Center for Networked Systems

WINTER 2015 NEWSLETTER

TECHNOLOGIES AND FOUNDATIONS FOR ROBUST AND SECURE NETWORKED SYSTEMS

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Ph.D. Student Michael Conley (above) at
Poster Session of CNS Fall 2014 Research Review (below)

FEELING GOOD ABOUT SORTING THINGS OUT

In continuing what is becoming a tradition, Ph.D. student Michael Conley once again won world records for data sorting in multiple categories. He was competing in the annual Sort Benchmark competition with CNS associate director George Porter and former CNS director Amin Vahdat (now at Google). Conley employed an updated version of their sorting system, Tritonsort, which was designed not only to achieve record-breaking speeds but also to maximize system resource utilization. Tritonsort tied for the "Daytona Graysort" category and won outright in both the "Daytona" and "Indy" categories of the new "Cloudsort" competition.

The established metric to win the GraySort competition is to achieve the highest sort rate (terabytes of data per minute) while sorting a minimum of 100 TB of data. The new CloudSort competition measures the minimum cost of sorting at least 100 TB of data on a public cloud. Conley's group achieved their victory with Tritonsort on Amazon's Elastic Compute (EC2) cloud computing platform.

Several factors underscore the effectiveness of Tritonsort's system resource utilization scheme as compared to the far more resource-intensive methods followed by their competitors. The team that tied with CNS for first place in the Daytona Graysort category, Databricks, used 10 percent more processors to achieve a slightly (though for the purpose of the contest, equivocal) slower sorting speed.

Tritonsort also continues to hold the world record for energy efficiency in data sorting. The 2011 iteration of Tritonsort still holds the record for the least amount of energy required to sort 100 TB of data in the "Joulesort" competition. The 2011 record set by Tritonsort still holds at 132 MJoules.

To see all the details about the competition: http://sortbenchmark.org

And to read more about the technical aspects of Tritonsort: http://sortbenchmark.org/TritonSort2014.pdf

INDUSTRY MEMEBERS

Google Microsoft CSR vmware

IN THIS ISSUE

A FELLOW HERE, A FELLOW THERE

The past year has seen recognition for the excellence and importance of CSE professor Yuanyuan (YY) Zhou's research by two of the premier professional associations in computer science and engineering.

In December 2013, the Association for Computing Machinery (ACM) announced its annual recognition of members from academic institutions, companies, and research labs all over the world who "have achieved advances in computing research and development that are accelerating the digital revolution and impacting every dimension of how we live, work, and play." Zhou is one of 50 ACM members elevated to be Fellows of the ACM in 2014. She was cited for her "contributions to software reliability and quality." Professor Zhou joins other CNS members Keith Marzullo and Stefan Savage as ACM Fellows.

More recently, in November 2014, professor Zhou was named an IEEE Fellow for "contributions to scalable algorithms and tools for computer reliability." For more than a century, IEEE has conferred the distinction of Fellow upon those of its members with extraordinary accomplishments in any of the fields of endeavor that are of interest to IEEE.

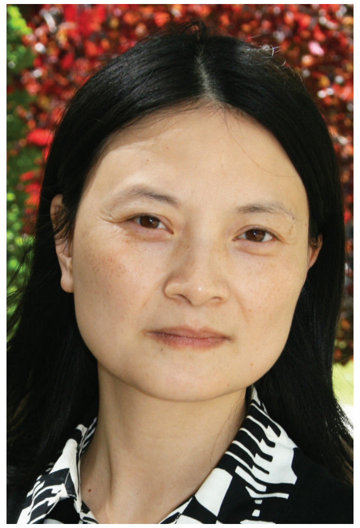

YY Zhou

kc claffy

RESEARCHER RECEIVES 2014 IEEE INTERNET AWARD

The head of the Center for Applied Internet Data Analysis (CAIDA) research group at the San Diego Supercomputer Center, kc claffy, was named one of the two recipients of this year's IEEE Internet Award. She was cited (along with UC Berkeley's Vern Paxson) for their "seminal contributions to the field of Internet measurement, including security and network data analysis, and for distinguished leadership in and service to the Internet community by providing open access data and tools."

"kc was a Ph.D. student here and, unable to escape the gravity of San Diego, created a completely independent research group, CAIDA, that single-handedly established UC San Diego as a worldwide leader in network measurement," said CNS director Stefan Savage, a professor in the Department of Computer Science and Engineering (CSE). "We're lucky to have her as a colleague, an adjunct in CSE, and as a member of the Center for Networked Systems."

Begun in 1999, the IEEE Internet Award is given to an individual or small number of collaborators who have provided exceptional contributions to the advancement of Internet technology for network architecture, mobility, and/or end-use applications.

RANJIT JHALA RECEIVES "TEST OF TIME" AWARD

Sometimes the best way to judge the importance of work is to view it in hindsight. An idea that might have seemed important or groundbreaking at the time of its publication can prove to be a dead end, while a seemingly modest proposal that was overlooked when it first came out eventually becomes the standard for how things are done in a particular field. With this phenomenon in mind, many sub-fields in computer science have taken to presenting annual "test of time" awards for work that has proven to be most influential.

Every year at the ACM Symposium on Principles of Programming Languages (POPL), the program committee announces the recipient of the POPL Most Influential Paper Award. The committee goes back to the papers presented at the symposium a decade earlier and determines which of the papers proved most influential in retrospect.

In January 2014, the POPL committee recognized the 2004 paper, "Abstractions from Proofs," co-authored by CNS member and CSE professor Ranjit Jhala while he was finishing up his doctoral work at UC Berkeley with colleagues Tom Henzinger, Rupak Majumdar, and Ken McMillan. The committee explained that, in this paper, Jhala and his team "demonstrated a fundamental generalization of Craig interpolation to program analysis by predicate abstraction, opening the door for interpolation to be applied to abstraction refinement for 'infinite-state' systems." Prior to the publication of Jhala's paper, interpolation had only been used by those researching programming languages in the model checking of "finite-state" systems. The award committee further praised how the 2004 paper "showed how interpolation offers a fundamental way to explain abstraction refinement in a logical framework, and has led to many extensions to increase the power of abstraction in program analysis."

Ranjit Jhala

IN THE NEWS

WHO WATCHES THE WATCHMEN?

If you have traveled by air in the United States in the last few years, chances are that you have passed through a security checkpoint that was monitored by some kind of advanced imaging technology (AIT). But how effective is this technology at scanning for weapons and explosives? Government agencies have assured the public and policy makers that the technology has been thoroughly tested for efficacy, but because details about these tests have not been forthcoming, and existing studies were only conducted by manufacturers, significant doubts about the machines have persisted.

In a groundbreaking paper presented at the 2014 USENIX Security Symposium, a group of CNS students and faculty at UC San Diego, partnering with researchers at Johns Hopkins University and the University of Michigan, conducted the first independent security evaluation of an AIT. Specifically, they tested whether individuals could smuggle weapons or explosives through security checkpoints by exploiting flaws in the design of an X-ray backscatter scanner, the scanner's software and hardware, and scanning protocols.

The U.S. Transportation Security Administration (TSA) adopted AITs in 2009 to screen for both metallic and non-metallic contraband such as knives, personal firearms, and explosive devices or materials that passengers might smuggle onto airplanes. The primary deployment of AITs from 2009 to 2013 took the form of an X-ray backscatter scanner called the Rapiscan Secure 1000. Though the Rapiscan was removed from service by TSA and replaced with a different kind of AIT in 2013, it has also been widely distributed to other secured venues throughout the U.S. (such as courthouses and prisons), so it remains widely though less visibly in use.

With the Rapiscan, travelers are asked to enter the machine and face it with their feet spread and their hands raised while they are illuminated with constant-spectrum X rays. Because different atomic compositions of materials reflect the radiation at different intensities, the pattern of radiation reflected back to the sensor creates an image. In this image, thin and more diffuse materials (such as clothing) become invisible, while less dense materials (e.g., flesh) can be discerned in contrast to very dense materials (such as steel or aluminum).

When first introduced in 2009, the scanners were presented as having been tested, effective, safe, and necessary components of the TSA's physical security system. However, the systems were almost immediately criticized for a number of perceived or hypothesized flaws. The technology, as well as the search protocols used to implement it, was considered by some to be ineffective and even potentially hazardous. What made the discussion difficult to resolve, however, was that while independent observers were not allowed to conduct their own evaluations of the technology, the existing studies by the manufacturer, if released at all, were presented in a highly redacted form that made informed criticism impossible. Perhaps more troubling to some observers, it was questionable how thoroughly the previous studies had tested the technology by trying to mimic motivated, malicious actors who might adapt their behaviors to circumvent the device.

(STORY CONTINUED ON PAGE 8)

usenix
THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

(Top) CSE Ph.D. student Keaton Mowery and professor Hovav Shacham, co-authors of the study presented at USENIX Security 2014; (below) Shacham in the Rapiscan Secure 1000 X-ray backscatter scanner

SCHOOLING AN ACADEMIC IN INDUSTRY SOLUTIONS

Mr. Jang Goes to Mountain View

When Computer Science and Engineering Ph.D. student Dongseok (Don) Jang was given the opportunity to complete a CNS internship at Google with Software Engineer Charlie Reis, he jumped at the chance. Jang knew Reis by reputation as a leader in the creation of innovative browser architectures that improve security. For a doctoral student passionate about security issues, the prospect of working under Reis on a project to improve the security of Google's web browser Chrome was impossible to pass up.

Web browser security is a constantly evolving problem of longstanding and great importance. Creators of browsers have to develop products that provide an excellent user interface while ensuring that their users are protected from the less savory elements on the Web. The company whose browser is best able to deliver to Internet users a fast and safe experience on the Web becomes the preferred provider of a window on the Internet. That company can then charge the most from advertisers who wish to place themselves into that window frame. However, creating a browser that is safe for users to operate while also rating high in ease of usability is a daunting task.

Problems with Web Browser Security

In order to understand the work upon which his internship was based, Jang provided some background into how Reis had already changed Chrome to make it more secure. Historically, browsers have been vulnerable to low-level security attacks. To prevent this from happening, Reis proposed separating the browser into smaller, self-contained pieces, each of which runs a different process. Because of this innovative measure, called site isolation, even if one process is compromised, it will not affect any of the other processes that are running in the browser. For example, a browser can have several tabs open simultaneously, such as a shopping site in one and an email account in the other. Since the operation of each is handled separately, if the user's shopping tab is compromised, site isolation prevents the email account from being affected.

As time passed and Web sites evolved, a couple of factors made it clear that site isolation needed to be extended beyond just separating the operation of tabs. Web sites became more complex, and this increasing complexity opened more doors to hackers.

It is the nature of modern Web sites that a number of documents are embedded within the Web page. For example, Web pages often include several inline frames (called iframes), which are HTML documents embedded within another HTML document on a Web site. Iframes allow documents from other sources to be part of a site's presentation to the viewer and are familiar to anyone who has ever seen a news site with a section advertising a product that, if clicked on, leads the client to a different URL. The iframe is a way for the host Web site to provide a window to other places on the Internet. However, the insecurity of iframes allows malicious parties to use low-level memory attacks to access browsers through the embedded websites. A compromised tab in a browser can then steal data from any subsequent site the user chooses to visit.

Don Jang's summer project focused on how to expand site isolation so that the browser can restrict the access of documents embedded within a website to the other processes running within that tab. The immediate solution to this conundrum would be to set up a document filter within the browser. But right at the beginning, there was a significant technical obstacle that needed to be addressed: Jang and his colleagues could not simply tell their document blocker to filter out specific categories of problematic documents. This is because Web sites often mislabel content headers that provide the identifying data about the documents being transmitted. For example, something may claim that something is HTML when it is not, or vice versa. Therefore, building a simple filter based on content header information would block some content that would be acceptable while still accepting insecure documents, making for a disrupted and yet still unprotected user experience.

The question posed to Jang when he undertook the project was: How do we block the insecure data without breaking the existing website?

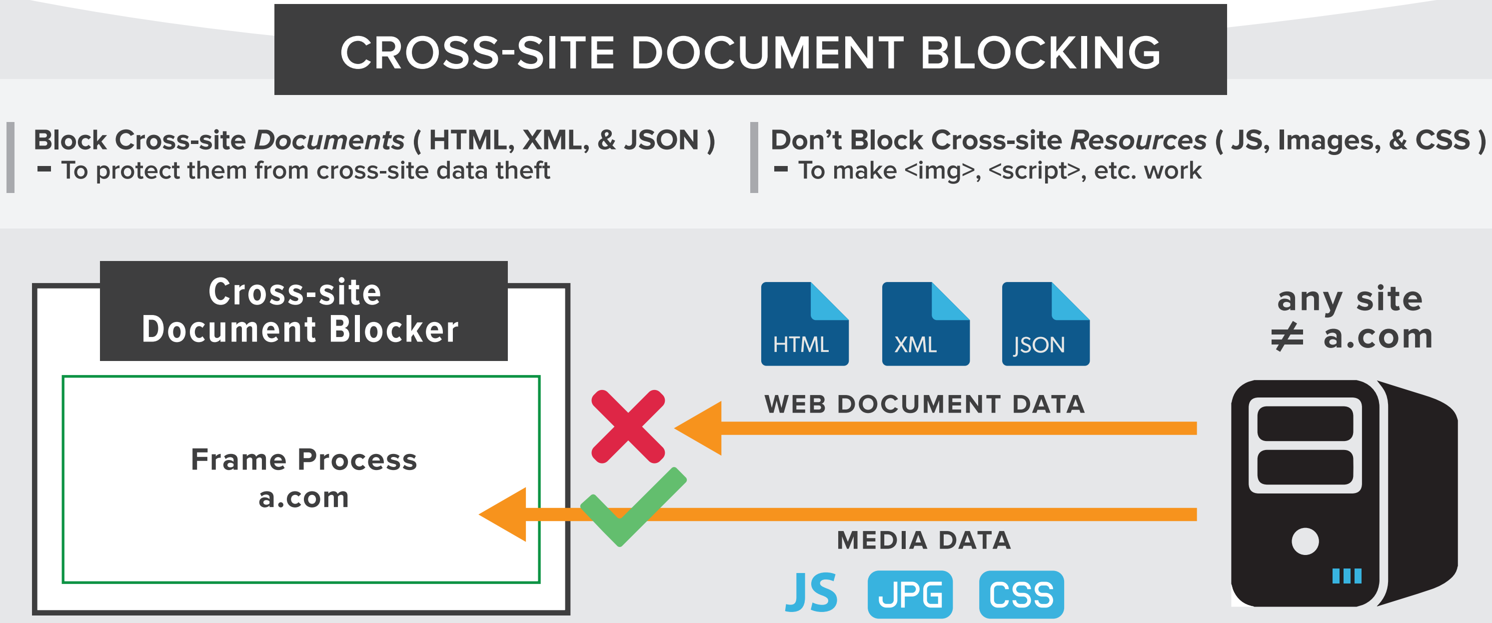

Cross-site blocking of HTML, XML and JSON documents protects them from cross-site data theft, but browsers should not block cross-site resources (JavaScript, images, CSS) because it would break the Web (in the lingo, "disruptive blocking").

The solution conceived by Jang and his research colleagues at Google was to create a cross-site document blocking policy. They designed a frame process that navigates to pages only within its own site and which will not read documents that are vulnerable to attack from other sites. For example, the frame process will read HTML files from its own site while blocking cross-site HTML files, thus preventing cross-site data theft.

Implementation and Evaluation

CNS's Jang and colleagues produced their modified version of Chrome and tested it while viewing Alexa Top 50K Web pages. During the testing process, they fine-tuned their policy by incrementally excluding more and more document types, thereby increasing browser security while also increasing the number of disruptive blocks. They then measured the likely number of definite and possible disruptive and non-disruptive blockings and used this information to estimate the probable effect the use of their implementation would have on the user experience.

After tinkering with some of the parameters of the process, Jang and colleagues lowered the disruptive blocking rate to 0.075%, a level deemed to be acceptable, especially considering the parallel increase in security that it conferred. The modification is now an active setting in Chrome.

For more information about this project, visit the Wikipedia page for "Blocking Cross-Site Documents for Site Isolation" at http://www.chromium.org/developers/design-documents/blocking-cross-site-documents.

CNS SUMMER INTERNSHIP EXPERIENCE: A WIN-WIN-WIN PROPOSITION

Dongseok (Don) Jang

While reflecting on his internship experience at Google, Ph.D. student Don Jang recalls that working on the project was "overwhelming at first. It was hard to figure out what I could contribute because Chrome is a really gigantic project. It consists of millions of lines of code with thousands of people who work on it. To make even a small change in Chrome is really hard – just to make a patch of five lines of code takes about two or three days because there are many people working on that."

However, in the end, Jang derived from his internship experience a satisfaction that everyone involved – himself, his fellow team members at Google, and users of the Chrome Web browser – walked away with something positive.

In the first place, notes Jang, the team's implementation increased the security of the browser without degrading the user's experience of ease and rapidity, thus bolstering Chrome's reputation as a premier Web browser and supporting Google's reputation in the market. Second,Chrome users gained a much safer browsing experience with only negligible effects on their viewing of content. And for Jang personally, the lessons he learned from his summer experience are likely to extend throughout his career.

The value of the internship to Jang was not merely the experience of working on a novel solution to a serious problem. Instead, it changed his whole approach to formulating research problems and solutions.

"I went to Google with a really academically-oriented mindset that I could do whatever I wanted to do," explains Jang. "So I proposed really aggressive ideas to secure Chrome." However, he soon learned that the most obvious solutions to problems are often incompatible or unworkable when seen in the context of a system's bigger picture. For example, some of Jang's ideas would have resulted in substantially slower response times for Web browsing in Chrome. In a business environment where one of the main priorities is speed, this would be unacceptable – "an example of the cure being worse than the disease," he adds. Other proposals would have resulted in a high percentage of media files being prevented from displaying on any given page. Once again, this would have caused an unacceptable degradation in the user experience, which would have rendered Chrome a less attractive product in comparison to its principal competitors.

The experienced engineers on Jang's team helped to teach him how to come up with solutions that improved security while also keeping in mind how the changes would influence other aspects of Chrome. "They helped me a lot to come up with reasonable and realistic ideas," recalls Jang. "I learned how to meet the project goal in a way that was feasible." Thus, he began to understand that what makes a solution to a given problem the "best" depends on how that solution operates within the broader context of the problem.

Additionally, Jang learned how to work with a large team to implement a far-reaching change within an enterprise-level system. "The solution [we came up with] consisted of just two hundred lines of code, [but] every single network request goes through these two hundred lines, so instead of implementing some ad hoc features of the browser, we made a big decision in the browser big picture," explains Jang. "I learned how to evaluate the possible big changes in the browser, because there are lots of different factors. I had to talk to more than 50 or 60 engineers scattered throughout the world. Someone was working in Germany, someone was in Zurich, another was in Asia." Collaborating with such a large number and more diverse group of people stood in stark contrast to Jang's prior programming experience in academia, in which he generally only teamed up with a handful of other students or faculty on a local basis.

"I had to talk to new people every day, and I became able to see different aspects of a problem," says Jang. "Beforehand, when I looked at a problem, I only looked at really small, academically-oriented ideas, things that people could see as an academic contribution to the solution. I wasn't able to think about industry-level impacts or how hard it would be to implement an idea in some corporate environments. Now I can measure how to do that."

Before his internship, Jang might have come up with a solution to a real-world problem and thought to himself, "The solution is so easy! Why don't they just implement this solution?" He now sees that some proposed solutions lack value because they cause more problems than they solve – something he now will take into consideration when proposing solutions.

Jang feels his experience as an intern at Google was invaluable, and the proof is that upon finishing his Ph.D., he accepted an offer to work full-time in Google's Los Angeles office. "With this experience, I can come up with a solution that is more easily adopted by a corporation such as Google," says Jang. He also believes that the combination of his academic credentials, CNS-based research and the internship at Google have ensured that he can produce high-quality technical work that is informed by real-world concerns and considerations.

CSE students Malveeka Tewari (l) and Soohyun Nam, recipients of CNS travel grants to attend 2014 Grace Hopper Celebration, part of a large delegation of UC San Diego alumni, faculty and students (below).

STUDENTS REPRESENT CNS AT GRACE HOPPER CELEBRATION

In an ongoing effort to support diversity in computer science and engineering, CNS selected two Ph.D. students to represent the center at the Grace Hopper Celebration of Women in Computing, organized by the Anita Borg Institute.

CNS awarded CNS 2014 Grace Hopper Travel Grants to two CSE Ph.D. students: Malveeka Tewari and Soohyun Nam. The students were chosen to attend the Phoenix, AZ, conference in November 2014. The Grace Hopper Celebration is the premier conference bringing together women in computing and technology to focus on research and careers. Tewari and Nam attended special sessions focusing on this year's theme of "the global prevalence of computer technology and the participation of one and all in its design, development, and deployment." The students were among approximately 40 young women from UC San Diego attending the 2014 conference, in a delegation led by CSE professor Christine Alvarado.

ANITA BORG INSTITUTE
GRACE HOPPER CELEBRATION
OF WOMEN IN COMPUTING

PH.D. DEGREES AWARDED

FENG LU

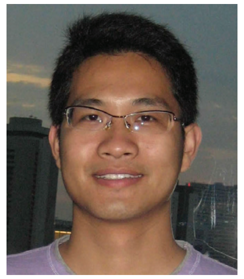

Feng Lu earned his Ph.D. in ECE in August 2014 after defending his dissertation, "Downclocking WiFi to Improve Energy Efficiency in Mobile Devices." Dr. Lu's advisors were Tara Javidi, Alex C Snoeren and Geoffrey M Voelker. He now works as a Software Engineer at Google.

DONGSEOK (DON) JANG

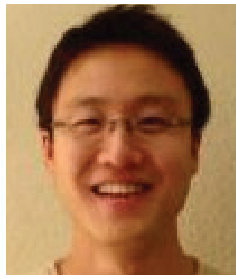

A Ph.D. student in CSE advised by Sorin Lerner, Dongseok (Don) Jang, defended his dissertation, "Language-Based Security for Web Browsers" in August 2014. Dr. Jang went on to a position at Google as a Software Engineer. [Editor's note: For more on Jang's research internship at Google prior to graduation, see page 4.]

DAVID WANG

David Wang, now a Software Engineer at Google, earned his Ph.D. in CSE after defending a dissertation on "A Comprehensive Approach to Undermining Search Result Poisoning" in September 2014. He was advised by Geoffrey M Voelker and Stefan Savage.

QING ZHANG

Qing Zhang, a CSE Ph.D. student advised by Geoffrey M. Voelker, did her dissertation on "Utilizing Source Information to Detect and Prevent Online Fraud." After graduating in October 2014, Dr. Zhang became a Software Engineer at Google.

DO-KYUM KIM

CSE Ph.D. student Do-kyum Kim became a Software Engineer at Google after presenting "Topic Modeling of Hierarchical Corpora" in September 2014. He was co-advised by Lawrence Saul and Geoffrey M. Voelker.

ARUP DE

In June 2014, CSE Ph.D. student Arup De, advised by Steven Swanson, became a member of the research staff at HGST, a Western Digital company, after defending his dissertation, "A Compute Capable SSD Architecture for Next-Generation Non-volatile Memories."

SARAH MEIKLEJOHN

CSE Ph.D. student Sarah Meiklejohn completed and defended her dissertation on "Flexible Models for Secure Systems" in April 2014. She was co-advised by Stefan Savage and Mihir Bellare. Dr. Meiklejohn is now an Assistant Professor at University College London in the departments of Computer Science as well as Security and Crime Science.

JIAQI ZHANG

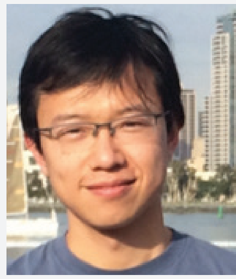

CSE professor YY Zhou advised CSE Ph.D. student Jiaqi Zhang, who defended his dissertation on "Software Configuration Learning and Recommendation" in October 2014. Dr. Zhang is a Software Engineer at Whova (a startup founded by Zhou).

SIVASANKAR RADHAKRISHNAN

In March 2014, CSE Ph.D. student Sivasankar Radhakrishnan presented "Network Performance Improvements for Web Services – An End-to-End View." He was jointly advised by Amin Vahdat and George Porter. Dr. Radhakrishnan is now a member of the technical staff at Forward Networks.

M.S. DEGREE AWARDED

DEVIN LUNDBERG

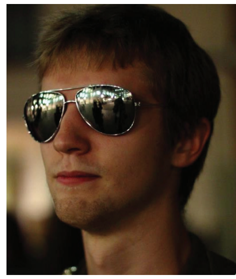

CSE M.S. student Devin Lundberg graduated in April 2014 and became the first Application Security Engineer at Pinterest. Lundberg works to protect 'pinners' by securing the company's web and mobile applications in addition to Pinterest's many internal tools.

SAVE THE DATE

The Center for Networked Systems (CNS) will hold its Spring 2015 Research Review **April 8-9, 2015**, in the Qualcomm Conference Center of Jacobs Hall on the UC San Diego campus. The review will focus on the data center, featuring exciting new data center-related work being carried out in CNS and at its industry members. Attendance is by registration only. Contact cns@ucsd.edu with questions or to request an invitation.

CNS WELCOMES NEW MEMBER

The Center for Networked Systems (CNS) welcomes its newest member, CSR.

CSR solves the challenges and delivers the core innovations that enable their customers to win in the global consumer electronics market. Their technologists create innovative and integrated platforms, helping their customers turn great ideas into market-leading products.

CSR is a highly successful, UK-headquartered fabless semiconductor company employing over 2,000 people across 10 countries, offering a range of "technology platform solutions" and System-on-a-Chip devices for a multitude of consumer electronics applications, from Indoor Location to DSLR cameras and Bluetooth wireless audio. Founded in 1998, the company has ongoing collaborations with world-class centers of excellence established at particular universities. CSR also supports and funds exceptional PhD candidates pursuing leading-edge innovation in fields related to CSR's research interests including Wi-Fi, Bluetooth, audio, NFC and power management technologies.

In October 2014 Qualcomm agreed to acquire the UK-based CSR for an estimated $2.5 billion. The acquisition is expected to be completed in late summer 2015.

WHO WATCHES THE WATCHMEN? (CONTINUED FROM PAGE 3)

The current study by CNS researchers achieves a number of firsts by addressing these concerns. "Ours was the first analysis of an AIT that is independent of the device's manufacturer and its customers," said CSE professor Hovav Shacham, lead researcher on the paper. "It was also the first to assume an adaptive adversary, and the first to consider software as well as hardware in their security analysis."

Shacham added that, regarding a malicious actor approaching the security apparatus, "the researchers experimented with a number of possible strategies to breach the security measures that safeguard the privacy, safety, and functionality of the scanner."

As a result, the researchers not only empirically tested a number of breaches previously proposed by other security researchers, but they also devised and tested privacy attacks heretofore not considered possible in connection with the use of X-ray backscatter technology. In the study, Shacham and his team first demonstrated how they could exploit their knowledge of the scanning technology and of security protocols to hide objects such as a handgun simply through clever positioning of the weapon in their clothing or taped to their bodies. For example, moving a weapon to the side of the body means that it can blend into the image's background – no longer standing out against the paleness of flesh. The researchers also showed how plastic explosives can be shaped and molded in such a way that they no longer appear in the scan.

The researchers also considered a number of cyber-physical attacks. The machine that the researchers studied was protected by a simple and easily picked lock that provided only a 10-second delay to gaining physical access to the machine's interface. The software on the tested machine lacked any electronic access controls or user verification, meaning that once access to the controls was achieved, any malicious operator or personnel could manually install malware to compromise the functioning of the Rapiscan device.

The researchers also conceived of a novel privacy hack. They used an external detection device to pick up the scattered X-rays from the scan and then recreated the scanned image in the attacker's device. Though this kind of attack on the security of the scanner does not enable individuals to smuggle contraband through checkpoints, it does point up unforeseen privacy breaches (for example, as a novel way to obtain a nude celebrity photo) that the use of this technology provides.

Shacham's team proposed simple fixes to the system to strengthen the effectiveness of the devices. First of all, checkpoint screeners could conduct side view scans or could pair backscatter scanners with magnetometers to increase the likelihood of detecting hidden objects. However, other complications ensue when making these fixes. Adding extra scans into the screening process would increase airport delays to a point that is logistically unacceptable, while increasing radiation exposures in a way that might pose an unsupportable health risk.

CSE Ph.D. student Keaton Mowery briefs CNS members on the vulnerabilities of X-ray backscatter scanning technology

Perhaps the most important security measure that the researchers verified as key to the use of the Rapiscan was in maintaining the physical unavailability of the machine itself to anyone who might want to smuggle contraband through it. In their study, the researchers showed that even a well-reasoned and well-planned circumvention of the Rapiscan required extensive trial-and-error with an actual machine. For this reason, the researchers recommend that the availability of these machines be highly regulated so that bad actors (such as a well-funded terrorist organization) cannot gain access to them. Simply banning the manufacturers from selling to other than approved buyers is not sufficient, since criminal elements can currently acquire a Rapiscan device using the same method that the CNS researchers used to obtain one: by purchasing it on eBay.

The study urged manufacturers of gate-keeping devices such as the Rapiscan to pursue two design objectives: first, to build the idea of an adaptive threat into the design and testing of their devices; and second, to agree on a thorough and independent review of their security claims. According to Shacham, the latter is the most important conclusion for manufacturers as well as policy makers who recommend or approve the purchasing of such devices – that all security claims be independently reviewed prior to allowing the purchase or deployment of these technologies on a large scale.