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 – Pittsburgh, PA - USA
Email contact: 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 services they support are increasingly becoming central in both personal and business life. The dramatic market growth of smartphones, portable storage devices, and other mobile 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 a difficult problem. Even in desktop computing environments, end-users have great difficulty using the Windows XP file permission system to create security policies for file access [MR05]. 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 traditionally 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 interface technologies 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 scenarios 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 smartphones that track their location. They interact with their devices to inquire about the locations of others (e.g., colleagues, friends, spouses) subject to privacy policies.
3. **Access Control to Resources:** Smartphones act as the token by which users access both physical and digital resources. Users interact with their smartphones to create and manage their *security policies*, and (via the smartphones) with each other to obtain credentials to access different resources.

A fundamental challenge here 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 for making a decision, or should policies take into account other factors such as relationship with the requester or requester's location? 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 supporting functionality that lets 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", "Why couldn't my boss access the quarterly report", and "What if all people in my department could see my location?"

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.

For control mechanisms, we decided to use a group-based approach to configure the contextual IM privacy settings, based on prior lab studies by Patil and Lai [PL05]. 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 (see Figure 1b), a grounding and social translucency mechanism that facilitates conversation by letting users know what others know about them (see Figure 1c), and a disclosure history letting users know whether the right information has been disclosed to the right person (see Figure 1d).

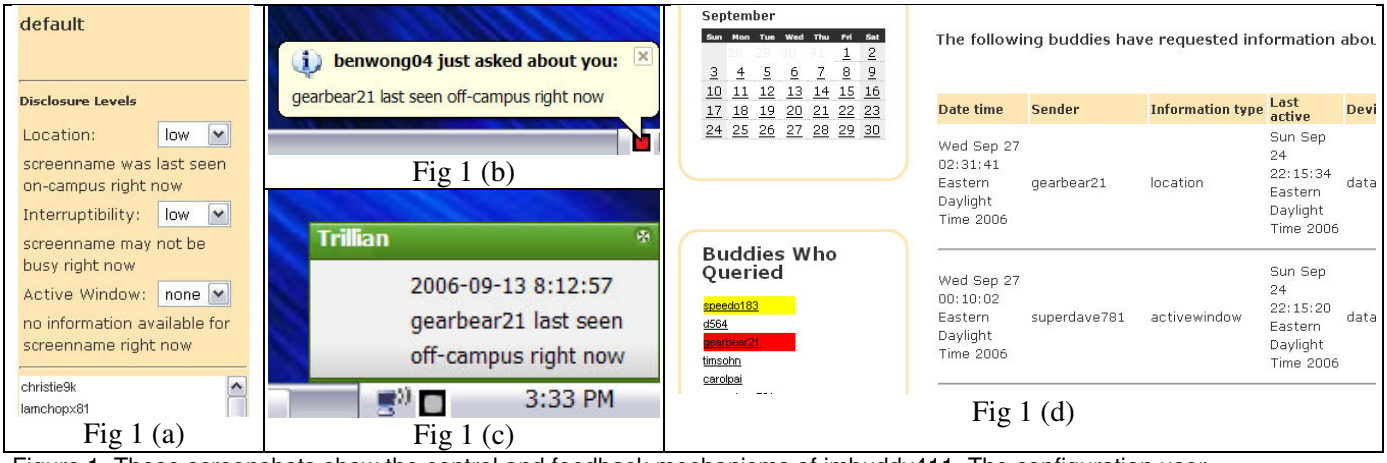

Figure 1. These screenshots show the control and feedback mechanisms of imbuddy411. The configuration user interface (a) shows what information will be disclosed by default. A notification (b) lets users known when someone is requesting information. A grounding and social translucency mechanism (c) lets 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.

Our *imbuddy411* service was implemented using 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 [FH07]), location (using PlaceLab [LCC+05]), 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 field study with ten IM users. There were 193 individual instances of use (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 screennames 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 participants said they were comfortable with their privacy settings for *imbuddy411* (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 (e.g. tracking your children), enterprise applications, and social applications. Experiments conducted with some of these applications in the context of the MyCampus environment show that adoption of these services often depends on whether users feel they can adequately control when their location information is shared (e.g. [SGK06]). 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, a total of 19 participants were presented with situations simulating queries from others. The queries were customized to capture elements of their daily activities involving their friends, colleagues, and family members. Each participant was asked to specify her privacy preferences in the form of 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, even for a seemingly simple application like people finder, 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 they had 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 rerunning these rules on all 30 (or 45) scenarios, decisions were only correct 70% of the time.

We are experimenting with learning technologies to see if we can do better. In particular, our results with case-based reasoning suggest that it is possible to train a system to learn a user's policies that can be more accurate than those directly specified by the users (Fig 2c and 2d) – 82% accuracy using case-based reasoning. While additional experiments are required to validate the statistical significance of these results, these preliminary findings seem to indicate that requiring users to fully specify their policies on their own 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).

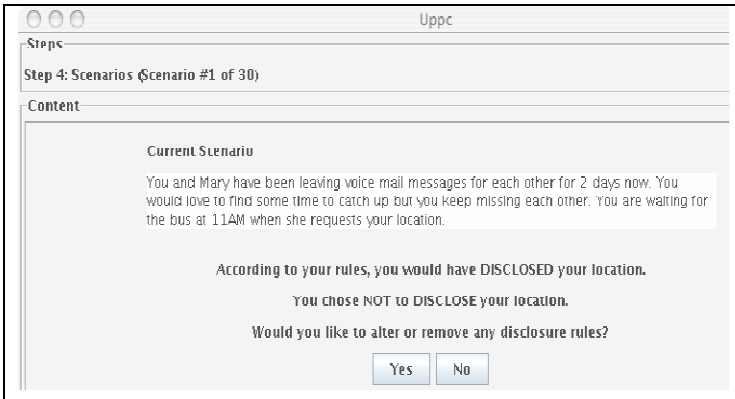

Fig 2(a)

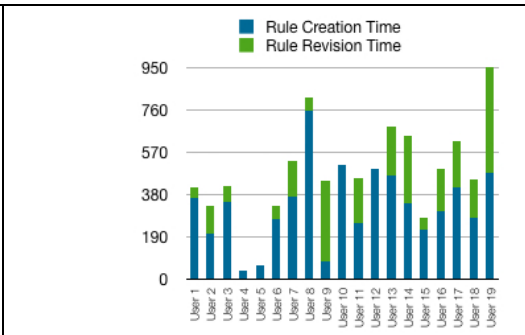

	Mean (sec)	Standard Deviation (sec)
Rule Creation	326.53	175.41
Rule Maintenance	143.0	132.74
Total	469.53	223.14

Fig 2(b)

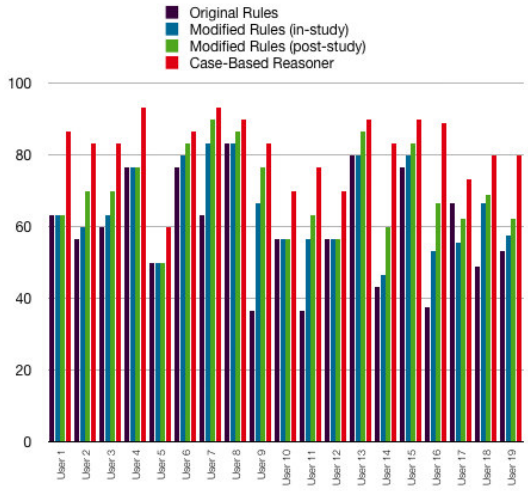

Fig 2(c)

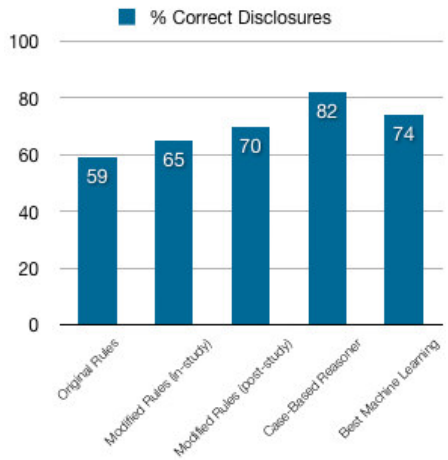

Fig 2(d)

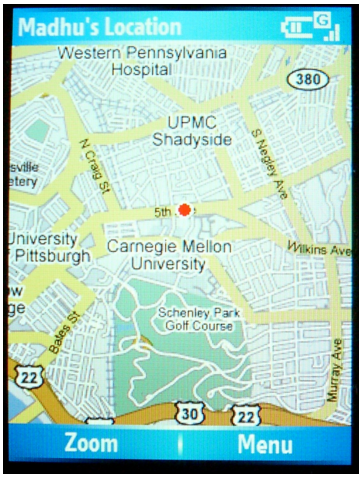

Fig 2(e)

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.

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 [BGM+05, BCR+06]. 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 can be accessed when an individual's smart phone presents a proof that access is permitted. Proofs are assembled from a set of credentials that Grey users can delegate to each other using their phones. Thus, there is no central access control list. Instead, end-users are empowered to create flexible access control policies for the resources they manage.

Grey users can delegate their credentials *proactively* by manually creating access credentials that let a particular 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. The user who has the needed credentials is prompted and given the option of helping the other user or denying the request. If she decides to help, she may have a number of options including generating an access credential for the other user or adding the other user to a group that already has access to the desired resource. If she decides to create an access credential she must choose the length of time for which the credential will be valid.

We have outfitted over three dozen doors in our building with Grey-enabled Bluetooth door locks and gave smartphones with Grey software installed to 16 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. Rest rooms, elevators, and other offices are located outside the perimeter doors. All Grey users were given credentials to unlock the perimeter doors. Grey 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. Users can select the name of the resource they want to access on their phone's menu and their phone will attempt to establish a Bluetooth connection with that resource. Once a connection is made, the phone sends the credentials to the door, the door verifies the proof, and if everything is in order it opens. If a user does not have credentials to access the resource, their phone prompts them to ask another Grey user to delegate the necessary credentials to them.

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 the system was slow and system failures that caused users to get locked out at inconvenient times. 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 never mentioned security concerns when we interviewed them.
- We were hoping to be able to observe frequent delegation, but since Grey relies on network effects, we found that 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.
- The small display on our smartphones limited the amount of information we could present. We initially tried to use abbreviations and let menu items to scroll off the screen. However, we discovered that users ignored some menu items because they did not understand our notation. We developed a new "wizard" interface for proactive delegations that breaks the task down into several steps that can be displayed fully on the small screen. Although this approach requires more steps, our users seem to prefer it because it is more understandable and proactive delegation is a relatively infrequent activity.
- 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 exactly mirror the policies they had created when they distributed physical keys, and we found that the low overhead for creating and changing policies with Grey encourages policy change.
- 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

Developing functionality that lets users effectively specify security and privacy policies is not an easy proposition, as illustrated by the work reported in this article. Initial solutions often need to be iteratively refined a number of times before they become effective. Our experiments suggest, however, that users can find value in new applications such as the ones we have researched and do not necessarily object to their new 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).

What it takes to achieve this level of comfort seems to vary, however. For *imbuddy411*, which involves the selective disclosure of contextual information among IM buddies, a clean interface design with a relatively small number of options seems to do the trick – at least based on our initial experiments. We believe that the ability of specifying conservative defaults probably plays an important role here. The

informal nature of IM as well as the peer nature of the group of users involved in the *imbuddy411* study might also have a played a role in the higher level of acceptance. Users did not seem to have any preconceived expectations about what information they should be allowed to see about others.

Expectations could explain the difference observed with users of our people finder, who seemed to have a much more difficult time articulating privacy policies they felt comfortable with. In follow-up interviews, participants indicated that their decisions about disclosing location included thinking about the reaction other users would have if denied (e.g hiding from your parents, boss, or spouse). A variation of the people finder that obfuscates the reason why a particular request is denied, or embodies notions of plausible deniability, might significantly help and allow for simpler privacy rules.

Our work also strongly suggests that users have a hard time articulating policies, especially as these policies become slightly complex (e.g. more than 5 rules in people finder). In these cases, simple policy editors are not sufficient. Instead, learning, dialog, and explanation functionality offer the prospect of richer interactions with user, leading to both higher fidelity policies and lower user burden. Ultimately, this should also include the development of auditing tools that enable users to analyze past situations and ask questions such as *"Why did you do X?"* (where *X* might be allowing my boss to access my location) or *"Why didn't you do X?"* (where *X* might be allowing your boss to access an important project document).

4. Acknowledgements

This work is supported by NSF Cyber Trust Grant CNS-0627513 and in part under ARO research grant DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab.

5. References

[BCR+06] L. Bauer, L. F. Cranor, M. K. Reiter, and K. Vaniea. Lessons Learned from the Deployment of a Smartphone-Based Access-Control System. Technical Report CMU-CyLab-06-016, CyLab, Carnegie Mellon University, October 2006.

[BKK+05] Brodie, C., Karat, C., Karat, J., and Feng, J. Usable security and privacy: a case study of developing privacy management tools. In *Proceedings of the 2005 Symposium on Usable Privacy and Security* (SOUPS '05). pp. 35-43. 2005

[BGM+05] 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.

[CGA06] Cranor, L., Guduru, P., and Arjula, M. User Interfaces for Privacy Agents. To appear in ACM Transactions on Computer-Human Interaction, 2006. http://lorrie.cranor.org/pubs/privacy-bird-20050714.pdf

[FH07] 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*

[KBK06] Karat, C., Brodie, C., and Karat, J. Usable privacy and security for personal information management. In *Commun. ACM* 49, 1 (Jan. 2006), pp. 56-57. 2006.

[LCC+05] LaMarca, A., et al. (2005) Place Lab: Device Positioning Using Radio Beacons in the Wild. In *Proc of Pervasive 2005*.

[MR05] Maxion, R. A. and Reeder, R. W. Improving user-interface dependability through mitigation of human error, International Journal of Human-Computer Studies, Volume 63, Issues 1-2, HCI research in privacy and security, July 2005, Pages 25-50.

[PL05] Patil, S. and Lai, J. (2005) Who gets to know what when: configuring privacy permissions in an awareness application. In *Proc of CHI 2005*.

[SGK06] 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.