The Saucer

The Newsletter of the
CyLab Usable Privacy and Security (CUPS) Laboratory

http://cups.cs.cmu.edu

Issue 1, Summer 2010

Contents

CUPS research is sponsored in part by the following:

- Alcatel-Lucent
- Army Research Office (ARO)
- Fundacao para a Ciencia e Tecnologia (FCT)
- Google
- IBM
- Microsoft
- National Science Foundation (NSF)

Carnegie Mellon
CyLab
CONFIDENCE FOR A NETWORKED WORLD

New CUPS Doctoral Training Program

Thanks to an NSF Integrative Graduate Education and Research Traineeship (IGERT) grant, CUPS has started a new usable privacy and security doctoral training program. Students in PhD programs from throughout Carnegie Mellon University have the opportunity to participate in the CUPS doctoral training program and earn a CyLab Usable Privacy and Security Meritorious Achievement Certificate through the CMU Information Networking Institute (INI). Students who are US citizens can also apply for an IGERT traineeship, which provides two years of tuition and stipend support for students in the doctoral training program.

The CUPS doctoral training program offers students a cross-disciplinary training experience that prepares them to produce the key research advances necessary to reconcile ostensible tensions between security, privacy and usability, moving away from an "either-or" view of these goals to a deeper understanding of underlying tradeoffs and eventually towards solutions where security, privacy and usability are configured to reinforce each other. The goal of this program is to serve as a catalyst to shape the field of usable privacy and security by developing and training a new generation of researchers in methodologies, principles, and approaches that can be applied across systems and applications, in contrast to one-off solutions.

This program leverages CMU's strong research programs in security, privacy, human computer interaction (HCI), behavioral economics, computer systems, artificial intelligence, and decision making, as well as a long tradition and strong commitment to interdisciplinary research.

First class of IGERT students.

[Continued on page 12]

From the Director

Welcome to the inaugural issue of *The Saucer*, the CyLab Usable Privacy and Security Laboratory's annual newsletter. We've come a long way since I founded the CUPS Lab with just a couple of students in 2004. I am grateful for those early CUPS students who encouraged me to turn our informal lab meetings into regular events and who contributed to making those meetings worthwhile and interesting. Now six years later, our cube space in the CIC building is crowded with PhD students, post-docs, and summer visitors, and our Wednesday CUPS lunches are regularly attended by 15-20 people from six or more departments. We've finally gotten around to starting a newsletter to share news about our activities with our friends. And this year we have lots of good news to share….

We were awarded an NSF IGERT grant just before the start of the fall semester. IGERT funding provides student fellowships as well as funds for curriculum development and new equipment (we'll be shopping for an eye tracker soon!). Thanks to the many CUPS faculty who assisted in writing the proposals required to secure this funding.

Our privacy research has touched on some timely public policy issues. Alessandro Acquisti and I were invited to speak about privacy at the FTC privacy roundtable series last December. In February I was invited to discuss our paper on location sharing technologies at a US House subcommittee hearing on the collection and use of location information. This was my first time testifying before a Congressional committee, and it was quite an interesting experience. I enjoyed hearing Representative Boucher tell the wireless industry representative that "It might be helpful if you review Professor Cranor's comments" about their privacy guidelines. I struggled a bit when Representative Stearns asked me, "If you were me and you were doing a privacy bill, what would you suggest as being part of location-based privacy?" He did not accept my first response that fortunately it was his job, not mine. The last question was from our own Representative Doyle, who acknowledged CMU as one of the great universities in America. At a previous hearing last Fall, Representative Doyle specifically mentioned CyLab and Alessandro Acquisti's privacy research.

In April I returned to Capitol Hill to talk about location privacy at the Congressional Internet Caucus' State of the Mobile Net event. I was back again in May to speak on a behavioral advertising and privacy panel and to meet with Congressional staff working on privacy legislation. I've also talked about our privacy nutrition label on a number of these occasions and have heard a lot of enthusiasm for this approach to privacy notices.

In February Alessandro Acquisti, Norman Sadeh, and I received a Google Focused Research Award, providing a substantial gift to fund a new project on "privacy nudges." In June we received a large NSF grant for privacy nudges research. Our research focuses on designing and testing systems that anticipate

CyLab
Usable
Privacy *and*
Security
Laboratory

http://cups.cs.cmu.edu

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Director
Lorrie Faith Cranor
lorrie@cmu.edu

Core Faculty

Alessandro Acquisti
Lujo Bauer
Nicolas Christin
Julie Downs
Jason Hong
Norman Sadeh
Marios Savvides

Supporting Faculty

Wandi Bruine de Bruin
David Brumley
Kathleen Carley
Laura Dabbish
Anupam Datta
Baruch Fischhoff
Greg Ganger
Virgil Gligor
Jim Herbsleb
Bonnie John
Robert Kraut
Ramayya Krishnan
George Lowenstein
Brad Myers
Adrian Perrig
Michael Shamos
Latanya Sweeney
Rahul Telang

and counter cognitive and behavioral biases that hamper users' privacy and security decision making. Inspired by behavioral economics research on "soft" paternalism and research in behavioral decision making and usability, we are designing and studying systems that "nudge" users towards certain behaviors that the users themselves have stated to prefer, or which empirical evidence has demonstrated to be preferable from a privacy perspective.

As you can tell from the paper abstracts included in this issue, CUPS Lab members have been busy with lots of research projects this year. These are all great papers, and if the abstracts pique your interest, I encourage you to get the full papers from our website. I am especially pleased with the growing number of researchers from different departments who have been involved in our research. This may be why our papers seem to have an increasing number of co-authors–including a CHI paper with 16 co-authors!

One paper I'm particularly excited about is "Encountering Stronger Password Requirements: User Attitudes and Behaviors," which will be presented at SOUPS 2010. This project came about from a series of CUPS lunch discussions after the Andrew system password requirements changed. From discussions with the CMU Information Security Office (ISO) we learned that organizations are increasingly under pressure to increase password strength, but that there is little empirical data on what password requirements actually result in stronger passwords. This paper estimates password strength based on survey results. In our ongoing research we expect to gather password strength data through online studies as well as from actual password systems. This project has been a large collaborative effort to tackle a real-world problem that went from nothing to a submitted paper in two months, despite having no dedicated funding. This sort of collaborative interdisciplinary research with real-world impact is what the CUPS Lab is all about. This project will continue on with support from Microsoft Research.

I ended the 2008-2009 academic year by attending the graduation of my first three PhD students–Rob Reeder, Serge Egelman, and Ponnurangam Komaraguru. In 2010 I graduated two more PhD students–Janice Tsai and Steve Sheng. I am thrilled at the wide variety of career paths these students have taken. Rob is developing trustworthy user experiences at Microsoft. Serge is finishing up a post-doc at Brown and will soon start as a usable security researcher at NIST. PK is doing cyber security research as an assistant professor in India. Janice is writing smart grid privacy legislation for the California legislature. And Steve is doing technology policy work for ICANN.

Finally, as we're gearing up for our 6th Symposium On Usable Privacy and Security (SOUPS) conference, I would like to thank Google and Microsoft, our corporate partners who have hosted SOUPS 2009 and 2010. Thanks to their support we were able to take SOUPS out of Pittsburgh for two years, making the conference more accessible to the developers and corporate decision makers who are working to make usable security a reality in their products. It has been great to see increasing interest in usable privacy and security from academic researchers, industry, and government.

Lorrie

Students

Idris Adjerid
Hazim Saleh Almuhimedi
Cristian Bravo-Lillo
Justin Cranshaw
Nipun Gupta
Pedro Giovanni Leon
Hanan Hibshi
Patrick Kelley
Peter Klemperer
Saranga Komanduri
Michelle Mazurek
Aleecia McDonald
Bryan Pendleton
Sasha Romanosky
Rich Shay
Kami Vaniea
Kai Wang

Post-Docs, Visiting Researchers & Staff

Anna Maria Berta
Naoko Hayashida
Mandy Holbrook
Janne Lindqvist
Jennifer Lucas
Marty McGuire
Eran Toch
Yang Wang

Visiting Summer Students

Robin Brewer,
University of Maryland
Yael Mayer,
Harvey Mudd College
Greg Norcie,
University of Pittsburgh

Shane, Lorrie, Maya, and Nina Cranor playing in the snow. February 2010 was Pittsburgh's snowiest February on record.

Lorrie Cranor speaking on Capitol Hill in June 2010.

Alessandro Acquisti speaks on NPR in October 2009.

Norman Sadeh presenting Locaccino.

Year in Review

July 2009

- CUPS director Lorrie Cranor presents at a National Academy of Sciences workshop on Usability, Security, and Privacy of Information Systems. A workshop report will be published in Summer 2010.

August 2009

- CUPS director Lorrie Cranor and her Carnegie Mellon University colleagues receive a five-year, $3 million IGERT grant from the National Science Foundation (NSF) to establish a Ph.D. program in usable privacy and security.
- Josh Sunshine presents the paper "Crying Wolf: An Empirical Study of SSL Warning Effectiveness" at USENIX Security 2009. Press coverage included *ABC News*, *The Tech Herald*, *Computerworld*, and *SC Magazine*.

September 2009

- Janice Tsai presents a paper on location-sharing technologies at the Telecommunications Policy Research Conference (TPRC).

October 2009

- On October 26, CUPS faculty Alessandro Acquisti is featured on the *NPR* program *All Things Considered* speaking about his research on privacy through the lens of behavioral economics.

November 2009

- Two new technical reports are issued and submitted as public comments to the Federal Trade Commission's exploring privacy roundtable series: "Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach" and "An Empirical Study of How People Perceive Online Behavioral Advertising."

December 2009

- CUPS director Lorrie Cranor and CUPS faculty Alessandro Acquisti speak at the Federal Trade Commission's series of day-long privacy roundtable discussions. The topic is about how to explore the privacy challenges posed by the vast array of 21st Century technology and business practices that collect and use user data.
- A Google Research Award is given to CUPS faculty Jason Hong and computer science faculty John Zimmerman for their project: Context-Aware Mobile Mash-ups. The project seeks to build tools for non-programmers to create location and context-aware mashups of data for mobile devices that can present time- and place-appropriate information.
- On December 19, CUPS faculty Jason Hong is featured in a *New York Times*'s article describing how Goggles uses location information to help identify objects, and how its ability to recognize millions of images opens up new possibilities.
- CUPS director Lorrie Cranor gives a keynote talk at the Annual Computer Security Applications Conference in Honolulu, Hawaii and receives the ACSAC Distinguished Practitioner Award.

Jason Hong receives the Alfred P. Sloan Foundation fellowship.

ACM Student Research Competition. Left to right: ACM President Wendy Hall, Second Place Winner Michael Tvarozek, First Place Winner Patrick Kelley, Judith Bishop of Microsoft Research.

Greg Ganger testifies at a July 1 hearing of the House Committee on Oversight and Government Reform and the Government Management, Organization, and Procurement Subcommittee.

- CUPS director Lorrie Cranor is featured in a Silver Bullet Security Podcast interview by Gary McGraw.

January 2010

- CUPS director Lorrie Cranor gives a keynote talk at Financial Cryptography and Data Security 2010 in Tenerife, Canary Islands.

February 2010

- CUPS director Lorrie Cranor discusses the paper "Location-Sharing Technologies: Privacy Risks and Controls" in her testimony at a Congressional hearing on February 24.
- The *Pittsburgh Post Gazette* quotes CUPS director Lorrie Cranor in an article warning about electronic Valentines Day cards.
- CUPS director Lorrie Cranor and CUPS faculty Norman Sadeh and Alessandro Acquisti receive a Google Focused Research Award to support research on "privacy nudges."
- CUPS faculty Jason Hong receives the Alfred P. Sloan Foundation fellowship.
- CUPS director Lorrie Cranor is featured in a *New York Times* article on "Redrawing the Route to Online Privacy."

March 2010

- Japan's Anti-Phishing Council and JPCERT/CC release a customized version of Wombat's Anti-Phishing Phil Training Game to educate the Japanese public. Anti-phishing Phil was developed in the CUPS Lab.
- On March 17, CUPS faculty Alessandro Acquisti is featured in a *New York Times* article on how privacy vanishes online.
- CUPS director Lorrie Cranor appears in a live BBC Newshour interview with Claire Bolderson about privacy and location-sharing applications.

April 2010

- Three CUPS papers, 1 note, 1 workshop paper, and 2 student posters are presented at CHI2010 in Atlanta, Georgia.

May 2010

- CUPS director Lorrie Cranor and CUPS faculty Jason Hong are featured in a *Pittsburgh Business Times* article about cyber security best practices.

June 2010

- CUPS PhD student Patrick Kelley wins first place in the grand finals of the ACM Student Research Competition. He receives his award at the ACM awards banquet on June 26.

July 2010

- CUPS faculty Greg Ganger discusses the government-wide transition to cloud computing in his testimony at a Congressional hearing on July 1.
- SOUPS 2010 is held July 14-16 at Microsoft, in Redmond, Washington.
- "Engineering Privacy," a paper coauthored by CUPS director Lorrie Cranor and Sarah Spiekermann, is recognized as a runner up for the 2010 Award for Outstanding Research in Privacy Enhancing Technologies. This paper was published in the January/February 2009 issue of *IEEE Transactions on Software Engineering*.

Building a Better Password

Tough to remember but easy to crack, passwords are the weak link in computer security. Billions hang in the balance.

By Nick Summers

Excerpted from a Newsweek article, published October 9, 2009

Among computer researchers, passwords are a key aspect of a burgeoning field known as "usable security." At Carnegie Mellon, the scientists who've pioneered the discipline work not in a lab but upstairs in a wing that looks no different from most universities' English or history departments. Look closer, though, and you'll see signs that this is no ordinary place. The doors are all marked with 2-D bar codes; a professor enters his office by snapping a photo with his cell phone. Click! goes the phone; thunk! slides the bolt. It's more secure than a physical key, which can be stolen and copied, and no less handy.

The academics here are rethinking basic questions about what makes something—an office, a Web site—secure, without driving its owner crazy. And their findings call into question many of the recent security advances in the banking, e-mail, and other critical systems you log into every day. Researchers here fault virtually everything your corporate IT department tells you about strong passwords. And they take the radical stance that you, the user, should be listened to when passwords become overbearing, not yelled at when you forget them.

As an academic discipline, usable security—a blend of computer science and psychology—is only about five years old. "When we first started waving the flag, not many people paid attention," says Carnegie Mellon professor Lorrie Cranor. "It's gratifying that people are starting to." Cranor may be more responsible than anyone else for establishing the field.

She founded CyLab's Usable Privacy and Security Laboratory and an annual symposium; she also edited the major textbook on the subject and teaches one of the few usable-security-specific courses in the nation. Polite and warm, Cranor strives to be user-friendly herself: when she gets too technical while describing her work to a decidedly non-Ph.D. *Newsweek* reporter, she pauses, laughs ("Were you expecting a more usable definition?"), and resumes the discussion in geek-free English. Her interest in patterns and complexity extends outside the lab: she's a master quilter whose designs have been featured on the covers of textbooks and journals.

. . . .

One way humans deal with password overload is to rely on a single password and simple variants for nearly every electronic interface in their lives—as I did. That's highly problematic because if that all-powerful password is cracked at just one site, it gives a hacker the keys to the kingdom. That's why Adrian Perrig, the technical director at Carnegie Mellon's CyLab, promotes disposable passwords: generated by special devices, people use these passwords once, then throw them away.

. . . .

CUPS students, faculty and staff pose for a photo with their private birds in September 2009.

Recent Publications

School of Phish: A Real-Word Evaluation of Anti-Phishing Training

Ponnurangam Kumaraguru, Justin Cranshaw, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, Mary Ann Blair, and Theodore Pham

SOUPS 2009

PhishGuru is an embedded training system that teaches users to avoid falling for phishing attacks by delivering a training message when the user clicks on the URL in a simulated phishing email. In previous lab and real-world experiments, we validated the effectiveness of this approach. Here, we extend our previous work with a 515-participant, real-world study in which we focus on long-term retention and the effect of two training messages. We also investigate demographic factors that influence training and general phishing susceptibility. Results of this study show that (1) users trained with PhishGuru retain knowledge even after 28 days; (2) adding a second training message to reinforce the original training decreases the likelihood of people giving information to phishing websites; and (3) training does not decrease users' willingness to click on links in legitimate messages. We found no significant difference between males and females in the tendency to fall for phishing emails both before and after the training. We found that participants in the 18-25 age group were consistently more vulnerable to phishing attacks on all days of the study than older participants. Finally, our exit survey results indicate that most participants enjoyed receiving training during their normal use of email.

A Comparative Study of Online Privacy Policies and Formats

Aleecia M.McDonald, Robert W. Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor

Privacy Enhancing Technologies Symposium 2009

Online privacy policies are difficult to understand. Most privacy policies require a college reading level and an ability to decode legalistic, confusing, or jargon-laden phrases. Privacy researchers and industry groups have devised several standardized privacy policy formats to address these issues and help people compare policies. We evaluated three formats in this paper: layered policies, which present a short form with standardized components in addition to a full policy; the Privacy Finder privacy report, which standardizes the text descriptions of privacy practices in a brief bulleted format; and conventional non-standardized human-readable policies. We contrasted six companies' policies, deliberately selected to span the range from unusually readable to challenging. Based on the results of our online study of 749 Internet users, we found participants were not able to reliably understand companies' privacy practices with any of the formats. Compared to natural language, participants were faster with standardized formats but at the expense of accuracy for layered policies. Privacy Finder formats supported accuracy more than natural language for harder questions. Improved readability scores did not translate to improved performance. All formats and policies were similarly disliked. We discuss our findings as well as public policy implications.

Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application

Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker,and Jinghai Rao.

Personal and Ubiquitous Computing, August 2009

A number of mobile applications have emerged that allow users to locate one another.
However, people have expressed concerns about the privacy implications associated with this class of software, suggesting that broad adoption may only happen to the extent that these concerns are adequately addressed. In this article, we report on our work on PeopleFinder, an application that enables cell phone and laptop users to selectively share their locations with others (e.g. friends, family, and colleagues). The objective of our work has been to better understand people's attitudes and behaviors towards privacy as they interact with such an application, and to explore technologies that empower users to more effectively and efficiently specify their privacy preferences (or "policies"). These technologies include user interfaces for specifying rules and auditing disclosures, as well as machine learning techniques to refine user policies based on their feedback. We present evaluations of these technologies in the context of one laboratory study and three field studies.

Crying Wolf: An Empirical Study of SSL Warning Effectiveness

Joshua Sunshine, Serge Egelman, Hazim Almuhimedi, Neha Atri, and Lorrie Faith Cranor

USENIX Security 2009

Web users are shown an invalid certificate warning when their browser cannot validate the identity of the websites they are visiting. While these warnings often appear in benign situations, they can also signal a man-in-the-middle attack. We conducted a survey of over 400 Internet users to examine their reactions to and understanding of current SSL warnings. We then designed two new warnings using warnings science principles and lessons learned from the survey. We evaluated warnings used in three popular web browsers and our two warnings in a 100- participant, between-subjects laboratory study. Our warnings performed significantly better than existing warnings, but far too many participants exhibited dangerous behavior in all warning conditions. Our results suggest that, while warnings can be improved, a better approach may be to minimize the use of SSL warnings altogether by blocking users from making unsafe connections and eliminating warnings in benign situations.

Access Control for Home Data Sharing: Attitudes, Needs and Practices

Michelle L. Mazurek, J.P. Arsenault, Joanna Bresee, Nitin Gupta, Iulia Iony, Christina Johns, Daniel Lee, Yuan Liang, Jenny Olsen, Brandon Salmon, Richard Shay, Kami Vaniea, Lujo Bauer, Lorrie Faith Cranor, Gregory R. Ganger, and Michael K. Reiter

CHI 2010

As digital content becomes more prevalent in the home, nontechnical users are increasingly interested in sharing that content with others and accessing it from multiple devices. Not much is known about how these users think about controlling access to this data. To better understand this, we conducted semi-structured, in-situ interviews with 33 users in 15 households. We found that users create ad-hoc access control mechanisms that do not always work; that their ideal policies are complex and multi-dimensional; that a priori policy specification is often insufficient; and that people's mental models of access control and security are often misaligned with current systems. We detail these findings and present a set of associated guidelines for designing usable access-control systems for the home environment.

Are Your Participants Gaming the System? Screening Mechanical Turk Workers

Julie S. Downs, Mandy B. Holbrook, Steve Sheng, and Lorrie Faith Cranor

CHI 2010

In this paper we discuss a screening process used in conjunction with a survey administered via Amazon.com's Mechanical Turk. We sought an easily implementable method to disqualify those people who participate but don't take the study tasks seriously. By using two previously pilot tested screening questions, we identified 764 of 1,962 people who did not answer conscientiously. Young men seem to be most likely to fail the qualification task. Those that are professionals, students, and non-workers seem to be more likely to take the task seriously than financial workers, hourly workers, and other workers. Men over 30 and women were more likely to answer seriously.

Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions

Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs

CHI 2010

In this paper we present the results of a roleplay survey instrument administered to 1001 online survey respondents to study both the relationship between demographics and phishing susceptibility and the effectiveness of several antiphishing educational materials. Our results suggest that women are more susceptible than men to phishing and participants between the ages of 18 and 25 are more susceptible to phishing than other age groups. We explain these demographic factors through a mediation analysis. Educational materials reduced users' tendency to enter information into phishing webpages by 40% percent; however, some of the educational materials we tested also slightly decreased participants' tendency to click on legitimate links.

Teaching Johnny Not to Fall for Phish

Ponnurangam K Kumaraguru, Steve Sheng, Alessandro Acquisti, Lorrie Faith Cranor, and Jason Hong

ACM Transactions on Internet Technology, May 2010

Phishing attacks, in which criminals lure Internet users to Web sites that spoof legitimate Web sites, are occurring with increasing frequency and are causing considerable harm to victims. While a great deal of effort has been devoted to solving the phishing problem by prevention and detection of phishing emails and phishing Web sites, little research has been done in the area of training users to recognize those attacks. Our research focuses on educating users about phishing and helping them make better trust decisions. We identified a number of challenges for end-user security education in general and anti-phishing education in particular: users are not motivated to learn about security; for most users, security is a secondary task; it is difficult to teach people to identify security threats without also increasing their tendency to misjudge non threats as threats.

Keeping these challenges in mind, we developed an email-based anti-phishing education system called "PhishGuru" and an online game called "Anti-Phishing Phil" that teaches users how to use cues in URLs to avoid falling for phishing attacks. We applied learning science instructional principles in the design of PhishGuru and Anti-Phishing Phil. In this article we present the results of PhishGuru and Anti-Phishing Phil user studies that demonstrate the effectiveness of these tools. Our results suggest that, while automated detection systems should be used as the first line of defense against phishing attacks, user education offers a complementary approach to help people better recognize fraudulent emails and websites.

Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach

Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor

CHI 2010

Earlier work has shown that consumers cannot effectively find information in privacy policies and that they do not enjoy using them. In our previous research on nutrition labeling and other similar consumer information design processes we developed a standardized table format for privacy policies. We compared this standardized format, and two short variants (one tabular, one text) with the current status quo: full text natural language policies and layered policies. We conducted an online user study of 789 participants to test if these three more intentionally designed, standardized privacy policy formats, assisted by consumer education, can benefit consumers. Our results show that providing standardized privacy policy presentations can have significant positive effects on accuracy of information finding, overall speed, and reader enjoyment with privacy policies.

RT @IWantPrivacy: Widespread Violation of Privacy Settings in the Twitter Social Network

Brendan Meeder, Jennifer Tam, Patrick Gage Kelley, and Lorrie Faith Cranor

Web 2.0 Security and Privacy 2010

Twitter is a social network that focuses on creating and sharing short 140 character messages know as *tweets*. Twitter's sole privacy policy is a binary option that either allows every message a user creates to be publicly available, or allows only a user's *followers* to see posted messages. As the Twitter community grew, conventions organically formed that allow for rich expressiveness with only 140 characters. Repeating what someone else says is called *retweeting*; this behavior facilitates the spread of information and commentary in real-time. We have performed a large-scale collection of data from the Twitter social network by means of the publicly available application programming interface they provide. Our data set contains over 2.7 billion messages, 80 million user profiles, and a 2.6 billion edge social network. We analyze these data and uncover the growing trend where users defeat Twitter's simple privacy mechanism of "protecting one's tweets" by simply retweeting a protected tweet. We have shown through quantitative and qualitative analysis that these privacy-violating retweets are a growing problem. More than 4.42 million tweets exist in our corpus that expose protected information. As Twitter gains popularity over time, we see an increasing trend in the number of privacy-violating retweets.

Institutional Review Boards and Your Research

Simson L. Garfinkel and Lorrie Faith Cranor

Communications of the ACM, June 2010

A proposal for improving the review procedures for research projects that involve human subjects and their associated identifiable private information.

Encountering Stronger Password Requirements: User Attitudes and Behaviors

Richard Shay, Saranga Komanduri, Patrick Gage Kelley, Pedro Giovanni Leon, Michelle L. Mazurek, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor

SOUPS 2010

Text-based passwords are still the most commonly used authentication mechanism in information systems. We took advantage of a unique opportunity presented by a significant change in the Carnegie Mellon University (CMU) computing services password policy that required users to change their passwords. Through our survey of 470 CMU computer users, we collected data about behaviors and practices related to the use and creation of passwords. We also captured users' opinions about the new, stronger policy requirements. Our analysis shows that, although most of the users were annoyed by the need to create a complex password, they believe that they are now more secure. Furthermore, we perform an entropy analysis and discuss how our findings relate to NIST recommendations for creating a password policy. We also examine how users answer specific questions related to their passwords. Our results can be helpful in designing better password policies that consider not only technical aspects of specific policy rules, but also users' behavior in response to those rules.

CUPS Research Showcased in Online Demos, Spinoff Companies

CUPS research projects have resulted in a number of tools that put our research findings into practice. The CUPS website currently features demos of three long-running projects, two of which are being developed further by start-up companies.

Anti-phishing Phil

Anti-Phishing Phil is an interactive game that teaches users how to identify phishing URLs, where to look for cues in web browsers, and how to use search engines to find legitimate sites. CUPS user studies have found that user education can help prevent people from falling for phishing attacks. However, it is hard to get users to read security tutorials, and many of the available online training materials make users aware of the phishing threat but do not provide them with enough information to protect themselves. Our lab and field studies demonstrate that Anti-Phishing Phil is an effective approach to end-user security education. After receiving inquiries from companies interested in licensing the game for their corporate security training, CUPS faculty Norman Sadeh, Jason Hong, and Lorrie Cranor started a company, Wombat Security Technologies, Inc. to commercialize Anti-Phishing Phil and other anti-phishing technologies. Wombat has now licensed Phil to large and small organizations around the world, and recently developed a new game called Anti-Phishing Phyllis.

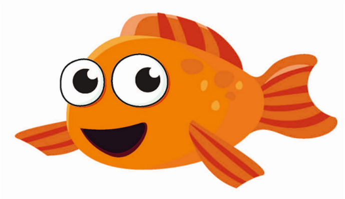

Visit http://wombatsecurity.com for demos.

Locaccino

As part of our user-controllable security and privacy research effort, we developed a location-sharing service called Locaccino to provide a platform for experimenting with privacy controls for end users. Locaccino provides users with precision control over who can see their location. While most location-based systems only allow users to list which contacts should or shouldn't be able to see their location, Locaccino offers more granular privacy settings. Locaccino allows users to reveal or hide their location based on time of day or location. For example, a user might allow her co-workers to locate her during business hours and her family to locate her 24x7. In addition, she might let her friends from the gym locate her only when she is at the gym. We are experimenting with approaches to simplify the configuration of privacy settings and offer users automated assistance in refining their settings over time. Locaccino requires a Facebook account to use and currently runs on wifi-enabled laptops, Android phones, and Symbian phones.

A new CMU start-up, Zipano Technologies, is commercializing the technology behind Locaccino.

Visit http://locaccino.org to try it out.

Privacy Finder

Reading and comparing the privacy policies of websites is difficult and time consuming. We developed a search engine called Privacy Finder that annotates search results with privacy meters to help users determine which sites have the best privacy policies. Users can click on the privacy meters to drill down and view a privacy "nutrition label" that summarizes each site's privacy practices. Privacy Finder reads web sites' machine-readable Platform for Privacy Preferences (P3P) privacy policies to generate the privacy meters and nutrition labels. Unfortunately, only a small fraction of websites have P3P policies, so you won't see this information for every site. We've developed the privacy nutrition label through an iterative design approach, which you can read about in two recent papers. The Privacy Finder search engine has also served as a platform for a number of our published studies on privacy-decision making behavior.

Visit http://privacyfinder.org to try it out.

Dissertations & Proposals

Thesis Proposal:
Using Proximity Information Displays and Audit Log Information to Motivate Users to View and Maintain Access-control Policy

Kami Vaniea

Thesis Committee:
Lorrie Cranor, Chair
Jeannette Wing
Lujo Bauer
Mike Reiter, University of North Carolina, Chapel Hill
Stuart Schechter, Microsoft Research

Carnegie Mellon University
Computer Science Department
June 1, 2010

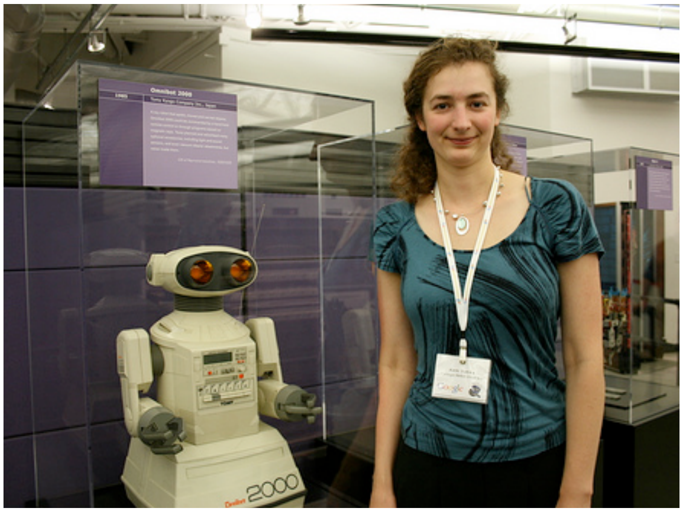
Kami Vaniea at the SOUPS 2009 reception in the Computer History Museum.

Managing access to shared information, such as photographs and documents, is an emerging difficult issue for end users who are accumulating an increasingly large and diverse collection of data that they want to share with others. Current policy management solutions require a user to proactively seek out and open a separate policy management interface when they want to review or change their access-control policy. This may be acceptable to users if they are engaged in a primary tasks such as ``Give Bob access to vacation pictures," but, for the majority of users, such tasks are uncommon. Consequently, access-control policies are rarely reviewed or even glanced at. Historically, security administrators and auditors helped fill this gap by actively checking for issues on behalf of users, but in the age of Facebook and Flickr users have no professionals double-checking their work. Users need a way to review their access-control policies that fits into their normal workflow.

To enable users to better understand the implications of their access-control policy as well as how it is used we need to provide greater transparency to end users. In this thesis I am proposing the use of proximity information displays to make users more aware of how their resources have been used in the past and how they could be used in the future. Proximity information displays are interface components that have the following properties:

- spatial proximity to resources,
- glanceable,
- show useful and interesting information,
- allow layered data exploration, and
- make it easy to segue to policy modification.

They are referred to as proximity information displays because information is always placed in close proximity to where resources are displayed on the interface. As the user navigates through their files, or other resources, the proximity information displays will update so that displayed information is always only about the displayed resources. How to best implement each of the above properties in proximity information displays in a way that best supports awareness of data usage and policy implications will be the focus of this thesis.

Janice Y. Tsai. *The Impact of Salient Privacy Information on Decision-Making*, PhD Thesis, Engineering & Public Policy Department, Carnegie Mellon University, Pittsburgh, PA, August 2009.

Steve Sheng. *A Policy Analysis of Phishing Countermeasures*, PhD Thesis, Engineering & Public Policy Department, Carnegie Mellon University, Pittsburgh, PA, August 2009.

New CUPS Doctoral Training Program

[Continued from page 1]

Students in the CUPS doctoral training program are required to take our graduate-level Usable Privacy and Security course, plus three approved full-semester courses from the CUPS course list. The CUPS course list includes courses in security, privacy, human computer interaction, social and decision sciences, and other areas. Students are also expected to participate in the weekly CUPS research seminar for at least two years and present their work at this seminar at least once per year. We also expect CUPS students to be engaged in usable privacy and security research.

CUPS faculty from a variety of disciplines participate in the CUPS research seminar and mentor students in the program. We expect most students will be mentored by at least two CUPS faculty members from complementary fields.

In August 2009 we admitted our first class of six PhD students from five different PhD programs in three schools across the university—I Computation, Organizations, and Society; Computer Science Department; Human Computer Interaction Institute; Electrical and Computer Engineering; and Public Policy and Management.

Students should apply directly to a CMU PhD program of their choice, and also send a letter of interest to the CUPS program administrator indicating which CMU doctoral program they have applied to and describing their interest in CUPS-related research. There is additional information on the CUPS website at http://cups.cs.cmu.edu/igert/ — including information for current CMU students interested in participating in the program.

CUPS summer student Greg Norcie collaborates with CUPS PhD students Saranga Kumanduri and Rich Shay on password research.

CUPS summer students Robin Brewer and Yael Mayer are conducting research on how social network users divide their friends into groups.

CUPS post-doc Yang Wang is doing research on privacy and cross-system personalization.