Research Statement

Robert W. Reeder

My primary research interests are at the intersection of human-computer interaction (HCI), computer security, and online privacy, in an area known as *usable security and privacy*. Through scientifically rigorous studies, I seek to understand individual users' and enterprises' needs for security and privacy and to develop user interface designs that allow them to better manage the security and privacy of their data. I believe this work is increasingly important as we share more data – especially personal data – in a connected world.

Research vision

My vision is for individual computer users and enterprises to be able to receive the productivity and entertainment benefits of modern computing without worrying about security and privacy risks.

Computing has become incredibly useful and fun for many people and enterprises since the advent of the Internet. Unfortunately, the usefulness and fun of computing is easily and often ruined by viruses, botnets, undesired data disclosure, and intellectual property theft. Protection mechanisms like firewalls, anti-virus scanners, passwords, and warning dialogs are useless when the humans who use them can't understand them. With my research, I aspire to enhance the usefulness and fun of computing by making the security and privacy risks easily understandable and manageable through simplified protection mechanisms.

Thesis work

Part of making security risks manageable is allowing people to clearly see the security state of their system and change it. In my thesis work, I created an information visualization technique, called Expandable Grids, for creating, viewing, and editing – in a word, *authoring* – security and privacy policies, such as file system access control policies and website privacy policies. Expandable Grids are a solution to problems with the inadequate but dominant paradigm for policy-authoring interfaces, the "list-of-rules" paradigm. List-of-rules interfaces are centered around a list of a policy's rules, which state what users have what access to what data. For

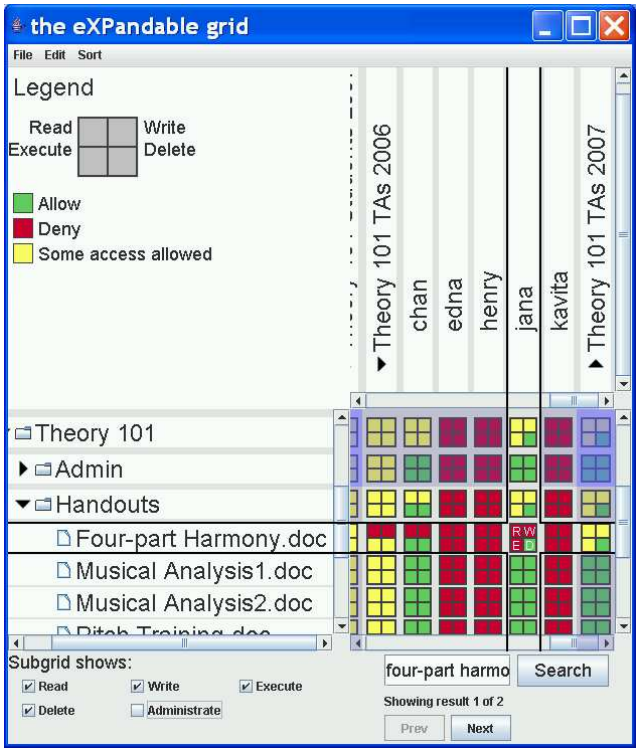

Figure 1. A user interface I built for setting Windows XP file permissions based on the Expandable Grids concept. The interface shows files and folders in the tree on the left, users and groups in the tree on top, and effective policy in the colored squares within the grid.

example, a rule may state that a user "jsmith" can "execute" the "calc.exe" program. Users and data may be grouped, and a rule may apply to a whole group; for example, a rule might state that the group "Employees" cannot "read" the "salaries.xls" file. In list-of-rules interfaces, rules may be selected one-at-a-time from the list for viewing and editing. The list-of-rules model has at least two major drawbacks. First, group membership information is not displayed in context with the rules. Second, rules may interact with each other, and even conflict with each other. List-of-rules interfaces do not give any indication of these rule interactions; instead it is left to the policy author to figure out how rules will interact.

In contrast to list-of-rules interfaces, Expandable Grids show group membership and show *effective policy;* that is, they show precisely what a policy allows or does not allow, rather than merely showing the rules from which effective policy is derived. Expandable Grids consist of a matrix with hierarchical axes that can be expanded or contracted to show more or less policy detail. Because they show effective policy directly, Expandable Grids do not require policy authors to determine subtle interactions amongst rules.

With collaborators, I implemented three user interfaces that use the Expandable Grids idea, including one for setting file permissions in a Windows file system (shown in Figure 1), one for viewing website P3P privacy policies, and one for managing physical access control policies in an office building. I conducted user studies testing the first two of these interfaces. In one study comparing the Expandable Grids interface for setting file permissions with the native Windows file permissions interface, users performed vastly more accurately and faster on a broad range of file-permissions tasks using the Expandable Grids interface. [3]

Current work

Since finishing my PhD in 2008, I have been working at Microsoft within Trustworthy Computing, the division of the company responsible for security engineering and security incident response. I work on the Usable Security team, whose mission is to make sure the security features in our products are usable by people. Our work entails, for example, helping Microsoft's engineers design warning dialogs that make sense and are actionable, authentication methods that are secure without placing an undue burden on users' memories, and configuration interfaces that help people avoid making mistakes.

In my role on the Usable Security team, I have two key challenges. First is gathering the knowledge about how to develop better security and privacy user experiences. Second is spreading this knowledge to the thousands of engineers within Microsoft's product teams so that they know the right things at the right times. To address these two challenges, my job is a combination of conducting research, writing guidance for engineers, teaching classes and giving talks about usable security, and consulting with product teams.

My role at Microsoft has been a great opportunity to both put my expertise in usable security into practice and to broaden the scope of my interests in usable security. Since my team's mission is to gather and disseminate knowledge about how to design for usable security and privacy, I have continued to do research in areas in which we don't yet have all the answers. This has allowed me to both deepen my work in usable access control and broaden my interests into authentication techniques and security warnings.

In usable access control, collaborators and I tested the Expandable Grids idea in the social-networking privacy domain by comparing it with another social-networking privacy-setting user interface [2]. This study and other internal studies have shown that while the Expandable Grid has strengths I showed in my thesis, it is best complemented by other text- and view-based approaches. Also in usable access control, Microsoft collaborators and I explored the idea of making permissions setting more usable by integrating it into a natural file sharing experience – sending links by email [1] – and tested a series of design variants for granting permissions to applications to use data from a social networking profile [7].

In authentication techniques, we developed and evaluated a "social authentication" system as a secondary authentication (i.e., password reset) mechanism. In our social authentication system, a user specifies several trustees in advance, and to authenticate, contacts those trustees in person or over the phone. The trustees then attest to the web service that the user really is who they say they are and is locked out of their account. If enough trustees attest, the user regains access. To test the security of the system, we had study participants come in pairs and "attack" their partners' accounts by trying to deceive their partners' trustees. The system was a big improvement in both reliability (probability of authenticating a legitimate account holder) and security (probability of rejecting an impostor) over the usual "secret questions" approach to secondary authentication, at the price of some time investment (most participants took an hour or longer to get in touch with 3 trustees) [5, 6].

In security warnings, I led an effort to produce guidance for designing security warnings based on the latest research of the usable security community. I packaged it into a concise form for sharing easily with product teams. The guidance is called NEAT, which stands for the four things a good security warning should be: Necessary, Explained, Actionable, and Tested. The guidance is in use now by engineering teams at Microsoft, and I am collaborating with Lorrie Cranor's lab at Carnegie Mellon to evaluate its effectiveness in user studies. In my role as usable security consultant to teams with questions about their specific scenarios, I have dealt with problems that are fascinating in their complexity to which I have no easy answers. Dealing with these problems has alerted me to the need for more research on security warnings with higher-fidelity data than we have today.

Planned future directions

In future work, I plan to apply my expertise to applications that are likely to have high social value and impact. In particular, my expertise is in user needs analysis, data collection, functional interface design, experimental design, and information visualization. I will continue to apply this expertise to the many research problems on the human aspects of computer security and privacy. In the short term, I have three areas in which I plan to continue my work: studying security warnings in context, verifying online identities, and usable access control.

Studying security warnings in context

Computing online has become dangerous enough that we all experience incessant warnings from our software: "Update immediately." "Software from an unidentified publisher may harm your computer. Do you wish to continue?" or my favorite, "Revocation information for the security certificate for this site is not available. Do you want to proceed?" There are many problems with these warnings: they interrupt us from the more important things we're

trying to do, they use incomprehensible language to explain what's going on, they fail to provide information needed to make a good decision, and they occur so frequently in benign situations (i.e., they "cry wolf") that we simply start ignoring them. They are annoying and vaguely disconcerting, but non-actionable.

At the same time, data published by Microsoft (see Microsoft's Security Intelligence Report, vol. 11) show that user interaction is now the dominant means by which computers are infected with malware. That is, attackers are increasingly using social engineering – attacking people, rather than software – to infect machines. It thus appears that helping people make better security decisions – perhaps, by designing better warnings – may be the greatest untapped means for mitigating online security and privacy risks.

Research on warnings so far has shown that warnings are rarely, if ever, effective at helping users prevent malware infections, credential theft, and other undesirable outcomes. So, should software publishers continue to include warnings in their products, given how annoying and ineffective warnings are?

The answer to this question is not yet clear. Warnings research has been conducted mostly in laboratory settings with small numbers of participants, but software from a publisher like Microsoft goes out to hundreds of millions of users in all kinds of usage contexts. It's entirely possible that warnings that are ineffective in contrived laboratory studies are actually quite effective for a certain audience in a certain context. Or, it may be the case that warnings really never work for anybody. I plan to conduct research to get at the answer to this question by studying many users confronting warnings in their actual context of use. Only by gathering data outside the lab can we determine which warnings are effective and with which audiences in which situations. Then we can answer the question of whether to continue designing software with warnings, and if so, what the warnings should look like. I am currently collaborating with Lorrie Cranor's lab at Carnegie Mellon on the beginnings of such real-world warnings research.

Verifying online identities

Phishing, in which an attacker lures a victim to give up login credentials to a spoofed version of a legitimate website, has been around for several years and remains a problem, despite numerous attempted technical solutions. However, phishing is only one manifestation of the larger difficulty of determining who's who and what's what online. Recent attacks on Web users that exploit this difficulty include luring users to install seemingly legitimate software that is in fact malware and luring users to buy fake anti-virus software by telling them their machine has been infected with viruses. The lures in these attacks often use the same security icons and language that legitimate operating system and anti-virus publishers use in their products. Users simply can't tell one shield icon from another.

The solution to these fake-identity attacks, I think, is two-fold: first, to educate users to be alert that not everything online is legitimate; second, to provide a technical mechanism to give users information about the source of all the icons, text, and warnings on their screen at a given time. Pieces of this solution exist, but they have not yet been brought together in a way that prevents these fake-source attacks on a large scale. I plan to assemble the educational materials, build the technical solution, and evaluate the total solution that will help users stay away from these attacks.

Usable access control

Numerous interesting research questions remain in the area of usable access control. The Expandable Grids idea, the focus of my thesis work, is a user interface technique for allowing for accurate authoring of fine-grained, large policies by novice policy authors. It does not necessarily address authoring of small policies (e.g., policies involving a handful of friends, rather than hundreds or thousands of system users), authoring of coarse-grained policies (e.g., policies in which access is either allowed or denied to all resources or all users, rather than being given out resource-by-resource or user-by-user), or authoring by experts (e.g., system administrators). Moreover, we do not even know whether these contexts are important for usability researchers to address. Important research questions include:

1. In what contexts are small and coarse-grained policies useful and sufficient?
2. In these contexts, are simple text-based or list-of-rules interfaces sufficient for supporting accurate policy authoring? Or are more advanced visual interfaces needed even for small and coarse-grained policy authoring?
3. Do Expandable Grids interfaces help policy-authoring experts, or do experts perform best with text-based or scripting-language interfaces? If the latter, can we combine visualization with text-based or scripting-language approaches to improve expert performance?

References

* For a complete list of publications, please see my C.V.

1. Johnson, M.L., Bellovin, S.M., Reeder, R.W., and Schechter, S.E. Laissez-Faire File Sharing. Paper accepted to New Security Paradigms Workshop (NSPW 2009). 2009.

2. Lipford, H.R., Watson, J., Whitney, M., Froiland, K., and Reeder, R.W. Visual vs. Compact: A Comparison of Privacy Policy Interfaces. Conference short paper accepted to *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. 2010.

3. Reeder, R.W., Bauer, L., Cranor, L.F., Reiter, M.K., Bacon, K., How, K., and Strong, H. Expandable Grids for Visualizing and Authoring Computer Security Policies. Conference paper accepted to *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. 2008.

4. Reeder, R.W., Karat, C.-M., Karat, J., and Brodie, C. Usability Challenges in Security and Privacy Policy-Authoring Interfaces. Conference paper presented at *INTERACT 2007*. Published in Springer *Lecture Notes in Computer Science* (LNCS 4663, Part II, pp. 141-155). 2007.

5. Reeder, R.W. and Schechter, S. When the password doesn't work: secondary authentication for websites. To appear in *IEEE Security & Privacy* magazine special issue on usable security, 2011.

6. Schechter, S., Egelman, S., and Reeder, R.W. It's Not What You Know, but Who You Know: A Social Approach to Last-Resort Authentication. Conference paper accepted to *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. 2009.

7. Tam, J., Reeder, R.W., and Schechter, S. I'm Allowing What? Disclosing the authority applications demand of users as a condition of installation. Microsoft Research technical report available at http://research.microsoft.com/apps/pubs/default.aspx?id=131517. 2010.