cost of solving CAPTCHAs falls, owners of high value web properties should employ more secondary defenses to further increase the cost of breaching their security. Conversely, owners of lower value web resources can continue to rely upon CAPTCHAs..

The full paper can be found at: *http://www.usenix.org/events/sec10/tech/full_papers/Motoyama.pdf*

## CNS Researchers Architect Future of the Internet

In August 2010, the Directorate of the NSF announced awards for four new projects, each worth up to $8 million over three years, as part of its Future Internet Architecture (FIA) program. The purpose of FIA is to challenge researchers to overcome the current limitations of the Internet through the proposal and development of "new Internet architectures that hold promise for the future." These awards will enable researchers at dozens of institutions across the country to pursue new ways to build a more trustworthy and robust Internet.

CNS researchers affiliated with the Cooperative Association for Internet Data Analysis (CAIDA), KC Claffy and Dmitri Krioukov, will be participating in one of the four projects, "Named Data Networking" (NDN). The project will be led by Lixia Zhang of UCLA and will also involve work by researchers working at Colorado State University, PARC, University of Arizona, University of Illinois at Urbana-Champaign, UC Irvine, University of Memphis, Washington University, and Yale University.

*By the time these students are in college, the Internet will be dramatically different [NSF photo by Karen Randall]*

The current approach for handling communication over the Internet is based upon an older client-server use model where parties transferred information contained in IP packets transported along a single path. Today's model focuses upon the "where" of data- i.e., addresses, servers, and hosts. While this approach enables the development of any number of successful applications and services, the use pattern of the Internet has largely shifted in the last decade. The Internet today is used largely for content creation, dissemination, and delivery, a paradigm that is only likely to expand in the future. Since the traditional client-server model does not incorporate adequate mechanisms to support secure content-oriented functionality, regardless of the specific physical location where the content resides, there are inherent limitations placed upon routing scalability, fast forwarding, trust models, network security, and content privacy and protection. NGN will create a new model, based on the "what" of communication- i.e., the content that users and applications care about- that will allow Internet users to communicate with greater privacy, reliability, and efficiency.

"One research effort to which the CAIDA team will contribute is how to efficiently find optimal communication paths directly through the hierarchical content name space," said Dmitri Krioukov. This approach allows the decoupling of trust in data from trust in hosts and servers, enabling trustworthiness as well as several radically scalable communication mechanisms, for example, automatic caching to optimize bandwidth and the potential to move content along multiple paths to the destination. NGN will develop an Internet architecture that will allow users to optimize bandwidth and move data along multiple paths to a destination.

*CAIDA researchers Dmitri Krioukov (left) and KC Claffy*

CAIDA is at *http://www.caida.org/home/*. For more on the Future Internet Architecture intiative, go to: *http://www.nsf.gov/news/news_summ.jsp?cntn_id=117611&org=OLPA* .

## CNS Database Researchers At Work on Four NSF Projects

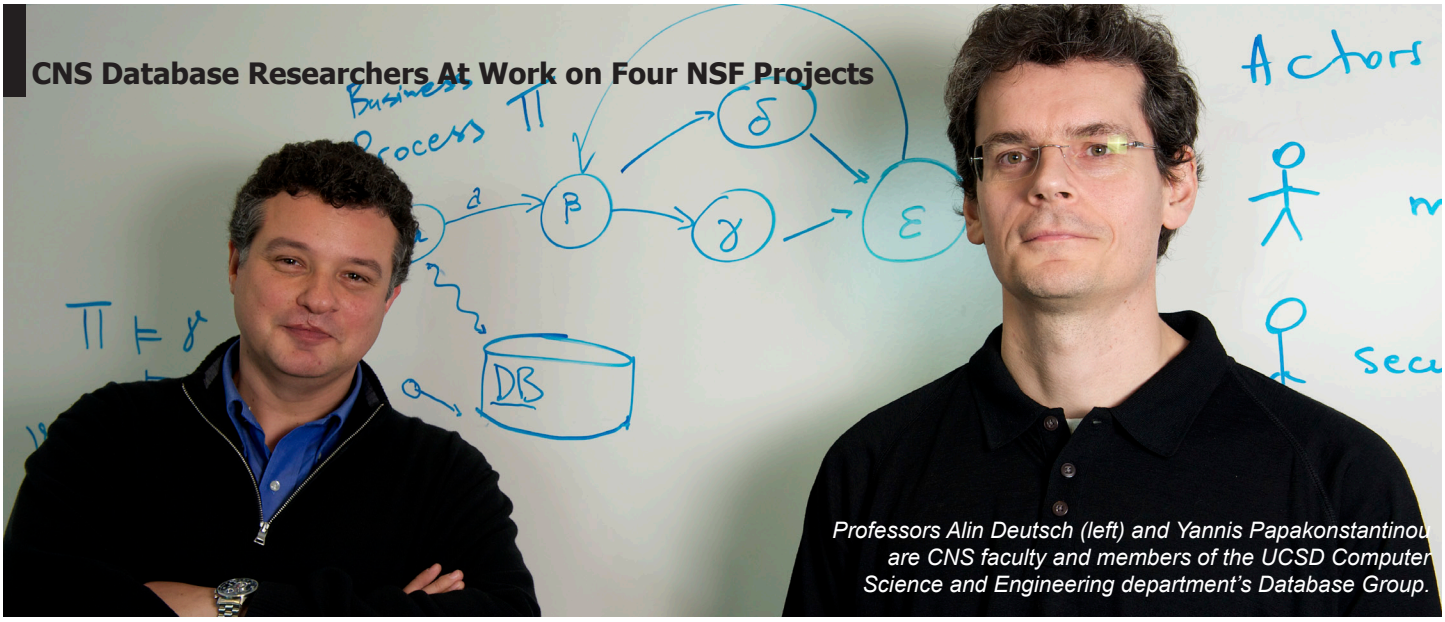


*Professors Alin Deutsch (left) and Yannis Papakonstantinou are CNS faculty and members of the UCSD Computer Science and Engineering department's Database Group.*

CNS faculty from the Database Group in the Jacobs School of Engineering's Computer Science and Engineering department—Professor Yannis Papakonstantinou and Associate Professor Alin Deutsch—are working on four NSF grants through summer 2012 at the intersection of web and data management systems.

**ASTERIX: A Highly Scalable Parallel Platform for Semistructured Data Management and Analysis:** This is a collaboration among UC researchers in San Diego, Irvine, and Riverside. ASTERIX is being built to meet the challenge of how to apprehend the growing sea of semi-structured data on the web. The project combines ideas from three distinct research areas—semi-structured data, parallel databases, and data-intensive computing—to create a next-generation, open source software platform that scales by running on large, shared-nothing computing clusters. The project—led on the UCSD end by PI Deutsch (below) and co-PI Papakonstantinou (above right)—will develop new technologies to manage, index, query, analyze, and subscribe to vast quantities of semi-structured information. [NSF#0910820] *http://db.ucsd.edu/asterix/*

**Database-Driven AJAX Applications:** AJAX programming enables superior performance and interface quality that is comparable to desktop applications, despite the fact that the user's browsers interact with a remote server on the cloud. The challenge is that AJAX programs are hard to write since they are essentially distributed programs combining three different languages: browser-side Javascript, a server side programming language (e.g., Java, PHP, etc) and database access with SQL. PI Yannis Papakonstantinou aims to deliver a development framework for the rapid creation of fully-fledged AJAX-based web application pages from declarative, data-driven, SQL-based specifications. [NSF #1018961] *http://db.ucsd.edu/NSF10FWD/*

**FORWARD: Do-It-Yourself Forms-Driven Workflow Web Applications:** Organizations that cannot afford the time and money needed to engage in the conventional code development process can benefit increasingly from emerging "do-it-yourself" (DIY) database-driven web application platforms. Yet these platforms must maximize two metrics that present an inherent trade-off: the simplicity of specification, and the application scope. PI Yannis Papakonstantinou and co-PI Alin Deutsch are working on an unlimited model of web applications, where programmers introduce code components and interface them with the 'limited' part via queries (reports) and updates (forms and actions). The proposed models of database-driven web applications will impact the education of both Computer Science (CS) and non-CS students who need to comprehend web applications at a high conceptual level. [NSF #0917379] *http://db.ucsd.edu/NSF10FWD/*

**Data-Centric Business Processes: Specification and Static Analysis:** Classical software verification techniques applicable to static analysis (notably model checking and theorem proving) have serious limitations. An alternative approach, introduced by IBM, focuses on data records, known as "business artifacts," that correspond to key business-relevant objects, their life cycles, and the services or tasks that are invoked on the artifacts. The main objective of this UCSD project is to develop new tools for the high-level specification and verification of data-centric business processes. Co-PI Deutsch and PI Victor Vianu are investigating the trade-offs between the expressiveness of the specification language and the feasibility of the analysis tasks. They are developing tools to carry out analysis tasks—tools that will be integrated with Siena, IBM's prototype for compiling artifact-centric business process specifications into workflow support code. This research is partly supported by IBM. [NSF#0916515]
*http://db.ucsd.edu/artifacts/*

The FORWARD effort has also received support from CNS member company Google.

## GentleCool (continued from p. 1)

To address this issue, the Ph.D. student and his research team developed a new algorithm, which they dubbed 'GentleCool', that takes a proactive multi-tier approach to lower fan cooling costs in highly utilized systems through the intelligent allocation of workload across different physical machines.

The GentleCool scheduler works in two phases, spreading followed by refinement. During the spreading phase, GentleCool balances the power density among the PMs/CPU and so lowers the fan speed. During the refinement phase, GentleCool further reduces cooling costs by balancing the load of hot VMs/threads so that the use of fast fans is maximized. And because the algorithm only runs once every ten minutes, the overhead cost for its run time is negligible- a range of a few tenths of milliseconds depending upon the number of machines in the system.

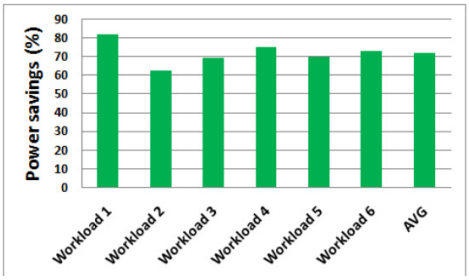


*GentleCool offers fan power savings (above) averaging more than 70%, and total power savings including CPU costs (below) in excess of 15%.*

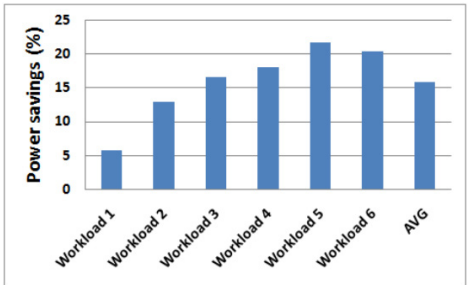


The use of GentleCool achieves results that are both significant in terms of environmental impact and the bottom line: up to a 72 percent reduction in cooling costs and an improvement of mean time between failures of the fans by 2.3 times compared to the state of the art.

The GentleCool article is available at *http://seelab.ucsd.edu/papers/rayoub_date10.pdf*

## Cisco Supports CNS Data Center Research

CNS member company Cisco Systems recently provided two gifts in support of research in data-intensive computer and computer security.

The first funded project, "Efficient Data-intensive Scalable Computing," is led by CNS Assistant Research Scientist George Porter. Porter's team will work to build a Data- intensive Scalable Computing (DISC) system that will maximize the efficiency and cross-resource balance. DISC is an increasingly important approach to solving web-scale problems faced by search engines that answer queries based on hundreds of terabytes of input data. The purpose of these search engines can be broad-ranging, from biological applications that look for subsequence matches within long DNA strands to scientific applications for radio astronomers that process deep space radio telescope data. The key to achieving this scale of computing is harnessing a large number of individual commodity nodes, so the goal of the project is to design DISC systems that focus on efficiency as a first-class requirement. To achieve this, Dr. Porter will develop programming models and tools to build balanced systems in which the key resources of networking, computing, and storage are maximally utilized while also providing fault tolerance and maintaining efficient use of computing and storage resources.

The second project, "Identifying Malicious URLs Using Machine Learning," is led by CSE Project Scientist Kirill Levchenko (pictured at right). Currently, the Web is a platform that hosts a wide range of criminal enterprises such as spam-advertised commerce, financial fraud, and as a vector for propagating malware. One common method for compromising individual's security is by luring unsuspecting users to malicious web sites. These visits can be driven by email, Web search results or links from other Web pages, but all require the user to take some action, such as clicking, that specifies the desired URL.

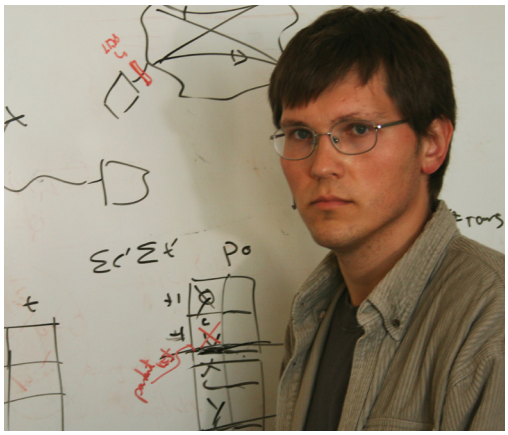

*CSE project scientist Kirill Levchenko*

One approach to protecting Web users from visiting malicious URLs has been the development of blacklisting services that provide users with warnings about the danger of visiting particular sites. These blacklists are in turn constructed by a range of techniques including manual reporting, honeypots, and Web crawlers combined with site analysis heuristics. Unfortunately, blacklisting services are only as good as their intelligence gathering, and so many malicious sites are not blacklisted. The reasons for this are many: the pages can be too new to have been discovered, have never been evaluated, or have been evaluated incorrectly (e.g., due to "cloaking"). To address this problem, some client-side systems analyze the content or behavior of a Web site as it is visited. But, in addition to run-time overhead, these approaches can expose the user to the very browser- based vulnerabilities that we seek to avoid.

In this project, Levchenko's team will focus on a complementary part of the design space: classifying the reputation of a Web site entirely based on its inbound URL. The motivation is to provide inherently better coverage than blacklisting-based approaches while avoiding the client-side overhead and risk of approaches that analyze Web content on demand. Sites will be classified by examining the relationship between URLs and the lexical and domain name features that characterize them. The justification for using lexical features is that URLs to malicious sites tend to "look different" in the eyes of the users who see them. Hence, including lexical features will allow the researchers methodically to capture this property for classification purposes, and perhaps infer patterns in malicious URLs that would otherwise be missed through ad-hoc inspection.

This support from Cisco is only the latest in a number of successful research collaborations that Cisco and CNS have initiated since Cisco joined the Center in 2006.

## CNS Lecture Series

CNS Lectures feature key players in the research and development of systems and networking technology and are free and open to the public. Fall lectures included:

• **Katerina Argyraki**, Network Researcher at EPFL Switzerland, presented "Verifiable Network-Performance Measurements" on November 16, 2010.

• **Xiaowei Yang**, Assistant Professor of Computer Science at Duke University, spoke October 29, 2010 on "Bootstrapping Accountability in the Internet We Have."

• **Benjamin Reed**, Research Scientist at Yahoo!, spoke October 7, 2010 on "Zookeeper: Because Distributed Computing Is a Zoo."

• **Roger Dingledine**, Project Leader of The Tor Project, spoke August 23, 2010 on "Tor: Anonymous Communications for the Department of Defense... and You."

To view webcasts or Powerpoint presentations of past lectures, visit the Lectures Archives page under the "News and Events" tab on the CNS Web site.

### Get Connected

Stay up-to-date about upcoming CNS events, including lectures and Research Reviews, by signing up for the CNS Events RSS feed. To do so, visit the CNS website at

**http://cns.ucsd.edu**

and click on the link "CNS Events RSS Feed."

## Upcoming Events

February 2 and 3, 2011
**CNS Winter Research Review**
Faculty Club, UC San Diego Campus in La Jolla, Calif.

The winter Research Review will feature keynotes, including talks by two of our industry affiliates; a roundtable panel hosted by our member companies to discuss their top current research concerns; progress reports from recipients of the CNS Research Grants Program; a graduate student research poster session and reception; and numerous opportunities for informal interactions among member companies and CNS faculty, researchers and graduate students.

Attendance at the Winter 2011 Research Review is limited to industry sponsors and invited guests.
For more information, please contact *kkrane@ucsd.edu* or call 858-822-5964.

## 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.