Towards Everyday Privacy for Ubiquitous Computing

Scott Lederer[1], Jason I. Hong[1], Xiaodong Jiang[1], Anind K. Dey[1,2], James A. Landay[3,4], Jennifer Mankoff[1]

[1]Department of Computer Science, University of California, Berkeley
[2]Intel Research Berkeley, 2150 Shattuck Ave. Suite 1300, Berkeley, CA 94704
[3]Intel Research Seattle, 1100 NE 45th Street, Seattle, WA 98105
[4]Department of Computer Science, University of Washington

ABSTRACT

The goal of everyday privacy is to make it easy for end-users to share information with the right people at the right level of detail in ubiquitous computing environments. In this paper, we describe a conceptual model we have developed for everyday privacy, consisting of control over and feedback about disclosure. We also describe a prototype we have created for helping end-users manage their personal privacy, an evaluation of that prototype, and a revised prototype based on feedback from the evaluation.

Keywords

Privacy, Ubiquitous Computing, Privacy User Interfaces, Conceptual Models, Everyday Privacy

INTRODUCTION

A great deal has been written about information privacy and ubiquitous computing, including analyses of the many potential risks (cf. [1-6]); legal, social, and economic treatments (cf. [7-14]); and technological approaches for managing information flow (cf. [15-17]). In this paper, we focus on *everyday privacy*, that is, how end-users might manage their privacy in such environments.

Given the many social, technical, and regulatory phenomena commonly lumped under the term "privacy," we want to be clear about what we mean by *everyday privacy*. We are concerned herein with individuals' *regular, voluntary disclosure of their dynamic contextual information to remote identifiable parties through ubicomp systems.* We focus on cases where users are interested in sharing their personal information with the right people at the right level of detail. Here, people *want* to disclose contextual information to others, perhaps to facilitate micro-coordination of arrivals at a meeting place or to convey a sense of presence to intimate companions. The key here is that people choose to selectively disclose their information to people in their social and professional networks. This view of everyday privacy positions it as a nuanced social practice rather than a functional goal.

This does not mean that secrecy and anonymity are not important. There are many reasonable cases for not disclosing information to colleagues, friends, and family. The problem, however, is that secrecy and anonymity only focus on an extreme case, and cannot cover the many situations in everyday life where people *do* want to share information with others. Again, the question is how to share personal information with the right people at the right level of detail, with secrecy and anonymity possibly being the solution in some cases.

The overall goal of our research is threefold: (1) to understand the nature and range of privacy concerns with respect to ubicomp systems; (2) to elucidate a conceptual model of *everyday privacy* in ubiquitous computing environments; and (3) to use this model to develop user interfaces that provide end-users with appropriate control and feedback mechanisms for managing their privacy.

In previous work, we performed a questionnaire-based study to investigate the relative importance of key factors in determining individuals' privacy preferences in ubicomp [18] and proposed a design framework for minimizing information asymmetry in the software architecture of privacy-aware systems [19]. In this workshop paper, we focus on a conceptual framework and on interfaces for end-users to manage everyday privacy. The core of this framework is providing feedback about inquirers and control over the *precision* of context information disclosed to them. We present two user interfaces that instantiate aspects of this framework and discuss the results of a formative evaluation.

EXAMPLE SCENARIOS

Rather than discussing privacy in the abstract, we describe some scenarios that we are trying to support and believe would be realistic within 5-10 years.

- Alice's workplace has small beacons setup in each room that chirp out the room number. Alice's PDA can be informed what room it is in without anyone else knowing. She is only somewhat concerned about her privacy in this environment, and so sets things up to let her co-workers see where she is at the floor level.
- Bob has purchased a new location-enabled phone that knows where he is within 20 meters. He wants to set it up so that his family and close friends can see where he is.

- Carol is going out to a block party with some people. Some of them are her friends, but others she has just met. She wants to make her location available to her new acquaintances to make it easier to coordinate if anyone gets lost, but only wants to provide access for a limited amount of time since she has only just met them.
- David is throwing a surprise party in his home for Edith. However, he wants to ensure that Edith cannot use a ubicomp system to determine that all of her friends are there prior to the surprise.

END-USER CONTROL OVER DISCLOSURES

Central to our model of everyday privacy in ubicomp is the ability to adjust the *precision* of dynamic contextual information disclosed to other people. Control over precision provides control over the information density of a given disclosure. The more precise the disclosure, the more information disclosed. For example, disclosing location and duration as "Last seen in downtown San Francisco" is decidedly less revealing than disclosing "At a strip club on Market Street for the past two hours."

Adjusting the precision can be useful when disclosing information to *institutions*. For example, in the first scenario where Alice is sharing her location with co-workers, she might not feel comfortable letting them know precisely where she is right now. She can adjust the level of disclosure by modifying the *spatial precision* of her location information (*e.g.* "5th floor" or "in" rather than "room 525") as well as the *temporal precision* (*e.g.* "last seen 15 minutes ago" rather than "right now") until she finds a level she is comfortable with.

Adjusting precision can also be useful when disclosing information to *individuals*, as it can help people manage the persona they wish to reveal to others. As Goffman explains, people behave in such a way as to impress upon observers that they are acting in accordance with the roles they occupy in relation to those observers [20]. They work, usually intuitively, to convey that their activities are socially and contextually appropriate. Similarly, control over precision allows people to manage how they present themselves in ubiquitous computing environments. Users might not be able to convey that they are acting appropriately with respect to a given observer, but the lack of information conveys a likelihood that they are *not acting inappropriately*. The observer is left to infer some range of possible activities. If the observer's trust in the subject is sufficiently low to inspire the assumption of inappropriate activity, then this is a matter for the observer and subject to work out together, a social process that no amount of ubicomp technology can replace.

This approach relies on ubiquitous computing applications and infrastructures having some level of inherent ambiguity for purposes of plausible deniability. As an analogy, if a person does not answer a cell phone call, it could be for technical reasons—such as being outside of a cell, not having the phone with them, or that the phone is off—or for social reasons, such as being busy or not wanting to talk to the caller right now. The result is that the person being called has a simple model for protecting their privacy, while the caller cannot tell why that person is not answering. A similar situation will likely exist in ubiquitous computing systems, where ambiguity exists due to sensor noise, incomplete wireless coverage, and forgetting to carry devices, as well as situational propriety and the subjective desire to be let alone. The range and validity of reasons for both intentional and unintentional imprecision creates an ambiguity around the conditions of disclosure and conveyance that is too thick for observers to bother deconstructing. Indeed, cultural accommodation to imprecise knowledge on the part of remote observers can engender a norm of plausible deniability [21, 22].

Imprecision Desensitizes Information

Our focus on adjustable precision of contextual information is based on two different ideas. The first is the Principle of Minimum Asymmetry [19]. At a very high level, one can think of ubicomp systems having two knobs, one controlling information going out about you and the other knob controlling the feedback coming back in about how your personal information is being used. The main idea behind minimum asymmetry is that systems should be built such that the first knob is minimized to what is needed (*i.e.* as little information as necessary goes out) and the second knob is maximized (*i.e.* as much information about how your information is used comes back in).

The second is Adams' work on perceptions of privacy in multimedia environments [23]. Adams identified four key factors that determine a subject's perception of privacy in sensed environments:

- the subject's trust in the perceived *recipients*,
- the subject's cost/benefit analysis of the perceived *usage* of the information,
- the subject's subjective judgment of the information's *sensitivity*, and
- the *context* in which the information was collected.

We also found supporting evidence for the importance of these factors in an earlier questionnaire-based study [18]. Subjects reported a stronger likelihood to adjust the precision of a disclosure depending on the identity of the inquirer than on their situation at the time of inquiry.

In developing our model, we sought to provide end-users appropriate feedback and control over these factors. However, since usage is difficult to represent, we focused only on recipients, sensitivity, and context. Hence our model for managing everyday privacy in ubicomp emphasizes feedback about *inquirers*' identities, control over the precision of disclosed context (thereby adjusting

sensitivity), and optional *context*ual triggers to automatically adjust precision in specific contexts.

Ordinal Precision Scale

User interfaces for everyday privacy have to be effective in providing control and feedback, but also simple enough so that they will be actually used. Thus, we developed a simple ordinal scale to uniformly represent various precision levels, because it combined the greatest amount of control with a simple interaction.[1] The actual values disclosed at these precision levels would depend on technical factors, standards, and the subject's true context at the time of inquiry. Using location as an example, the scale is:

- *Precise* (ex. Barstucks Café at 123 New Montgomery)
- *Approximate* (ex. San Francisco Financial District)
- *Vague* (ex. San Francisco)
- *Undisclosed* (ex. Unknown)[2]

At a conceptual level, this ordinal scale is similar to the privacy slider in Microsoft's Internet Explorer (MSIE). The main difference here is that MSIE's privacy model focuses on the transfer of cookies, which can be indirectly linked with static information such as identity and shopping history[3], whereas our everyday privacy model focuses on the disclosure of dynamic information, with more direct control over the information that is shared with others. Further, we do not claim that our scale adjusts privacy directly. It adjusts the *precision* of disclosed context.

Privacy is not a slider bar; not on the desktop, and not in ubicomp. But we believe that one-dimensional controls like sliders or dials can be useful for adjusting information precision, thereby helping people manage their privacy.

USER INTERFACE PROTOTYPES

In this section, we discuss our initial prototype that instantiated aspects of this model, the formative evaluation of that user interface, as well as our revised user interface for managing privacy.

Prototype 1 – Faces Metaphor

Our first prototype was a desktop interface that allowed users to specify arbitrary inquirers, contexts, and privacy preferences, and to assemble them into arbitrary rules for managing disclosure (See Figure 1). Additionally, a log recorded information about all inquiries and disclosures, and a handheld interface provided real-time feedback and limited control over disclosure. We represented context as the sum of location, activity, companions, and time, which we called *situations*. So, for example, the situation "studying" might consist of "location = library" and "activity = reading or computing".

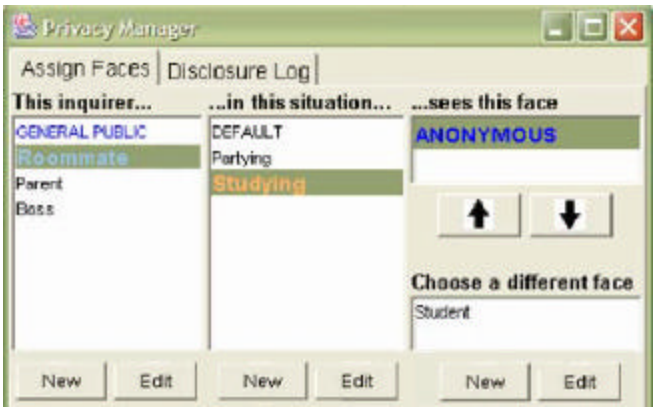

Figure 1. PC-based user interface for creating and assigning faces. Each face represents a set of privacy preferences that is used when a given inquirer makes a request when the user is in a given situation.

Preferences took the form of *faces*, which encapsulated the precision preferences for each of various information dimensions, including name, location, activity, and companions. This metaphor was inspired by Goffman's analysis that individuals present different "faces" to different people in different social situations [20]. So, for example, the face "student" might specify "precise" disclosure of identity, location, and activity but only "approximate" disclosure of companions.

Users could assign a face to handle disclosures for each inquirer and for each situation the user might be in. For example, the configuration in Figure 1 can be read as "if a Roommate makes a request while I am studying, show my Anonymous face."

An essential design issue here is that control and feedback can be overwhelming if always employed in real-time. To address this, the faces interface time-shifts control and feedback to before and after an information request, respectively. Control takes place when users set up their face preferences. Feedback takes place in an integrated log of disclosures reviewable after disclosure. Users could navigate their log to help them understand what information is flowing to whom, and they can reconfigure their preferences in response to unfavorable disclosures, to ensure they do not happen again[4].

However, a formative Wizard of Oz evaluation with five participants revealed several flaws with this prototype.

---

[1] Finer-grained, advanced options may be appropriate at a deeper layer of the interface.

[2] A *custom* setting would enable the disclosure of false information. We welcome investigations of this option, but do not address it in this paper.

[3] Identity and shopping history are not disclosed *per se*, but rather cookies that might be linked to that information are disclosed.

[4] This may seem naïve, especially considering that once some information is disclosed, that disclosure can never be erased. Nonetheless, people make unwanted disclosures in everyday life and employ social mechanisms (*e.g.* apologies) to amend them. Ubicomp will be no different.

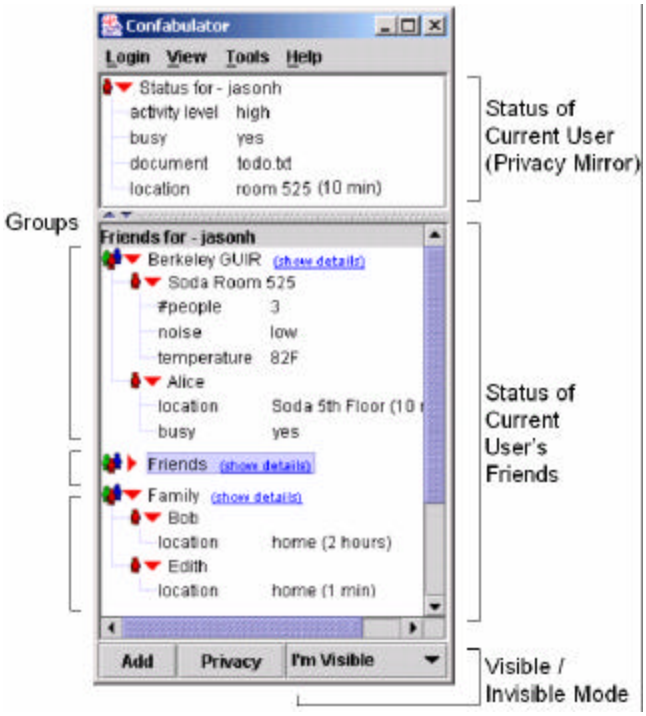

Figure 2. Our second prototype for managing everyday privacy makes use of an instant messenger interface style. The top half of this mockup shows the current user's status, while the bottom half shows the status of the user's friends.

infospaceinteractive: bob@cs.berkeley.edu at guir.berkeley.edu is requesting your location
(1) My location is 525
(2) My location is 5th Floor
(3) My location is Soda Hall
(4) Manually correct my location
(5) My location is Unknown (default in 45 seconds)
(6) Go invisible
(7) More choices
Please select an option
jas0nh0ng: 1
infospaceinteractive: Let bob@cs.berkeley.edu see that my location is 525
(1) Just this once
(2) Just for today (default in 20 seconds)
(3) Always do this
(4) More choices
(0) Go back
Please select an option
jas0nh0ng: 2
infospaceinteractive: Automatically saving settings in 5 seconds, send any message to cancel
infospaceinteractive: Setting saved

Figure 3. An example of *in situ* configuration. People can receive requests for personal information and be able to set their preferences as an activity happens rather than trying to predict abstract situations beforehand. Here, Bob is not already part of the user's list of friends. When Bob makes a request for the user's location, the request is

First, participants typically used only a single aspect of context to create situations. For example, they would only care about location or activity, but not both. So, encapsulating contextual variables into *situations* proved superfluous and confusing. Additional confusion occurred because subjects saw no difference between a situation and an activity (and indeed, there is none). Second, the face metaphor obfuscated the true meaning of the preferences, namely, information precision. Subjects made it clear that their "face" (metaphorical or not) had little to do with the precision of, say, their disclosed location. Third, the fact that static and dynamic information were combined within a face caused confusion. Faces encapsulated both dynamic information—such as location and activity—and static information, like name and profile. However, static information was more sensitive with unfamiliar inquirers, while dynamic was more sensitive with respect to familiar inquirers. Combining both static and dynamic information caused problems in defining appropriate faces. Lastly, requiring an explicit and complex structure for managing end-user privacy was untenable. Privacy is a highly subjective process and is not easily captured by, and is arguably disrupted by, a system requiring explicit preferences set *a priori*.

In sum, the faces interface modeled Goffman's theory literally *in* the interface, but did not facilitate the intuitive impression management practice *through* the interface. We believe that a more lightweight, flexible approach is warranted, possibly involving on-the-fly adjustment of disclosure precision to a small set of inquirer groups.

Prototype 2 – Instant Messenger Metaphor

Because of the complexity involved with the faces metaphor, we are making six major changes in our second prototype (Figure 2 is a screenshot of our work in progress). First, we are basing the overall design on instant messenger clients. IM is an interface style many end-users are already familiar with, and shares several characteristics with what we are trying to accomplish, including adding and removing friends, customizing status messages (sometimes used to disclose location or activity), and organizing friends into groups.

Second, we are focusing exclusively on managing dynamic contextual information rather than static information. Our formative studies indicated that people found it difficult to combine both of these into a single conceptual model. Since much of ubiquitous computing focuses on representing dynamic state, we decided to focus on managing that aspect. This also made sense because inquirers and friends would presumably already know some static information about you, such as who you are.

Third, we are modifying the interface to show real-time data reflecting the status of the current user (top half of Figure 2) as well as the status of the user's friends (bottom half of Figure 2). The rationale for this is that the new interface can cover more basic tasks, such as adding and removing friends, setting preferences, and viewing someone's current

status. Displaying the current user's status also provides feedback about what the system knows about that user, making it a sort of privacy mirror [24].

Fourth, we are modifying the preferences to be set for *groups*, regardless of situation, rather than for individuals. We made this decision for scaling purposes. We felt that it was likely that a person would have similar preferences for people within the same group, and that it would be tedious to have to set individual preferences. Situations were eliminated altogether because our earlier study suggested that identity was more important than situation when disclosing information [18].

Fifth, we have added *invisible mode*, a familiar interaction technique from instant messenger. This provides a clear and simple mechanism for end-users when they want to hide all information from all people.

Sixth, we have added the ability to do *in situ* configurations (see Figure 3). One of the problems with setting up all of one's preferences *a priori* is that it is difficult to predict all possible circumstances one might be in. To address this, we created a messaging interface that would make it easy to adjust preferences on the fly. The end-user is notified when someone makes a request that cannot be handled by existing preferences (in an instant message window or through SMS, for example). The end-user can then make decisions on what kind of information to share. Figure 3 shows how *in situ* configurations can be used to share information temporarily with others. Bob makes a request for the user's current location. If Bob is not on the user's friends list, or is requesting more detailed location than he is currently allowed to see, then a message is sent to the user. The user can then decide what level of information to disclose and for how long.

The user interface for *in situ* configurations is currently implemented using Yahoo Instant Messenger. It is part of the second prototype, but represents a distinct user interface for managing different aspects of one's privacy. The interface in Figure 2 is meant for setting up preferences beforehand on a desktop computer, whereas the interface in Figure 3 is meant for managing configurations on the fly on mobile devices.

The second prototype is built on top of Context Fabric infrastructure [25]. We are currently in the process of connecting the user interface in Figure 2 to a backend, and have completed development of the interface in Figure 3. We are also preparing for another user evaluation of these interfaces.

DISCUSSION

Since this is work in progress, rather than ending with conclusions, we look at how this work fits into the larger picture of analyzing, designing, constructing, and evaluating privacy-sensitive applications for ubiquitous computing. These include:

- Is a centralized nexus for managing privacy the right approach? Whitten and Tygar [26] point out that, in the security domain, "Security is usually a secondary goal. People do not generally sit down at their computers wanting to manage their security; rather, they want to send email, browse web pages, or download software, and they want security in place to protect them while they do those things." On the other hand, decentralizing control and feedback across multiple applications can lead to inconsistency, confusion, and inefficiency. One possibility may be a hybrid solution, having a centralized privacy control panel as well as application-specific controls, similar to how printers are managed in modem operating systems.
- One issue that arose from our evaluation is that it takes a great deal of work to setup privacy preferences. This is because we are making explicit something that people do implicitly and intuitively, and because privacy is a secondary goal. What are better approaches for helping end-users do this initial setup? One possibility is to use web crawling combined with data mining on address books, instant messenger contacts, and emails sent and received to come up with a rough cut of groups and settings, and then let people manually twiddle with the settings. What other possibilities are there?
- Currently, giving someone your mobile phone number is all or nothing, *i.e.* there is no way of lowering the precision. However, there is still a certain level of plausible deniability. For example, if you cannot contact someone, it could be because the phone is turned off, they are not carrying it with them, they are outside of any cells, or they simply do not or cannot talk to you right now. On the other hand, access to your location is not all or nothing: you can provide it at various granularities and for various durations. Is it socially acceptable to give someone only temporary access to one's information, or in providing less precise information? If so, are there ways of adding some level of ambiguity for plausible deniability purposes to smooth things out? At a higher level, what kinds of intentional ambiguity are needed for ubicomp? Can we build these directly into systems, perhaps even as intentional flaws? And should we?
- In this paper, we have focused primarily on a *pull-model* for everyday privacy, where you give people limited access to dynamic information about you. An alternative is the *push-model*, where you pass your information along with a request (for example, passing your current location to E911 only when you call). What are the different tradeoffs of these architectural styles in terms of privacy, utility, deployment, and robustness? What kinds of services are supported by

each, and consequently, which architectural style will be more common in the future?

REFERENCES

1. Agre, P.E. and M. Rotenberg, *Technology and Privacy: The New Landscape*. 1997, Cambridge MA: MIT Press.
2. Grudin, J., *Desituating Action: Digital Representation of Context.* Human-Computer Interaction (HCI) Journal, 2001. **16**(2-4).
3. Doheny-Farina, S., *The Last Link: Default = Offline Or Why Ubicomp Scares Me*, in *Computer-mediated Communication*. 1994. p. 18-20.
4. Garfinkel, S., *Database Nation: The Death of Privacy in the 21st Century.* 2001: O'Reilly & Associates.
5. Talbott, S., *The Trouble with Ubiquitous Technology Pushers, or: Why We'd Be Better Off without the MIT Media Lab*. 2000.
6. Marx, G., *The Iron Fist and the Velvet Glove: Totalitarian Potentials Within Democratic Structures*, in *The Social Fabric: Dimensions and Issues*, J. James E. Short, Editor. 1987, Sage Publications: Beverly Hills.
7. Marx, G., *Murky Conceptual Waters: The Public and the Private.* Ethics and Information Technology, 2001. **3**(3): p. 157-169.
8. Marx, G., *Identity and Anonymity: Some Conceptual Distinctions and Issues for Research*, in *Documenting Individual Identity: The Development Of State Practices In The Modern World*, J. Torpey, Editor. 2001, Princeton University Press.
9. Langheinrich, M. *Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems*. in *Ubicomp*. 2001. Atlanta, GA.
10. Acquisti, A. *Protecting Privacy with Economics: Economic Incentives for Preventive Technologies in Ubiquitous Computing Environments*. in *Workshop on Socially-informed Design of Privacy-enhancing Solutions, 4th International Conference on Ubiquitous Computing (UBICOMP 02)*. 2002. Goteborg, Sweden.
11. Bellotti, V. and A. Sellen. *Design for Privacy in Ubiquitous Computing Environments*. in *The Third European Conference on Computer Supported Cooperative Work (ECSCW'93)*. 1993. Milan, Italy: Kluwer Academic Publishers.
12. Harper, R.H.R., *Why Do People Wear Active Badges?* 1993, Rank Xerox: Cambridge. p. 10.
13. Lessig, L. *The Architecture of Privacy*. in *Taiwan NET'98*. 1998. Taipei, Taiwan.
14. Lessig, L., *Code and Other Laws of Cyberspace*. 1999, New York NY: Basic Books.
15. Beresford, A. and F. Stajano, *Location Privacy in Pervasive Computing*, in *IEEE Pervasive Computing*. 2003. p. 46-55.
16. Fasbender, A., D. Kesdogan, and O. Kubitz. *Analysis of Security and Privacy in Mobile IP*. in *4th International Conference on Telecommunication Systems, Modeling and Analysis*. 1996. Nashville, TN.
17. Langheinrich, M. *A Privacy Awareness System for Ubiquitous Computing Environments*. in *Ubicomp*. 2002. Goteberg, Sweden.
18. Lederer, S., J. Mankoff, and A.K. Dey. *Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing*. in *Extended Abstracts of CHI 2003, ACM Conference on Human Factors in Computing Systems*. 2003. Fort Lauderdale, FL.
19. Jiang, X., J.I. Hong, and J.A. Landay. *Approximate Information Flows: Socially-based Modeling of Privacy in Ubiquitous Computing*. in *Ubicomp 2002*. 2002. Göteborg, Sweden.
20. Goffman, E., *The Presentation of Self in Everyday Life*. 1956, New York, NY: Doubleday.
21. Nardi, B.A., S. Whittaker, and E. Bradner. *Interaction and Outeraction: Instant Messaging in Action*. in *Proc. ACM CSCW Conf.* 2000. New York, NY: ACM.
22. Woodruff, A. and P.M. Aoki. *How Push-to-Talk Makes Talk Less Pushy*. in *Proc. ACM SIGGROUP Conf. on Supporting Group Work (GROUP '03)*. 2003. Sanibel Island, FL: ACM.
23. Adams, A. *Multimedia Information Changes the Whole Privacy Ball Game.* in *Computers, Freedom, and Privacy*. 2000. Toronto, Canada: ACM Press.
24. Nguyen, D.H. and E.D. Mynatt. *Privacy Mirrors: Making Ubicomp Visible*. in *Human Factors in Computing Systems: CHI 2001 (Workshop on Building the User Experience in Ubiquitous Computing)*. 2001. Seattle, WA: ACM Press.
25. Hong, J.I. and J.A. Landay, *An Infrastructure Approach to Context-Aware Computing.* Human-Computer Interaction (HCI) Journal, 2001. **16**(2-3).
26. Whitten, A. and J.D. Tygar. *Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0*. in *8th USENIX Security Symposium*. 1999.