
Protecting Private Data in Public

Peter Tarasewich

HCI Laboratory
College of Computer &
Information Science
Northeastern University
360 Huntington Ave., 202WVH
Boston, MA 02115 USA
tarase@ccs.neu.edu

Jun Gong

HCI Laboratory
College of Computer &
Information Science
Northeastern University
360 Huntington Ave., 202WVH
Boston, MA 02115 USA
gjoliver@ccs.neu.edu

Richard Conlan

HCI Laboratory
College of Computer &
Information Science
Northeastern University
360 Huntington Ave., 202WVH
Boston, MA 02115 USA
kaige@ccs.neu.edu

Abstract

Current technologies allow users to access information in virtually any public setting. This creates situations where sensitive information, both organizational and personal in nature, can be seen and captured by nearby people and technology. Therefore, methods are necessary to ensure the privacy and security of information displayed in public spaces. The authors have developed Web browser privacy blinders, which hide sensitive information from view while leaving other information unobscured. Results of two pilot studies supported the viability and potential usefulness of the privacy blinder concept, and have set the stage for continued development of the technique through large-scale controlled studies and field tests.

Keywords

Public information display, privacy, security, mobility

ACM Classification Keywords

H5.2. Information interfaces and presentation (e.g., HCI): User Interfaces.

Introduction

Maintaining privacy in the mobile environment remains difficult because the context of a device or application can change rapidly and without notice. This is in sharp contrast to a fixed environment, like an office, where people can consistently control the way that

Copyright is held by the author/owner(s).

CHI 2006, April 22–27, 2006, Montréal, Québec, Canada.

ACM 1-59593-298-4/06/0004.

information is handled to minimize the chance of divulging sensitive information to unauthorized parties. For example, computer screens can be pointed towards a user such that other people who enter their office cannot easily read them [3]. When the user's environment is not fixed, and can potentially change in a fraction of a second, privacy management becomes a challenge, but one that must be met to ensure the trustworthiness of ubiquitous computing applications.

People are not intentionally careless when it comes to protecting information in public places, but normal human behavior makes it easy for unsafe conditions to exist. As an example, laptop computers are often used whenever and wherever needed or desired (e.g., in an airplane). Technology is rapidly reaching a point of allowing unrestricted Internet access to information. But in these situations, the user can become more focused on the task at hand rather than the fact that information might be overseen or recorded by someone close by. While current technology makes it easy to access information anywhere and anytime, it does not provide adequate protection of that information at the same time. Effective solutions to privacy protection problems must not only be technically sound, but usable and understandable from a social perspective. Mobile users need interaction methods that work well with multiple and varied tasks, and in environments that can change rapidly. In addition, information must be readily accessible and presented in a clear and comprehensible manner. Therefore, innovations are needed to protect the information that users access on their mobile devices without adding to the existing complexity of the mobile environment. If this is not accomplished, users must accept certain tradeoffs

between the pervasive availability of information and the potential loss of privacy and security [3].

Our research does not seek to create new ways of securing information as it travels or is stored on various devices. While privacy is often maintained through methods (e.g., encryption) that keep data from being read by unauthorized parties, this research looks at the relatively unexplored but equally important problem of maintaining the privacy of displayed information. Our overall goal is to create technically sound but practical methods of maintaining privacy of sensitive information that is displayed in public and mobile environments. The underlying systems used to access information cannot be affected by these solutions, and the user should experience little or no additional cognitive or task overhead from these measures. Any solution must also be resilient enough to work in any context (i.e., location and task independent). Ultimately, the solutions should be able to adapt to a changing context, such as when a user moves from a private office to a company conference room to an airport terminal.

Background

Privacy is valued and expected by most people to varying degrees. Usually an individual expects reasonable access to personal information while restricting strangers' access to this same information. Privacy requirements will also vary based on the type of information, and on the preferences of the information's owner (e.g., individual or organization). For example, results from a recent study [4] that looked at public information access showed participants felt most strongly about protecting the privacy of medical and financial information. This is consistent with a study on e-commerce privacy that shows people

are not comfortable divulging social security numbers, credit card and phone numbers, income, and medical information [1]. However, it has also been shown that attitudes towards privacy can differ quite broadly from actual behavior, and that people may not be willing to use technology that protects their privacy [2]. Therefore, if privacy management in the mobile environment requires too much effort on the part of the user, it will most likely be ignored or used incorrectly.

Several hardware-based solutions have been explored to solve the problem of maintaining information privacy on mobile displays. Privacy covers have been developed for laptop screens that provide a clear view of the screen's contents to the user but obscure the view to anyone looking at the screen from an angle. Contents of screens can also be blurred, readable only through devices such as special eyeglasses, or becoming partially clear only when eye-tracking hardware senses the user looking directly at a specific area on the screen. While such hardware-based methods are potentially valuable in protecting the privacy of information, they may have potential drawbacks in terms of 1) additional cost; 2) additional weight, bulk, and power consumption; 3) increased complexity of use; and 4) distortion or degradation of the displayed information, which could affect user task performance.

Our ongoing research concerns the development and testing of privacy blinders. Privacy blinders are removable software tiles that automatically cover sensitive information on a screen but can be temporarily removed by the user at their discretion to view information hidden underneath them. Through two pilot studies, we have started to validate the usefulness of the blinder concept, and have discovered additional

uses for blinders as well. We briefly discuss our work to date, and describe our ongoing testing.

Hiding Sensitive Information with Blinders

Privacy blinders mimic the use of yellow sticky-notes to cover parts of a larger document so that they are not viewable by others. Blinders can be used to provide a mixed display in which sensitive information is hidden (covered) but information not considered private is displayed normally. If the user decides to view the sensitive information, they can temporarily remove the blinder. For example, blinders on a tablet PC or PDA might be removed by touching them with a stylus. When the stylus is removed from the screen, the blinders reappear. It is also possible to create blinders that can only be removed with a certain gesture (e.g., writing a letter on them), thereby creating a level of security along with information privacy.

Furthermore, privacy blinders can automatically respond to a predefined organizational and/or personal "privacy policy," which specifies what types of information are covered under different circumstances. An organizational policy might be dictated by the company a person works for, while a personal policy is customized to a user's own comfort level and privacy requirements. This method can also account for user context changes; if a person moves to a less public space, they might turn off the blinder feature and view all information without obstruction. Context might also be taken into account automatically by the system. For example, a change in location from a private office to a public meeting room might modify privacy settings by design. This flexibility allows adaptation to the changing environment of the mobile device user.

First pilot study: Prototype and controlled experiment

To date we have completed two pilot studies with privacy blinders. The first study [5] tested the basic concept of the blinders in terms of usability and effectiveness. We used a limited-function Mozilla Web browser prototype in a controlled laboratory study with each subject searching three “canned” banking Web sites for specific information. Privacy blinders were displayed based on special HTML tags next to information that was defined to be sensitive in nature. The blinders were set to a predefined width and height of 80 pixels, and appeared centered directly over sensitive information. The user could reveal information protected by a blinder in one of two predetermined ways. For one, the blinder disappeared and revealed the information underneath it when the stylus was moved over it. The blinder reappeared when the stylus moved away. In the second, the privacy blinder disappeared for a total of 10 seconds when a special stylus gesture was made, then reappeared. Privacy blinders that overlapped one another functioned together as a group. Task times were significantly longer with blinders, and users found the gesture interface the most difficult to use, but most participants felt that the standard blinders would be useful in public settings. Figure 1 shows a sample Web page with blinders obscuring elements such as dollar amounts.

Second pilot study: Browser plug-in and field testing

In the second phase, a FireFox Web browser extension with user-configurable privacy settings allowed users to browse personal information from any Web site. Unlike the software used in the first study, this extension did not simply look for special predefined HTML tags. It preprocessed any Web page, located user-defined sensitive information, placed user-customized blinders

on top of the specified content, and presented the resulting “blinded” Web page to the user.

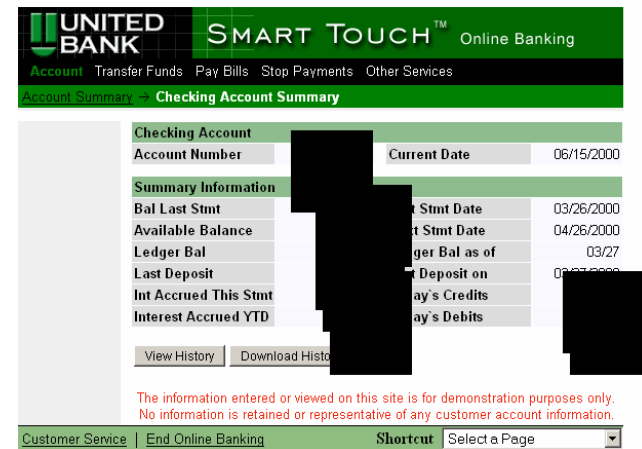


Figure 1. Sample banking screen with blinders enabled.

Properties of the privacy blinders that could be user-configured included 1) whether to group overlapping blinders, 2) to use fixed or variable blinder sizes, and 3) setting the opacity (transparency) of the blinders. Besides these blinder properties, the “privacy policy” could also be customized. Four classes of potentially sensitive content were supported. These were 1) monetary amounts (any number starting with a dollar sign or containing exactly two decimal points); 2) email addresses (in the format abc@xyz.dom); 3) telephone numbers (formats such as xxx-xxx-xxxx); and 4) phrases, which allowed the user to specify a delineated list of words and phrases that they wanted covered.

After showing each of seven participants how the software worked, they were instructed to spend 30

minutes in the public lobby of our College and perform tasks such as reading email and checking bank account balances. They were asked to take notes about where they went, what they did, which privacy settings they changed, and any comments they had. Subjects were then debriefed after the testing. Figure 2 shows a random webpage from the Internet viewed with the privacy blinders extension enabled.

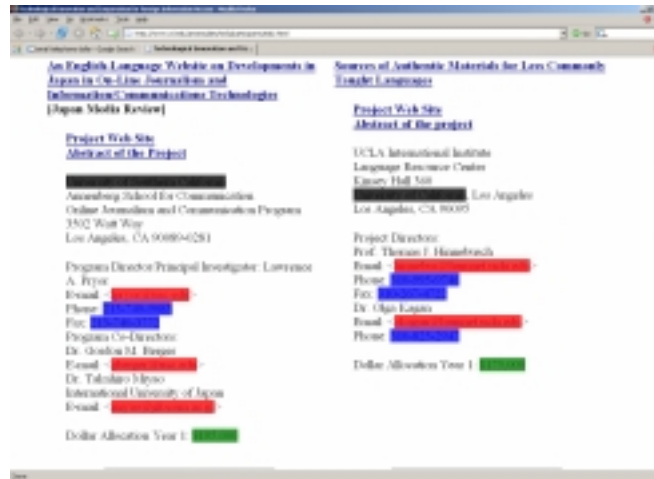


Figure 2. A random webpage with privacy blinders enabled.

All participants responded positively to using the privacy blinders, and felt that privacy blinders would be useful on a PDA or cell phone. The majority reported feeling more secure while surfing the pages with the blinders enabled than they would have otherwise. In general all participants reported the blinders were easy to use and did not impinge on their ability to use the site, although the majority indicated that it was somewhat annoying when too many items were

covered on a screen. Many participants felt that privacy blinders would be most useful on a laptop or desktop computer. In all cases they referred to the utility in office or lab settings where there were coworkers or peers able to see their screens. One participant brought up the process of logging into his cell phone voicemail, and noted that the password was displayed on the screen. This observation seems to hold with any type of cell phone interaction, and is a privacy issue that we will address in future research.

Continuing Work

The original intent of the privacy blinders had been to protect sensitive personal and financial information from onlookers, specifically on mobile devices. The results of our pilot studies indicate that while the blinders would likely help in this regard, there are many alternative uses that we had not originally envisioned, including parental filters, workplace privacy, highlighting, and color-coding of sensitive terms.

Parental filters came from the ability of a user to enter a list of words and phrases to be covered. A concerned parent could use this to block terms they deemed inappropriate for their children. Highlighting occurs by adjusting the opacity down until the underlying text is clearly visible but surrounded by the colored blinder. The notion of color-coding sensitive terms was suggested as a way to allow people to search for topics/words they might otherwise be embarrassed to have accidentally overseen by others. If the privacy blinders interface were extended so that the user could specify distinct colors for separate words then the user would not even have to uncover the blinder while navigating the site since the meaning of the colored box would be readily apparent to the user.

One difficulty is determining what content to cover. At current the plug-in relies upon a simple matching paradigm to determine what to cover. Given the limitations we decided to err on the side of matching too much rather than too little, but this proved frustrating because users did not like it when too much was covered. However, the paradigm is sufficient to cover the vast majority of cases.

This research has introduced the concept of privacy blinders as a software-based method for protecting sensitive information on mobile devices in public settings. Work continues on examining different sizes and shapes of blinders, alternate ways of placing and removing blinders, and degrees of user customization (along with the usability of the interface for defining the desired privacy levels). One idea is rather than completely removing a blinder, it might be set to drop its opacity to a certain level, allowing the user to view the underlying information through a translucent panel, but still discouraging onlookers who are at a distance from the screen. When this is done, the software can also be compared directly against hardware-based techniques for screen privacy. Integrating fixed organizational privacy settings with user customizable settings will also be investigated.

We are also planning to run longitudinal field tests where participants will be given a version of the software to run on their personal devices. The software will automatically track information about the privacy settings, and how often the blinders appear on various Web sites. The software will prompt the user for feedback after they perform tasks using the browser. We eventually hope to make a version of the plug-in

available for general user consumption. Versions of the privacy blinder software will also be created to run on PDA's and mobile phones.

Context data might also be used to automatically ensure that a user is interacting with a mobile information system in the safest possible manner, while still allowing for the greatest ease of use. We plan to look at using various types of context data (such as location, co-location, and scheduled events) to increase the effectiveness of privacy management by shifting more burden of environmental awareness from the user to the system. Context data will be used in conjunction with privacy preferences, rules, and guidelines from both a user and organizational point of view. A truly adaptive mobile system would take into account relevant changes in the user's environment on a real-time basis and modify privacy settings as appropriate.

References

- [1] Ackerman, M.S., Cranor, L.F., and Reagle, J. Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences. *Proc. ACM Conf. On Electronic Commerce* (1999), 1-8.
- [2] Acquisti, A. and Grossklags, J. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3, 1 (2005), 24-30.
- [3] Dourish, P., Grinter, R.E., Delgado de la Flor, J., and Joseph, M. Security in the Wild: User Strategies for Managing Security as an Everyday, Practical Problem. *Personal and Ubiquitous Computing*, 8, (2004), 391-401.
- [4] Grimes, A. and Tarasewich, P. Testing Privacy-Augmented Displays on a Mobile Device. *Proc. of HCI 2005*.
- [5] Tarasewich, P. and Campbell, C. What Are You Looking at? *Proc. SOUPS 2005*, ACM Press (2005).