Privacy Mechanisms for Context-Aware, Group-based Mobile Social Software

Karen P. Tang
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
kptang@cs.cmu.edu

ABSTRACT

With the recent surge of mobile social software and the increasing trend of mobile adoption, users are becoming inundated with opportunities to disclose their contextual information for these systems, but are often left without the proper tools to preserve their privacy when doing so. The goal of my research is to address this gap in privacy management by designing, developing, and evaluating a set of privacy mechanisms that are specifically targeted for protecting end-user personal privacy when disclosing contextual information for mobile social systems.

ACM CLASSIFICATION: H5.2 [Information interfaces and presentation]: User Interfaces.

General terms: Design, Human Factors

Keywords: Privacy, Context Aware, Mobile Social Software, Groupware

INTRODUCTION

It is now commonplace to find mobile devices that are networked, location-enabled, and support voice, text input, photos, and video. Moreover, many of these mobile devices are not only becoming ubiquitous, but they are also being carried around by their owners nearly everywhere they go.

Together these features have introduced a relatively recent surge in the deployment of mobile software which focuses, not on productivity, but on meeting people's social needs. Examples of such software include applications which aim to better facilitate awareness, communication, coordination, and sharing when users are on the go. In addition, many of these mobile social applications target groups of users, ranging from just a few people (e.g., coordinating small workgroups and families) to thousands of people who may not even know each other (e.g., a laptop-based system that detects and shares how crowded places such as cafés are, based on wireless network traffic).

While there are many factors that may impede mobile social systems from enjoying widespread deployment (such as cost of mobile bandwidth, difficulty of implementation, and critical mass), perhaps one of the more significant factors is the lack of design for implementing these systems in a privacy-sensitive manner. This is particularly problematic given how many of these services implicitly require users to disclosure a significant amount of contextual information. Thus, a critical challenge to developing and deploying mobile social software is to understand how to incorporate privacy mechanisms that are easy for end-user to use to manage their personal privacy.

To address this problem, I intend to design, develop, and evaluate various privacy-aware mechanisms in several mobile social system prototypes. In this paper, I describe both completed and planned research for three primary contributions related to this goal. First, I present a set of privacy mechanisms to help preserve end-users' personal privacy when disclosing their contextual information in mobile social services. Specifically, I investigate how to better support privacy-aware collection of contextual information, privacy configurations for sharing context, and awareness support for informing users the consequences of their privacy configurations. Second, I present a suite of mobile social system prototypes that I intend to use as a test bed for evaluating these privacy-aware mechanisms. Each system is intended to exercise different context sensors and different target users. Lastly, I plan to distill what I have learned as guidelines and design patterns for developers and practitioners to use when designing and building future mobile social software.

DEFINING MOBILE SOCIAL SOFTWARE

Because its main purpose is to support social needs, it is not surprising that most mobile social applications are targeted for groups. Some of these applications (like Microsoft's SLAM [3] and PlaceMail [8]) are best suited for relatively small groups, for example co-workers, families, or groups of friends. A few of these applications can serve larger groups, such as twitter [18], dodgeball [4], SWARM [6], and ContextContacts [11]. There are also a few mobile social applications, like GeoNotes [5], which are designed for much larger populations, such as those capable of supporting on the order of thousands of people.

Most mobile social applications provide some degree of support for collecting contextual information and providing a means to share and display this other group members. Examples for such context sharing are plentiful: there are systems that disclose information about each other's presence [11, 16], location [14, 15], motion [1], and proximity [10]. In addition, many of these mobile social applications support some form of direct communication

UIST'07, October 7-10, 2007, Newport, Rhode Island, USA.
ACM 978-1-59593-679-2/07/0010

between group members. Most applications utilize SMS as the communication medium of choice, though others also use more novel forms, like synchronous vibrations or blinking to signal mutual awareness. In my work, I have chosen to use commodity devices, and thus I focus more on traditional means of awareness and communication.

Based on this brief overview, we can describe a majority of existing mobile social services as having three key features:

- They target both small and large groups
- They support a rich range of mobile, contextual information disclosures between group members
- They support asynchronous communication between group members.

Yet, many of these mobile social systems do not explicitly address privacy, or, if they do, only support a subset of privacy concerns. In the next section, I describe the focus of my work: the design, development, and evaluation of mechanisms to provide these kinds of mobile social systems with better support for protecting end-user privacy during contextual information disclosures.

PRIVACY MECHANISMS FOR MOBILE SOCIAL APPS

To fully appreciate the utility of most mobile social software, users are required to share their contextual information with other users. However, without proper privacy-aware mechanisms, these systems often suffer from deployments that fall short of attaining critical mass, resulting in mobile social applications which cannot sustain beyond the initial novelty effect. In my work, I introduce several ways that context-aware mobile social software can better support end-user privacy: 1) provide privacy-aware collection of contextual information, 2) provide usable privacy controls for context sharing, and 3) provide feedback to ensure users are aware of the effects and consequences of their privacy settings.

Privacy-Aware Collection of Contextual Information

Many mobile social systems collect contextual information by requiring client devices to continually upload their data to a centralized server. Moreover, these systems usually employ a person-centric model, whereby collected context data is tagged with unique identifiers that can be traced back to the user who is continuously uploading their information. This clearly poses as a potential privacy threat, especially in the event that the server holding the collected information is maliciously overtaken.

To address this privacy threat, I developed *hitchhiking*, a new approach to anonymous and privacy-sensitive collection of sensed contextual data for location-based applications [17]. Hitchhiking supports applications that combine location information from many people to infer busyness and density estimates for a particular place. Examples of hitchhiking applications include live traffic monitoring, inferring the availability of seats at particular places, estimating the arrival time of a bus, monitoring the busyness of a popular place, or monitoring wait times (like with airport security lines or restaurant waiting lines). The insight for these location-centric applications is that it is irrelevant who is providing the contextual information (location). Instead, hitchhiking treat locations as the entity of interest, and not the person providing the context. By aggregating the contextual information for a given location, hitchhiking enables individual users to anonymously provide their context without compromising the application's utility and ensures the server does not require unique identifiers that can be tracked back to end-users.

Rule-Based Privacy Controls for Sharing Context

While anonymity is one way to support privacy in mobile social software, it is often difficult to apply these principles to collecting all sources of context data. Thus, these systems must also consider other means to make it easy for end-users to maintain their personal privacy. In this section, I describe work involving the design of privacy controls to regulate what information is shared with other end-users.

Based on results from a lab study, work by Patil and Lai suggests that group-based privacy controls are sufficient for contextual information disclosure scenarios [13]. To evaluate this claim in a field deployment, I developed *IMBuddy*, a contextual instant messaging service that allows AIM users to query an AOL Instant Messaging Robot (AIMBot) for different types of contextual information, including interruptibility, location, and current window in focus (as a proxy for the user's current task).

Any AIM user can request a user's information by typing a command in a chat window to the AIMBot. For example, the query "*howbusyis* X" asks for screenname X's interruptibility. The AIMBot passes this request to the server, which communicates with the appropriate *IMBuddy* client to retrieve the requested contextual information. Based on the user's privacy control settings, *IMBuddy* reports the privacy-filtered response back to the requester in the original chat window (via the AIMBot). Information requests are also stored in a database on the server, which lets *IMBuddy* share the most recent disclosure information in the event the user in question is offline.

To manage what response requesters obtain, *IMBuddy* uses group-based privacy controls where users categorize their AIM buddies into separate groups defined by the amount of information the users want to disclose to them. For example, given a particular context (e.g. location), users could create up to three groups: one to disclose no information ("location is not available"), another to disclose their location with low granularity ("I'm at home") and the remaining group to disclose the information with high granularity ("I'm at 50 Walker Dr"). The response that the requester obtains from *IMBuddy* will then depend on which group the requester falls under. Note that there is a default group with a customizable level of granularity for requesters who have not been explicitly categorized into a group at the time of the request.

Using Context to Inform Policies for Sharing Context

Group-based privacy controls represent one way that users can define how to disclose their contextual information to other group members. There is also the option to define rules using more parameters, such as by time (e.g. "allow my location to be disclosed from 9am-5pm") or by location (e.g. "allow my location to be disclosed whenever I'm in the vicinity of work"). While more parameters may produce more customized and fine-tuned settings, it also makes privacy management much more complicated for end-users.

Alternatively, Palen and Dourish argue that privacy is more than authoring rules; it is an ongoing "boundary definition process" in which boundaries of disclosure, identity, and time are fluidly negotiated [12]. Group-based rules leverage social relationship as the primary factor in determining whether to disclose information to others. However, it may be more appropriate to consider this social relationship within a certain context.

Consider the decision of opening up your calendar information to other group members. Group-based rules would only consider who is asking when disclosing information. A better solution might be to disclose the contextual information based not only on *who* the requester is, but also *when* the requester is asking. For example, sharing calendar information with others may be permissible, but only if requesters have meetings scheduled on the same day that they are asking about, or if requesters ask around the time they already have meetings scheduled with you. Note that using context to inform disclosure decisions does not preclude the need for using some type of rule-based privacy controls. Rather, by leveraging context, we see two benefits: 1) we minimize end-users' needs to continually revisit their privacy control settings for special situations, and 2) requesters can obtain useful information without sacrificing the end-user's personal privacy.

To demonstrate and evaluate this type of privacy control, I developed *inTouch*, a phone-based mobile social system that bundles a contextual awareness panel together with a novel contextual messaging interface as a means to ease mobile coordination burdens for small groups. *inTouch*'s awareness panels include contextual cues like location ("at school") and activity cues ("at soccer practice"). The awareness panel supports two types of view models. The first view is person-centric, showing a complete set of each group member's contextual information. The second type provides a task-oriented view, where users can sort chronologically (which task deadline is approaching the soonest), by responsibility (who's completing which task), or by participant (who is involved in each task). Because mobile coordination relies heavily on communication, *inTouch* also features a modified mobile messaging interface using forms (or templates) instead of free-text input. We posit that these templates can: 1) increase efficiency in completing and sending coordination messages, and 2) lower the number of rounds required to resolve mobile coordination issues. Also, form-based inputs may allow a shallower learning curve and provide lower error rates than free-text input, since forms typically require only selecting from familiar widgets (like drop-down boxes, radio buttons, etc) rather than using traditional text input techniques which can be frustrating for novice SMS users.

While *inTouch* may seem light on the "social" aspect of mobile social software, particularly when used for groups like families or workgroups, it should be noted that *inTouch* is also designed for use by casual small groups that can form serendipitously at conferences, reunions, or other get-togethers. During these chance encounters where unanticipated groups are formed, rule-based privacy controls may not be the most appropriate mechanisms for preserving end-user privacy. This observation plus the fact that successful mobile coordination often requires disclosing lots of contextual information (e.g. location, busyness, current activity, etc), makes *inTouch* a useful mobile social system to evaluate using contextual triggers to inform privacy disclosure policies.

Incorporating Feedback to Increase Privacy Awareness

While privacy controls enable end-users to define how to share their contextual information with others, feedback is also important because if users are not aware of how their information is being disclosed, then they will be unable to react appropriately to scenarios where potentially harmful requests may result in undesirable information disclosures. Feedback can be classified as providing either delayed or immediate feedback, as discussed in Nguyen and Mynatt's work on Privacy Mirrors [9].

In our *IMBuddy* evaluation [7], I explored five different mechanisms covering both delayed and immediate feedback. Our results show that, of these five, the two most highly-regarded feedback mechanisms are: 1) a simple bubble notification that popped up whenever someone asked for information and 2) a disclosure log showing a history of who had requested what information and what was disclosed to them. Even more interestingly, people felt better knowing that a disclosure log was there, even if they rarely looked at it. Many people viewed it as a safety net that they could rely on if there was a need, even if they never actually used it. Other feedback mechanisms we provided included social translucency (to provide common grounding after the disclosure occurs, so that end-users are aware of what requesters know about them), privacy transparency (to inform end-users what would be disclosed to requesters, before the actual disclosure occurs), and shorten disclosure logs (showing only disclosures that occurred in the past five hours). Our evaluation showed that these mechanisms were not as highly rated.

IMBuddy's privacy feedback interactions represent a starting point for evaluating other types of feedback to help end-users better understand the consequences of their privacy settings for disclosing their context information.

Evaluating Privacy Controls and Feedback

Thus far, I have described several mechanisms to better support privacy in context-aware group-based mobile social software. To evaluate these mechanisms, I developed several types of mobile social systems in order to test not only the utility of the privacy mechanism, but also its usability for providing end-users an easy way to control their personal privacy. By building these prototypes, I can better explore the kinds of privacy features that should be generally useful for mobile social applications, as well as better understand which privacy controls and feedback interfaces are best for end-users. Through iterative designs and both short-term and long-term evaluations, I hope to better understand the end-user's mental model of these privacy mechanisms and to use that to inform design which are not overly burdensome for the user or interfere with the overall utility of the mobile social system.

For example, one method of managing contextual information disclosures is to interrupt the end-user each time there is a request for their information. While this mechanism affords a tightly controlled privacy management system, it places a heavy burden on the user as it incurs significant interruption costs. Thus, there is a tradeoff between the amount of control provided to the user and the management burden that the systems requires of the user in order to preserve their personal privacy. My goal is to help design privacy mechanisms which minimize the number of actions or interventions required from the end-user while still allowing flexibility in protecting the user's personal privacy when disclosing their contextual information.

FUTURE WORK: PRIVACY DESIGN GUIDELINES

For the remainder of my dissertation, I plan to finish the evaluation and iterative designs for the aforementioned privacy mechanisms, and to demonstrate their usability and utility in various mobile social system prototypes. As a final deliverable, I hope to summarize my findings into: 1) a set of controls and feedback interfaces that can be easily consumed by developers so that they too can incorporate these mechanisms into existing mobile social systems, and 2) a set of design patterns or guidelines to help inform practitioners in building future mobile social software.

ACKNOWLEDGMENTS

I would like to thank my advisors Jason Hong and Dan Siewiorek. This work was supported in part by the National Science Foundation grant IIS 0121560, an AT&T Labs fellowship, and an Intel Research fellowship.

REFERENCES

1. Bentley, F.R. and Metcalf, C.J. (2007). Sharing motion information with close family and friends. In *Proc. of ACM Conference on Human factors in computing systems (CHI 2007)*, 1361-1370.
2. Cornwell, J., Fette, I., Hsieh, G., Prabaker, M., Rao, J., Tang, K.P., Vaniea, K., Bauer, L., Cranor, L., Hong, J.I., McLaren, B., Reiter, M. and Sadeh, N. (2007). User-Controllable Security and Privacy for Pervasive Computing. In *Proc. of The 8th IEEE Workshop on Mobile Computing Systems and Applications (HotMobile 2007)*, (to appear).
3. Counts, S. Group-Based Mobile Messaging in Support of the Social Side of Leisure. *Computer Supported Cooperative Work*, *16* (1-2), 75-97.
4. Dodgeball, http://www.dodgeball.com/.
5. Espinoza, F., Persson, P., Sandin, A., Nyström, H., Cacciatore, E. and Bylund, M. (2001). GeoNotes: Social and Navigational Aspects of Location-Based Information Systems. In *Proc. of Ubicomp 2001*, 2-17.
6. Farnham, S. and Keyani, P. (2006). Swarm: Hyper Awareness, Micro Coordination, and Smart Convergence through Mobile Group Text Messaging. In *Proc. of 39th Annual Hawaii International Conference on System Sciences (HICSS'06)*.
7. Hsieh, G., Tang, K.P., Low, W.Y. and Hong, J.I. Field Deployment of IMBuddy: A Study of Privacy Control and Feedback Mechanisms for Contextual Instant Messengers. In *Proc. of Ubicomp 2007*, (to appear).
8. Ludford, P.J., Frankowski, D., Reily, K., Wilms, K. and Terveen, L. (2006). Because I carry my cell phone anyway: functional location-based reminder applications. In *Proc. of ACM Conference on Human Factors in Computing Systems (CHI 2006)*, 889-898
9. Nguyen, D.H. and Mynatt, E.D. (2001). Privacy Mirrors: Making Ubicomp Visible. In *Proc. of ACM Conference on Human Factors in Computing Systems (CHI 2001)*, Workshop on Building the User Experience in Ubiquitous Computing.
10. Nokai Sensor, http://www.nokia.com/sensor.
11. Oulasvirta, A., Raento, M. and Tiitta, S. (2005). ContextContacts: re-designing SmartPhone's contact book to support mobile awareness and collaboration. In *Proc. of MobileHCI 2005*, 167-174.
12. Palen, L. and Dourish, P. (2003). Unpacking "Privacy" for a Networked World. In *Proc. of ACM Conference on Human Factors in Computing Systems (CHI 2003)*, 129-136.
13. Patil, S. and Lai, J. (2005). Who gets to know what when: configuring privacy permissions in an awareness application. In *Proc. of ACM Conference on Human Factors in Computing Systems (CHI)*, 101-110.
14. Plazes, http://plazes.com/.
15. Smith, I., Consolvo, S., Lamarca, A., Hightower, J., Scott, J., Sohn, T., Hughes, J., Iachello, G. and Abowd1, G.D. (2005). Social Disclosure of Place: From Location Technology to Communication Practices. In *Proc. of Pervasive 2005*, 134-151.
16. Tang, J.C., Yankelovich, N., Begole, J., Kleek, M.V., Li, F. and Bhalodia, J. (2001). ConNexus to awarenex: extending awareness to mobile users. In *Proc. of CHI 2001*, 221 - 228.
17. Tang, K.P., Keyani, P., Fogarty, J. and Hong, J.I. Putting people in their place: an anonymous and privacy-sensitive approach to collecting sensed data in location-based applications. In *Proc. of CHI 2006*, 93-102.
18. Twitter, http://twitter.com/.