Madhu Prabaker · Jinghai Rao · Ian Fette · Patrick Kelley · Lorrie Cranor · Jason Hong

Norman Sadeh*

Understanding and Capturing People's Privacy Policies in a People Finder Application

Received: date / Accepted: date

Abstract Over the past few years, a number of mobile applications have emerged that allow users to locate one another. Some of these applications are driven by a desire from enterprises to increase the productivity of their employees. Others are geared towards supporting social networking scenarios or security-oriented scenarios. The growing number of cell phones sold with location tracking technologies such as GPS or A-GPS along with the emergence of WiFi-based location tracking solutions could lead to mainstream adoption of some of these applications. At the same time, however, a number of people have expressed concerns about the privacy implications associated with this class of software, suggesting that broad adoption may only happen to the extent that these concerns are adequately addressed.

In this article, we report on work conducted at Carnegie Mellon University in the context of PEOPLEFINDER, an application that enables cell phone and laptop users to selectively share their locations with others (e.g. friends, family, and colleagues). The objective of our work has been to better understand people's attitudes and behaviors towards privacy as they interact with such an application, and to explore technologies that empower users to more effectively and efficiently specify their privacy preferences (or "policies").

1. Introduction

Over the past few years, a number of mobile applications have emerged that allow users to locate one another. Some of these applications are driven by a desire from enterprises to increase the productivity of their employees. Others are geared towards supporting social networking scenarios or security-oriented scenarios. The growing number of cell phones sold with location tracking technologies such as GPS or Assisted GPS ("A-GPS") along with the emergence of WiFi-based location tracking solutions could lead to mainstream adoption of some of these applications.

In this article, we report on work conducted at Carnegie Mellon University in the context of PEOPLEFINDER, an application that enables cell phone and laptop users to selectively share their locations with others (e.g. friends, family, and colleagues). This article extends a previous workshop paper in which we introduced PEOPLEFINDER [1], and provides a more thorough and detailed report.

Our objective has been to better understand people's attitudes and behaviors towards privacy as they interact with such an application, and to explore technologies that empower users to more effectively and efficiently specify their privacy preferences (or "policies").

The work presented in this article confirms that people are generally apprehensive about the privacy implications associated with location tracking. It also shows that privacy preferences tend to be complex and depend on a variety of contextual attributes (e.g. relationship with requester, time of the day, where they are located). Through a series of user studies, we have found that most users are not good at articulating these preferences. The accuracy of the policies they define increases only marginally over time unless they are given tools that help them better understand how their policies behave in practice.

Overall our studies, which included a combination of controlled lab experiments with 19 users and field studies involving a total of over 60 participants, suggest that functionality that increases user awareness can contribute to the definition of more accurate policies. In our field studies, as users grew more comfortable with PEOPLEFINDER and the way in which it was used by their acquaintances, they started

* Primary contact point (sadeh@cs.cmu.edu)

refining their preferences and relaxing some of their policies to allow for requests that would have been denied under their initial policies. Overall, these results suggest that functionality that empowers users to more effectively control their policies can contribute to the adoption of context-aware applications like PEOPLEFINDER.

This article also compares results obtained in the context of controlled lab studies with results from longitudinal studies spanning up to several weeks. While both types of studies show that users have a hard time defining policies, our results suggest that users tend to be significantly more careful when defining policies that will be used to make decisions in actual situations (rather than under simulated conditions). To the best of our knowledge, the results from our field studies are the first of this type to analyze the behavior of users and their policies in the context of a fully deployed application with actual users.

The remainder of this article is organized as follows. Section 2 provides an overview of our PEOPLEFINDER application. Section 3 discusses the privacy policy authoring functionality we have developed as well as several enhancements we are currently working on. An overview of PEOPLEFINDER's auditing functionality is provided in Section 4. Section 5 provides a summary of a first set of lab experiments we conducted in the Summer of 2006. Results and observations from a series of three pilots involving over 60 participants in the Spring of 2007 are presented in Section 6. Section 7 contains some concluding remarks and discusses future work.

2. Overview of PEOPLEFINDER

In PEOPLEFINDER, users rely on Policy Enforcing Agents (PEA) to handle queries about their locations. The user's PEA operates according to a policy, or set of rules, specified by the user, with each rule granting access to the user's location under a particular set of conditions (e.g. query coming from a particular group of users on one of several possible days and within one of several possible time windows).

Users can invite other people (e.g. friends, family members, or colleagues) to check their location with PEOPLEFINDER, using either a mobile phone client or the PEOPLEFINDER web site. Users can specify rules under which other people can access their location and define groups of people to which particular rules apply.

PEOPLEFINDER is available for cell phones and for laptops. The cell phone version relies on GPS technology to pinpoint the user's location. When no GPS reading is available (e.g. the user is indoors), the application falls back on a GSM triangulation solution developed by Intel Research Seattle [3]. While the GSM approach is not as accurate as GPS, it provides an estimate of the user's location (often within a few hundred yards) under a significantly wider set of conditions.

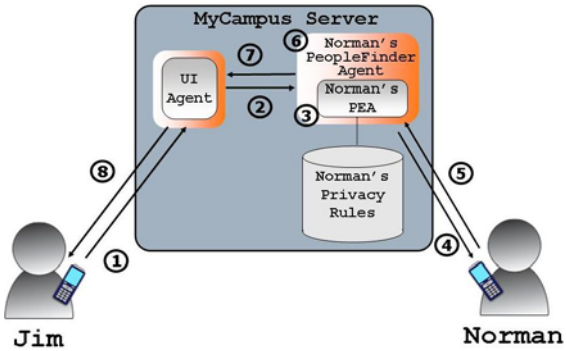

Fig. 1 Processing Jim's request for Norman's location.

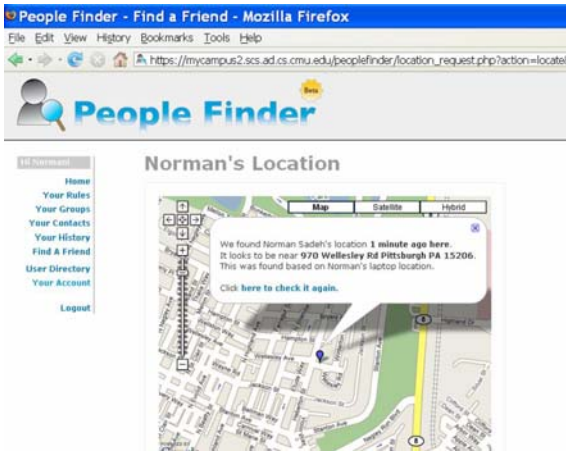

Fig. 2 The results of a location query displayed in a web browser.

The laptop version uses a WiFi positioning solution developed by Skyhook Wireless [5]. In urban areas, this solution tends to have an accuracy of about 30 yards. It is complemented by an ad-hoc WiFi-based solution developed specifically for Carnegie Mellon's campus. This latter solution, which uses a database of access points on campus, often provides readings that are even more accurate than the more general Skyhook Wireless solution.

We distinguish between *target users*, namely PEOPLEFINDER users who are willing to share their locations with others, and *requesting users*, namely users who can submit queries about the location of one or more target users. A user can be both a target user and a requesting user but does not have to be. Target users who rely on their laptops to track their location need to download a C# application on their laptops. J2ME and C# versions of the application have also been developed for target users who rely on their cell phones to track their location, though these versions only work on a

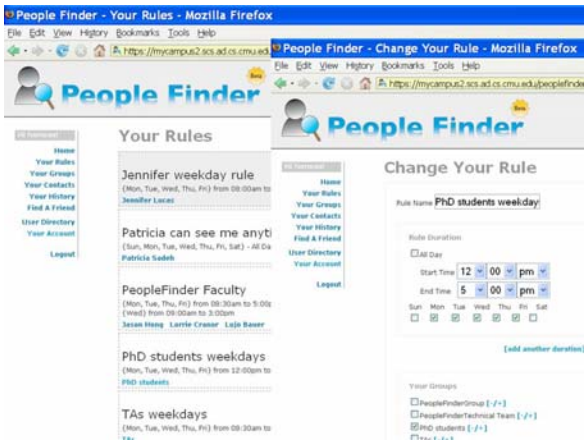

Fig. 3 User interface for defining simple privacy rules.

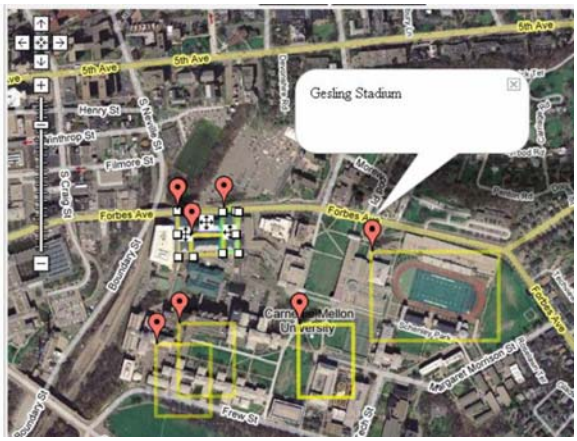

Fig. 4 Defining locations as combinations of rectangular areas for use in location-sensitive privacy rules.

Fig 5. Bubbles notifying users of incoming queries help maintain awareness while being minimally disruptive.

limited number of smartphone models. The smartphone version also lets users query for other people's locations.

Figure 1 outlines the main steps involved in processing a query from a user, say Jim, for the location of a target user, say Norman. The request submitted by Jim is forwarded by his User Interface (UI) Agent (e.g. Web browser or cell-phone application) to Norman's PEOPLEFINDER Agent. The agent invokes Norman's Policy Enforcing Agent (PEA) to check whether the query is consistent with the privacy rules specified in his policy. If it is, the request is forwarded to Norman's location tracking device, a cell phone in this example. Once returned, the location may need to be further processed by Norman's PEOPLEFINDER Agent (e.g. to combine multiple readings of Norman's location such as a GPS reading from a few minutes ago and a more recent reading based on GSM triangulation) before being forwarded to Jim. Finally, the results of the request are displayed on Jim's client, as shown in Figure 2.

In general, processing may be somewhat more complex and some privacy rules may in fact require checking Norman's location to determine whether or not to disclose his location. For instance, Norman may have specified that his colleagues can only access his location during weekdays and while he is on campus. Query processing could also involve the use of obfuscation rules that manipulate the accuracy of the response returned to a user [2].

PEOPLEFINDER is built on top of the MyCampus infrastructure, a semantic web environment in which policies are expressed using a rule extension of the OWL language [2]. The resulting language is capable of modeling a wide range of policies. Access to a user's location can be restricted according to conditions that refer to any number of concepts or instances of concepts defined in an open collection of ontologies (e.g. ontologies of locations, social relationships, and calendar activities). This includes capturing a variety of context-sensitive restrictions such as disclosing your location only when you are in a particular place, or enforcing obfuscation policies that allow users to specify how they want the application to manipulate the accuracy of their location before disclosing it (e.g. city-level versus street address).

Presently, PEOPLEFINDER only uses a small fraction of the policies that can be expressed in this framework. In fact, one of the questions our project is attempting to address has to do with how much expressiveness is actually required for users to feel comfortable using the application and to what extent adding more expressiveness enables users to more accurately specify their policies – in contrast to creating more confusion.

3. Privacy Policy Authoring

Users can define rules in which they grant access to their locations to individuals or groups of users. Each rule includes one or more restrictions such as the day(s) of the week or time(s) of day during which location queries from particular

individuals or groups of users will be granted, as shown in Figure 3. Users can belong to multiple groups.

Extensions of the rule interface also allow users to specify locations as collections of rectangles on a map (e.g. all buildings in the School of Computer Science) and specify rules that include location-based restrictions (e.g. only disclose my location when I am in a School of Computer Science building), as shown in Figure 4.

To avoid conflicts in rules, we currently only allow positive assertions. For example, a person can specify "Mary can see my location between 9AM and 5PM", but cannot specify, for example, "Colleagues can not see my location on weekends".

4. Auditing Functionality

The experiments reported in Sections 5 and 6 show that users often have difficulty anticipating how people they invite will use the application. To be effective, user interfaces have to be designed to increase user understanding of how the application is being used. We have found that simple bubbles that discreetly pop up (e.g. at the bottom of a laptop screen) to notify users that their location is being requested can go a long way in helping users feel more comfortable with the application (see Figure 5). This finding was also validated in imbuddy411 [4], a sister project of PEOPLEFINDER.

An even more important element is the design of auditing functionality that enables users to review requests that have been submitted, see how they were processed by the rules they currently have in place, and possibly request more detailed explanation to identify rules they may want to modify.

In PEOPLEFINDER, users have a number of options to audit previously submitted requests. This includes reviewing requests that were denied or requests that have not yet been audited, as shown in Figure 6. They can incrementally access additional details about a particular request, such as where they were when their location was requested or the way in which their location was estimated (e.g. GPS versus GSM), as shown in Figure 7.

The interface also supports *explanation functionality*. As Figure 7 illustrates, the system identifies for users what rules led to a particular disclosure/non-disclosure decision. By letting users indicate whether they are satisfied with the decision made based on their current policy, the system can try to help users refine their policies. Sections 5 and 6 present results obtained by running different learning algorithms on the feedback obtained from users to help refine their policies. The same type of feedback could also be used to initiate dialogues and offer suggestions on how they could improve the accuracy of their rules. Functionality aimed at doing this is currently under development.

Fig. 6 Auditing functionality helps users understand how their policies work and enables them to more effectively refine their policies.

Fig. 7 Explanation can help users better understand their policies. User feedback can also be used to make suggestions or learn the user's preferences.

5. Initial Lab Experiments

Our current version of PEOPLEFINDER reflects several design iterations with users. Initial work was conducted using a mockup application designed to present users with scenarios that captured elements of their daily routines and interactions with members of their social networks. In this section, we briefly summarize findings from this initial work, which revolved around lab experiments involving 19 participants. In Section 6, we present more recent results from 3 pilot studies conducted with users of a deployed version of PEOPLEFINDER. This second set of experiments involved a total of over 60 participants. We discuss how results from the latter studies reinforce most of our initial findings and also point to a few differences between these two sets of experiments.

In our laboratory experiments, users were asked to provide information about their daily routines and social networks (e.g. names of key family members, boyfriend/girlfriend/spouse, colleagues/classmates, and friends). 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

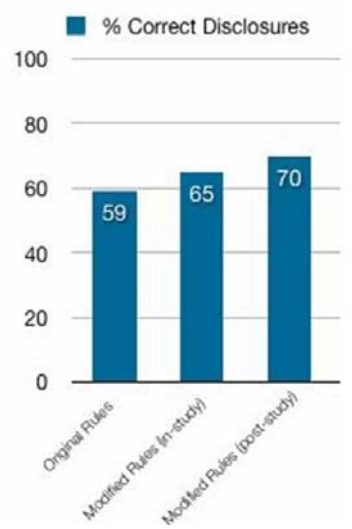

Fig. 8 Controlled lab experiments: Users are not very good at articulating their privacy policies – accuracy of initial rules versus rules modified after being presented with 30 customized usage scenarios.

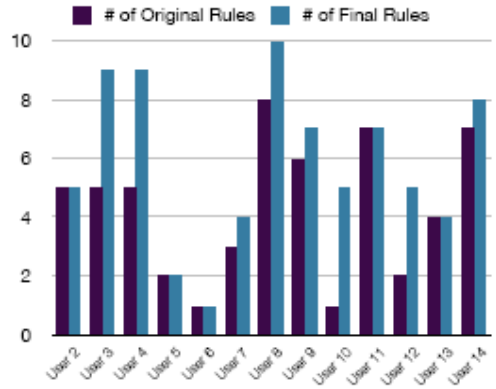

Fig. 9a Controlled lab experiments: initial number of rules versus final number of rules.

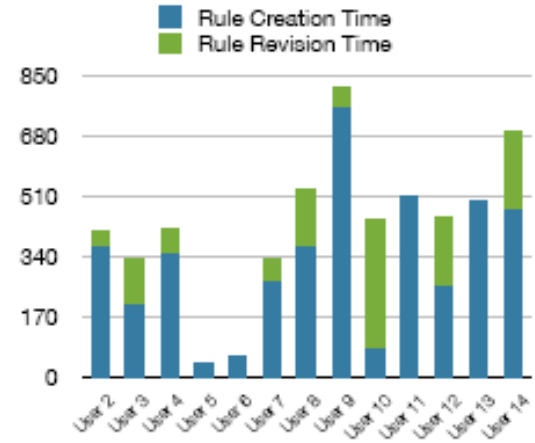

Fig. 9b Controlled lab experiments: time spent creating and modifying rules – the latter includes both changes to initial rules and addition of new rules

her current policies would do, and offering her a chance to refine her policies.

On average, subjects required a little over 5 minutes to specify their initial rules and nearly 8 minutes if one includes the time spent refining their rules as they were confronted with new situations. Several users ended up with 8 or more rules by the end of the experiments. Despite the time and effort spent specifying and refining their policies, participants were generally unable to achieve high levels of accuracy. Based on feedback provided as they were presented with individualized scenarios, subjects indicated they were only satisfied with 59% of the decisions made by their initial rules, as shown in Figure 8. As they refined 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.

During the course of the experiments, most users refined their existing policies and also added new ones, as shown in Fig. 9a and 9b. In other words, the relatively small increase in rule accuracy (from 59% to 70%) suggests that users were willing to refine their policies. Also, as indicated in Figure 10, most users thought that the interface they were provided with to modify their rules was easy to use – the interface had been carefully designed and refined through a number of evaluations with users.

In fact, there is relatively little correlation between policy accuracy and the number of rules specified by participants (Fig 11). Similarly, there is also little correlation between policy accuracy and the time spent by participants refining their rules (Fig. 12). Instead, it seems that users quickly reach a plateau and are often unable to articulate highly accurate policies.

While users seem to have a hard time accurately describing their privacy policies, their feedback tends to be fairly consistent and can be used as a basis for learning more accurate policies. Results displayed in Figure 13 compare the accuracy of policies defined by each of the 19 participants with policies obtained by applying case-based reasoning (CBR) using a k-nearest neighbor heuristic. In this approach, each new situation is compared with prior cases available for a given user. The *k* closest cases cast a vote on whether to disclose the user's location or not (computed individually for each user). CBR systematically improved the accuracy of the policies to 82% (versus 70% when re-applying the user's final policies to each of the scenarios).

6. Field Studies

In Spring 2007, we deployed a first version of PEOPLEFINDER and made it available to three groups of target users. Each target user was asked to invite members of their social network and set up rules so that others could query their locations. The three groups of target users included (1) 15 members of our research team, (2) a group of seven MBA students, and (3) a group of six people involved in organizing buggy races during the Spring Carnival week at Carnegie Mellon. With the requesting users they invited, this amounted to a total of over 60 active users.

The pilot with members of our team spanned a total of six weeks. The pilot with MBA students lasted two weeks and the pilot with Carnival organizers spanned a total of nine days. Usage of the system was rather uneven with some target users having as many as 25 or more requesting users in their list of contacts and others having as few as one or two. For this reason, we limit the results presented in this section to the set of 12 most active target users (and their fairly large social

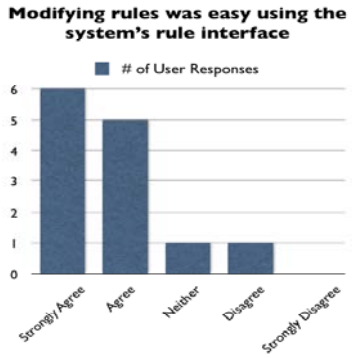

Fig. 10 Difficulty articulating policies is not due to a poorly designed rule interface.

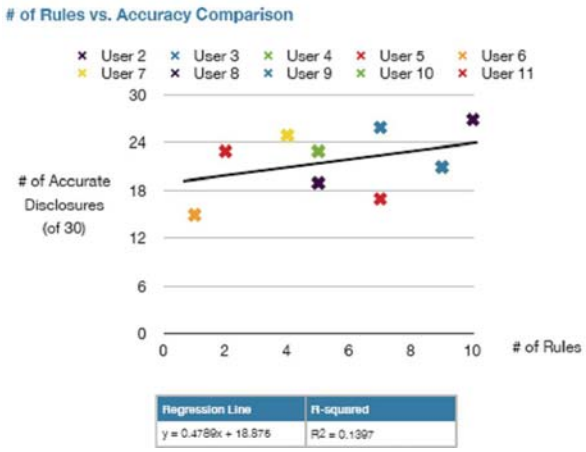

Fig 11: Users reach a plateau: little correlation between (post-hoc) accuracy and number of rules created

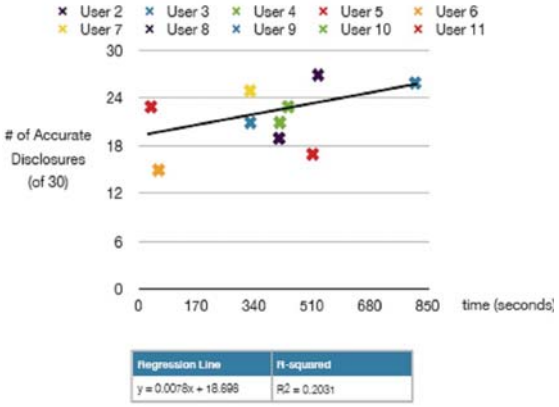

Fig 12: Users reach a plateau: little correlation between (post-hoc) accuracy and time spent defining and refining rules.

networks), as measured by the number of daily requests submitted for their locations. This includes four members of our research team, two MBA students and all six Carnival users. Collectively, these target users were the subject of 1,314 location queries.

Overall the accuracy of the rules defined by the 12 most active users in these 3 pilot studies, as measured by the feedback they provided when auditing their logs (which was generally done once per day) was 79% (Figure 14). This percentage is sensibly higher than the 65% accuracy measured in laboratory experiments involving our PEOPLEFINDER mockup (see Section 5). We believe that the difference can be attributed to several factors. In particular, it seems that users were probably more careful in defining their rules, as they knew they were going to be used to process actual queries from friends and colleagues. We also believe that several improvements in the design of our system played a significant role in helping users define more accurate policies. In particular, this includes the introduction of functionality that lets users see detailed information about the context of each query and get explanations that identify the particular rules behind each disclosure/non-diclosure decision. Other factors such as the significantly larger number of queries per user than in our laboratory experiments (over 100 queries per user versus 30 to 45 scenarios for users of our mockup application) may also have contributed to the increase in accuracy.

While these results are encouraging, post-hoc experiments conducted using a random forest classifier [6] to refine a user's rules based on his or her feedback show that accuracy can probably be further improved (Fig. 14). We are currently working on a new user interface that attempts to combine this insight with new dialogue functionality to help users refine their policies. The objective is to produce rules that are not just more accurate but that the user can also relate to – in contrast to rules obtained through a learning algorithm that acts as a "black box".

A more detailed analysis of user policies over time suggests that users tend to initially err on the safe side as they define their policies. As they become more comfortable with the application and the way in which it is used by their acquaintances, they refine their policies and start allowing requests that in the past would have been denied. This is illustrated in Figure 15, which compares disclosure/non-disclosure decisions made by the user's final rules with those the user had originally defined. While the majority of requests results in the same decision ("same"), the majority of decisions that are processed differently involve changing a non-

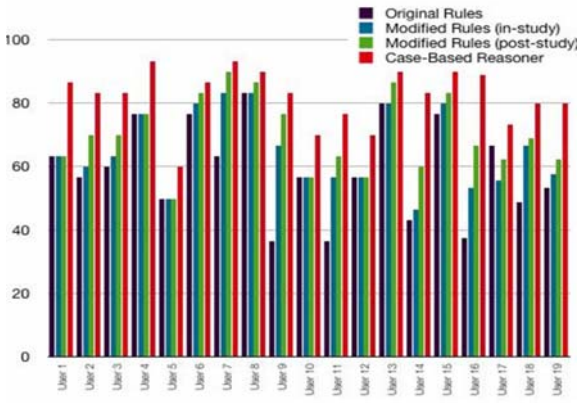

Fig. 13 User feedback can help the system learn the user's privacy policy.

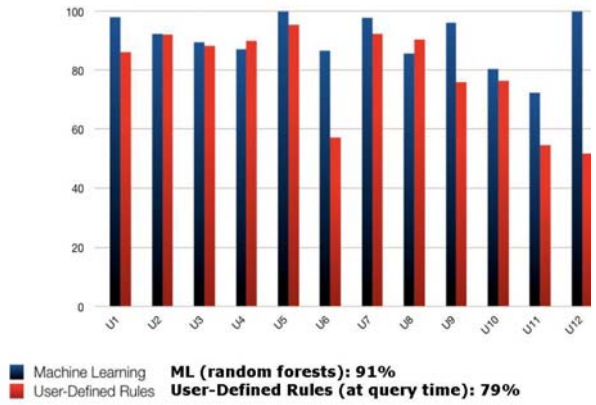

Fig. 14 Results for 12 most active target-users from 3 field pilots involving over 60 users

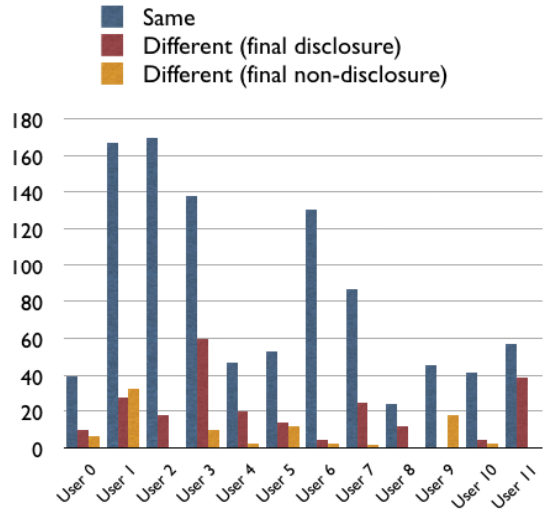

Fig. 15 Policy evolution – 12 most active target users.

disclosure decision into a disclosure decision ("Different: Final Disclosure"). This was the case for 10 out of the 12 most active users.

7. Concluding Remarks and Future Work

In this article, we presented our work on PEOPLEFINDER, an application that enables cell phone and laptop users to selectively share their locations with others. Our main objective has been to better understand people's attitudes and behaviors towards privacy with respect to one pervasive computing application, and to develop technologies and user interfaces that help users specify privacy preferences.

We conducted a laboratory study as well as three field trials involving a total of over 60 participants. One interesting finding is that people have a hard time articulating effective privacy preferences. Functionality that increases user awareness of how the application is used and assists users as they audit queries (e.g. through explanation and access to detailed information about the context of each query) seems to help users define more accurate policies. Early results also indicate that machine learning techniques can help further improve accuracy and be used. As part of our ongoing research, we are developing techniques that use machine learning to provide suggestions to users on how to refine their policies.

Another interesting finding is that people tend to be conservative about disclosures at first, but tend to relax their policies over time as they become more comfortable with PEOPLEFINDER and with how others are using it to find their location. This finding suggests that systems should help people stay in their comfort zones while also helping them evolve their policies over time.

Currently, we are continuing our work with PEOPLEFINDER, developing visualizations that can help people specify policies as well as see how their personal information is being accessed. We are also developing more sophisticated dialogues and explanations, to help people better understand the behaviors resulting from their policies and help them more effectively refine these policies.

8. 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. Additional support has been provided by FranceTelecom, Nokia and IBM. People Finder's WiFi-based location tracking functionality runs on top of technology developed by Skyhook Wireless.

The authors would like to thank all the other members of Carnegie Mellon University's project on "User-Controllable Security and Privacy for Pervasive Computing" for their help designing and evaluating the PEOPLEFINDER application, including Lujo Bauer, Bruce McLaren, Mike Reiter, Paul Drielsma, Alberto Sardinha, Wei Zhiqiang, Jason Cornwell, Gary Hsieh, Rob Reeder, Karen Tang, Kami Vaniea, Yue Zhang, Jacob Albertson, David Hacker, Justin Pincar, and Michael Weber.

References

[1] Cornwell, J., I. Fette, G. Hsieh, M. Prabaker, J. Rao, K. Tang, K. Vaniea, L. Bauer, L. Cranor, J. Hong, B. McLaren, M. Reiter, and N. Sadeh. User-Controllable Security and Privacy for Pervasive Computing. In Proceedings of *The 8th IEEE Workshop on Mobile Computing Systems and Applications (HotMobile 2007)* 2007.

[2] 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.

[3] Intel Research Seattle, POLS: Privacy Observant Location System. http://pols.sourceforge.net/

[4] Hsieh, G., K.P. Tang, W.Y. Low, and J.I. Hong. Field Deployment of IMBuddy: A Study of Privacy Control and Feedback Mechanisms for Contextual Instant Messengers. In Proceedings of *The Ninth International Conference on Ubiquitous Computing (Ubicomp 2007)*, To Appear

[5] Skyhook Wireless website, 2007 - http://www.skyhookwireless.com/

[6] Ho, Tin Kam. "Random Decision Forest". Proc. of the 3rd Int'l Conf. on Document Analysis and Recognition, Montreal, Canada, August 14-18, pp. 278-282, 1995