

User-Controllable Security and Privacy for Pervasive Computing

Jason Cornwell, Ian Fette, Gary Hsieh, Madhu Prabaker, Jinghai Rao, Karen Tang, Kami Vaniea, Lujo Bauer, Lorrie Cranor, Jason Hong, Bruce McLaren, Mike Reiter, Norman Sadeh
School of Computer Science – Carnegie Mellon University
sadeh@cs.cmu.edu

Abstract

We describe our current work in developing novel mechanisms for managing security and privacy in pervasive computing environments. More specifically, we have developed and evaluated three different applications, including a contextual instant messenger, a people finder application, and a phone-based application for access control. We also draw out some themes we have learned thus far for user-controllable security and privacy.

1. Introduction

Mobile devices and the services they support are increasingly becoming central in both personal and business life. The dramatic market growth of smartphones and portable storage devices suggests that the number of devices that contribute to personal, enterprise, and government computing environments will continue to increase. At the same time, the vast majority of these devices are unmanaged, and so with these new applications comes the need to enable lay users to handle the inherent security and privacy implications.

Managing security and privacy policies is known to be difficult. Even in desktop computing environments, end-users have great difficulty using the Windows XP file permission system [5]. In mobile and pervasive computing settings, this situation is often exacerbated by the limitations of devices and the numerous tasks users concurrently engage in. To make matters worse, desired security and privacy settings are not just difficult to articulate, but also tend to change over time. However, emerging demands for empowering end-users to set up policies are often unrealistic. This in turn may result in new sources of vulnerability and high levels of user frustration, if not outright distrust or even fear of pervasive computing technologies.

We believe it is important to develop new user interfaces to support lay users in understanding and managing security and privacy policies – their own as

well as those implemented by systems and individuals with whom they interact. Previous solutions have taken a narrow view, e.g. limiting the expressiveness of policy languages, or restricting some decisions to specific roles within the enterprise. As systems grow more pervasive and more complex, and as demands for increasing flexibility and delegation continue to grow, we argue it is imperative to take a more fundamental view that weaves together issues of security, privacy, and usability.

In this paper, we report on our initial work in designing and evaluating novel mechanisms for managing security and privacy in pervasive computing environments. Our research combines the development of new user interfaces with learning, dialog, and explanation functionality to empower users. We describe our current work with respect to three pervasive computing scenarios, and then draw out themes that we have learned thus far. Our three applications are:

1. **Contextual Instant Messaging:** Users can inquire about each other's context (e.g. interruptability, location and current task) through an instant messaging service
2. **People Finder Application:** Users are equipped with location-aware smartphones. They interact with their devices to inquire about the locations of others subject to privacy policies.
3. **Access Control to Resources:** Smartphones are used to access both physical and digital resources. Users can use their smartphones to create and manage their *security policies*, and to give others credentials to access different resources.

A fundamental challenge is capturing users' policies without being burdensome. One strand that connects these scenarios together is understanding how to balance the tradeoff between expressiveness and simplicity. For example, when creating policies about disclosing one's location, are current location, time, and requester's name sufficient, or do other factors such as relationship with the requester and requester's

location need to be taken into account? A tradeoff also exists between the frequency and timing of user prompts, and the tolerance users have for the system making incorrect decisions. A second strand that connects these scenarios is conveying to users what the capabilities of the system are, what policies are currently in effect, and the consequences of a policy change. This includes letting users author policies, audit the results of decisions made based on these policies, and ask simple questions such as “Why was the John allowed to enter my office” and “Why couldn’t my boss access the quarterly report?”

2. Ongoing Research and Preliminary Results

2.1. Contextual Instant Messaging

We have iteratively designed privacy controls and feedback mechanisms for *imbuddy411*, a contextual IM service that lets any AOL Instant Messenger (AIM) users query for three types of information: interruptibility, location, and current task (abstractly represented as the name of the current window being viewed). Currently, AIM users can only query information of AIM users who are running our client software, which collects and reports their contextual information.

We decided to use a group-based approach to configure the contextual IM privacy settings, based on prior lab studies by Patil and Lai [6]. Users can modify their privacy control setting via a web browser (see Figure 1a). All buddies are first classified under a ‘default’ privacy group that denies all disclosures.

Users can create as many groups as they want and move buddies from the default group to any of other group. Other AIM users who request information from *imbuddy411* but are not part of the user’s buddylist are dynamically added to the default group.

We also developed three feedback mechanisms: a notification letting users know when their information is being seen (Figure 1b), a grounding and social translucency mechanism that facilitates conversation by letting users know what others know about them (Figure 1c), and a history letting users know what information has been disclosed to others (Figure 1d).

We implemented *imbuddy411* as an AIM robot that could answer queries, such as “howbusyis alice” and “whereis bob”, and a Trillian plug-in that can sense contextual information such as interruptibility (using the SUBTLE toolkit [3]), location (using PlaceLab [4]), and current task. To introduce *imbuddy411* to our participants’ buddies, a short blurb was included in each participant’s profile. Our Trillian plug-in also advertised the *imbuddy411* service whenever a conversation starts between a user and their buddies.

We conducted a two week study with ten IM users. There were 193 queries not counting users querying themselves, including 54 interruptibility requests, 77 location requests, and 62 active window requests. Also, 63 queries were hits to the database (i.e. when users were not online). There were 46 distinct users who queried *imbuddy411* and 9 of those were repeat users.

More importantly, although all our participants agreed that the three information types being disclosed were all potentially sensitive (interruptibility: 3.6, location: 4.1, active window: 4.9, all out of 5), our

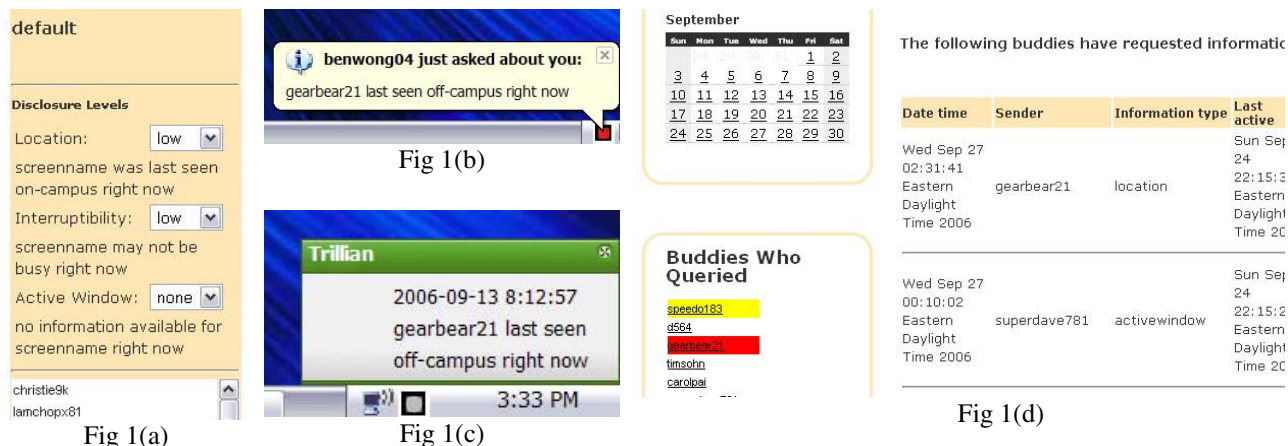


Figure 1. These screenshots show the control and feedback mechanisms of *imbuddy411*. The configuration user interface (a) shows what information will be disclosed by default. A notification (b) lets users know when someone is requesting information. A grounding and social translucency mechanism (c) lets a user know what the other person knows, and is shown at the start of a conversation. A disclosure history (d) lets people audit disclosures.

participants said they were comfortable with their privacy settings (4.1 / 5). We found this result particularly interesting, since as part of our experiment, we would occasionally use alternative screen names to make requests for personal information (i.e. fake probes). However, most of our participants' settings were set up to not reveal anything by default, and so they were unconcerned and did not mention this issue at a debriefing at the end of the study.

2.2 People Finder Application

The emergence of cell-phone-based location tracking opens the door to a number of new applications, including recommendations, navigation, safety, enterprise applications, and social applications.

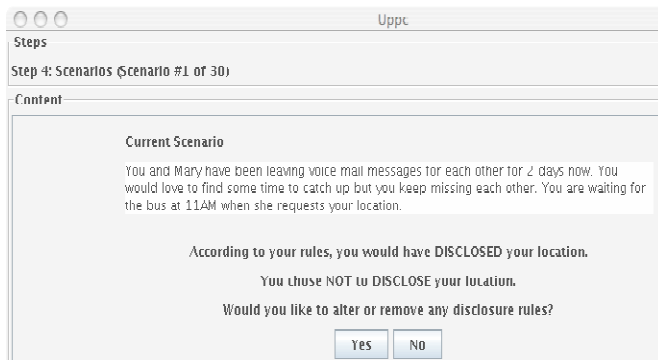



Fig 2(a)

Experiments conducted with some of these applications in the context of MyCampus show that adoption of these services often depends on whether users feel they can adequately control when their location is shared (e.g. [7]). To better understand the privacy preferences users have in the context of these applications, as well as what it takes to capture these preferences, our group is conducting a series of experiments involving a cell-phone-based people finder that lets users inquire about the location of their friends, family members, and colleagues.

In a first set of experiments, 19 participants were presented with situations simulating queries from others. The queries were customized to capture elements of their daily activities involving friends,



	Mean (sec)	Standard Deviation (sec)
Rule Creation	326.53	175.41
Rule Maintenance	143.0	132.74
Total	469.53	223.14

Fig 2(b)

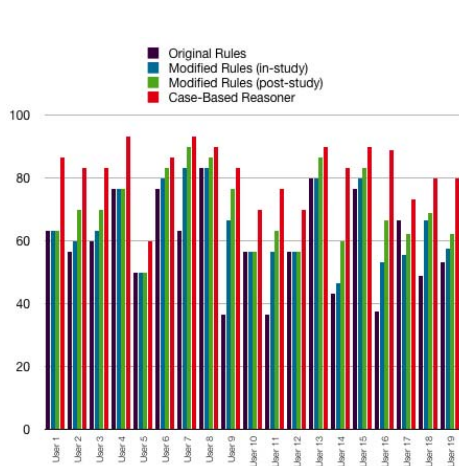


Fig 2(c)




Fig 2(d)

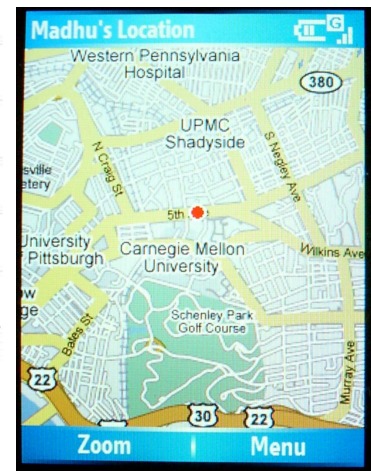


Fig 2(e)

Figure 2. People finder application: (a) screenshot of individualized scenario in simulation experiment, (b) Time spent by users creating and refining their privacy policies, (c) accuracy of policies captured through rule editing and learning measured as percentage of scenarios on which they make the right decision for each subject, (d) average accuracy for all 19 subjects, (e) screenshot of cell phone-based application.

family, and colleagues. Each participant was asked to specify rules indicating the conditions under which she would be willing to share her location information with others (e.g. “My colleagues can only see my location on weekdays and only between 8am and 6pm). The experiments involved presenting each participant with a total of 30 individualized scenarios (45 scenarios for each of the last 4 participants). Each individualized scenario included asking the participant whether she felt comfortable disclosing her location, showing her what her current policies would do, and offering her a chance to refine her policies – Figure 2(a).

Our experiments show that users often have fairly sophisticated privacy preferences, requiring over 5 minutes just to specify their initial rules and nearly 8 minutes if one adds time spent revising these rules as they get confronted with new situations. Several users ended up with 8 or more rules by the end of the experiments. More surprisingly, despite the time and effort spent specifying and refining their policies, participants were generally unable to achieve high levels of accuracy. Rules specified at the beginning of the experiments only captured their policies 59% of the time (Fig 2c). When given a chance to revise their rules over time, that percentage only went up to 65%. Even when using the rules that users ended up with at the end of the experiments and re-running these rules on all 30 (or 45) scenarios, decisions were only correct 70% of the time.

We are experimenting with machine learning to see if we can do better. Our results with case-based reasoning (CBR) suggest it is possible to train a system to learn a user’s policies that can be more accurate than those specified by users (Fig 2c and 2d) – 82% accuracy using CBR. While additional experiments are required to validate statistical significance, these preliminary findings suggest that requiring users to fully specify their policies may be unrealistic. Instead, learning as well as dialog and explanation technologies seem to have the potential of offering solutions that better capture user policies while also reducing user burden. At the time of writing, our group is finalizing steps to conduct another round of experiments with participants inquiring about each other’s location using actual cell phones in their daily routines (Fig 2e).

2.3 Access Control to Rooms in an Office Building

We have deployed a distributed, smartphone-based access-control system called Grey in a building on our campus [1, 2]. Grey can be used to control access to physical resources such as office doors, as well as electronic resources such as computer accounts or

electronic files. Grey-enabled resources allow access when an individual’s smartphone presents a proof that access is permitted. Proofs are assembled from a set of credentials that express authority. The credentials are created and managed by end-users on their Grey phones. Instead of relying on a central access-control list, in Grey end-users are empowered to create flexible access-control policies for the resources they manage.

Grey users can delegate their authority *proactively* by manually creating credentials that let a user or group of users access a specified resource during a specified time period. Grey users can also create credentials *reactively*, when another user asks for access. In this case, the user who may have the needed credentials is prompted to help the user who is trying to gain access. If she decides to help, Grey will forward the relevant credentials from her phone to the user trying to gain access or, if such credentials don’t yet exist, intelligently prompt her to first create such credentials, e.g. by adding the requestor to a group that already has access to the resource.

We have outfitted over three dozen doors in our building with Grey-enabled Bluetooth door locks and given smartphones with Grey software to 19 users. Grey is also used by nine members of the Grey project team. We have monitored Grey usage by collecting log files from phones and doors and by interviewing Grey users every four to eight weeks over a period of several months.

Our office building includes a shared workspace with open cubicles, as well as conference rooms, labs, storage closets, and offices. Locked perimeter doors secure the entire workspace in the evening and on weekends. Conference rooms, labs, storage closets, and offices can be individually locked. All Grey users were given credentials to unlock the perimeter doors, and users with offices were given credentials to unlock their own office doors. Some Grey users were also given additional credentials, e.g. to unlock a lab or a storage closet. A user accesses a resource (e.g. a door or a computer login) by selecting its name from the phone’s menu, after which the phone and the resource communicate via Bluetooth. The resource grants access (e.g. the door unlocks) when it has verified the credentials and proof submitted by the phone. If a user does not have credentials to access the resource, her phone prompts her to ask another Grey user to delegate the necessary authority.

We have learned a number of lessons from our initial deployment of Grey, many of which may be broadly applicable to other mobile-device applications and access-control technologies.

- We found a variety of obstacles to acceptance of Grey, including user perception that Grey was slow (even when it wasn't) and system failures that caused users to get locked out. While security usually focuses on keeping unauthorized users out, our users were more concerned with how easy it was for them to get in, and in interviews never mentioned security concerns
- We were hoping to observe frequent delegation, but since Grey relies on network effects, the small number of users and resources limited opportunities for delegation. We are investigating better ways to bootstrap so that Grey will be more useful, even for a small population.
- One of our objectives of this trial deployment was to study the types of access-control policies users would create when no longer constrained by the limitations imposed by difficult-to-obtain physical keys. We observed users creating policies that did not mirror the policies they had with physical keys, and we found that the low overhead for creating and changing policies with Grey encourages policy change and the creation of policies that better fit the users' needs.
- Finally, we were surprised at some of the unanticipated uses our users made of the Grey system. For example, some of our users routinely use Grey to unlock doors without having to get out of their chairs. We probably would not have discovered this usage without a field study.

3. Concluding Remarks

In this article, we have given an overview of our work in developing some core technologies for helping end-users manage their security and privacy. When done, we hope to have simple user interfaces and visualizations for specifying and understanding policies, better dialog and explanation technologies that can give details as to why a given policy did or did not do a desired action, and reliable machine learning to help shoulder the burden of managing these policies. We have also described three applications we are developing to evaluate these technologies.

One challenge in researching security and privacy is that people rarely talk about it in those terms, instead focusing on the value proposition of applications. For example, with Grey, our users were more concerned with how easy it was for them to get in than in the security issues, despite the fact that most Grey users are security researchers. Finding a meaningful value proposition has also been a difficulty with people finder. Some of our early participants could not foresee a situation where they would use our system, saying that they would just call the person instead. An upfront

value proposition is essential to drive adoption of applications, so that there are enough users and enough usage so that we can study the related security and privacy issues.

Another related challenge is in getting a large user base. For example, the more people there are that use contextual instant messaging and people finder, the greater the potential utility (good for users) and the more potential security and privacy risks there are (good for our research). Growing a large user base can be difficult, though, for applications that require special hardware, as in people finder and Grey. To address this problem, our strategy for the next iteration of our applications is to tie them to existing applications and resources. One thrust of our current work is to integrate contextual instant messaging and people finder, making it so that people can use laptops and smartphones to query each other's information.

Thus far, however, our experiments suggest that users can find some value in our applications and do not necessarily object to the privacy or security implications. It seems that, if given adequate control over the situations when information is shared with others or when access to resources is granted, they will adopt these solutions and sometimes come up with unexpected ways of using them (e.g. remotely unlocking doors with Grey).

4. Acknowledgements

This work is supported by NSF Cyber Trust grant CNS-0627513, NSF grant CNS-0433540, and ARO research grant DAAD19-02-1-0389 to Carnegie Mellon University's CyLab.

5. References

- [1] Bauer, L., Cranor, L. F., Reiter, M. K., and Vaniea, K. Lessons Learned from the Deployment of a Smartphone-Based Access-Control System. Technical Report CMU-CyLab-06-016, CyLab, Carnegie Mellon University, October 2006.
- [2] Bauer, L., Garriss, S., McCune, J. M., Reiter, M. K., Rouse, J., Rutenbar, P. Device-enabled authorization in the Grey system. In *Proceedings of the 8th Information Security Conference*, Singapore, September 2005.
- [3] Fogarty, J. and Hudson, S.E. (2007) Toolkit Support for Developing and Deploying Sensor-Based Statistical Models of Human Situations. Submitted for Review, *CHI 2007*.
- [4] LaMarca, A., et al. (2005) Place Lab: Device Positioning Using Radio Beacons in the Wild. In *Proc of Pervasive 2005*.
- [5] Maxion, R. A. and Reeder, R. W. Improving user-interface dependability through mitigation of human error.

International Journal of Human-Computer Studies 63(1-2),
HCI research in privacy and security, pp. 25-50, 2005.

[6] Patil, S. and Lai, J. (2005) Who gets to know what when:
configuring privacy permissions in an awareness application.
In *Proc of CHI 2005*.

[7] Sadeh, N., Gandon, F., and Kwon, O. B. "[Ambient Intelligence: The MyCampus Experience](#)", Chapter in
"Ambient Intelligence and Pervasive Computing", Eds. T.
Vasilakos and W. Pedrycz, ArTech House, 2006.