User-Controllable Location Privacy

Norman M. Sadeh, Patrick Gage Kelley, Lorrie F. Cranor and Jason I. Hong

Carnegie Mellon University
{sadeh;pkelley;lorrie;jasonh}@cs.cmu.edu

Abstract. Over the past decade, the Mobile Commerce Laboratory in the School of Computer Science at Carnegie Mellon University has been piloting a number of context-aware applications and services [1,5,9].[1] From the very beginning an important part of our work has revolved around reconciling the demands associated with context awareness and privacy [3,8]. This short paper summarizes some of the main findings of our work and provides a brief overview of how many of the results from our research have been encapsulated in a User-Controllable Privacy Platform for context-aware computing.

Briefly, when it comes to sharing their location with different applications and services as well as with other users, our research has shown that:

- Users often exhibit complex privacy preferences
- Users often do not fully understand the implications of many of their privacy decisions and require better interfaces to help them make better decisions
- Users prefer auditing functionality that helps them review instances when their location information has been shared. When given this log, they report significant increases in comfort and a greater sense of control.
- Users often do not manipulate default location privacy settings unless they are given functionality that prompts them to do so.
- Because different users often have different location privacy preferences, one-size-fits-all default policies are generally inadequate. Instead users should be helped to select from a small set of default privacy personas
- Simple dialogues and user-oriented machine learning functionality can also help users incrementally refine their privacy preferences over time

Users location privacy preferences are often complex

Figure 1 displays the location sharing preferences of 30 different users, when it comes to disclosing (green) or not disclosing (red) their location to other members of the Carnegie Mellon campus community (see [2,4] for additional details). Each square represents a different user. As can be seen, a small number of users (3) are never willing to share their locations with others, whereas an equally small number of users

[1] http://locaccino.org

(also 3) are always happy to share their locations with others. The majority of users fall in between, with location sharing preferences that vary by day and time during the course of a given day.

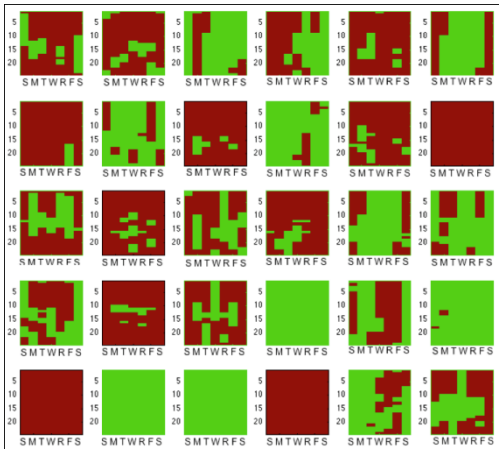

Figure 1. shows privacy preferences for sharing one's location with other members of the campus community. Data collected from 30 different users based on day of the week and time of the day. Green means "share" and red means "don't share." Each square represents a user.

The user's current location is yet another important factor. A common privacy preference is of the type "I am willing to share my location with my colleagues, but only on weekdays, during working hours and only when I am on company premises." Other than time, day, and location, important factors also include the granularity at which location information is disclosed (precise, neighborhood, city level) and the frequency of requests. Sadly, most current location services contain one or less of these control factors [7].

Figure 2 summarizes the accuracy with which different levels of expressiveness in underlying policy languages capture people's actual location privacy preferences (see [2] for additional details). This figure differentiates between the following levels of expressiveness:

- *White List:* unconditional list of entities that can access one's location
- *Time:* Differentiating solely based on time of the day
- *Time+:* Differentiating based on times between weekdays and weekends
- *Loc:* Restricting access based on one's current location
- *Loc/Time:* Restricting access to one's location subject to both time of the day and location restrictions
- *Loc/Time+:* Restricting access to one's location subject to both weekend vs. weekday, time of day, and location restrictions

As can be seen, with the exception of friends and family, simple white-lists are significantly less accurate. More expressive settings have a major impact on our ability to accurately capture people's privacy preferences. This is best illustrated in

the case of location-based advertising and sharing one's location with Facebook friends (see [2] for additional details).

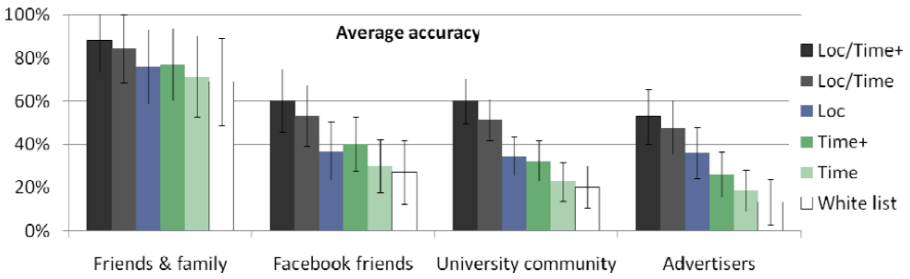

Figure 2. Accuracy with which one can capture a user's location privacy preferences when varying the expressiveness of location privacy settings exposed to that user. Sharing one's location with friends & family, Facebook friends, members of the University community, or advertisers.

Given that users are only willing to invest so much time in configuring their privacy preferences, it is legitimate to ask to what extent the above results hold when taking into account user burden considerations. As soon as users define at least 2 privacy rules (with the rules varied by mechanism), the accuracy with which one can capture their location privacy preferences nearly doubles. The increase is even more dramatic for users willing to define 4 or 5 rules (again, see [2] for additional details).

While a single one-size-fits-all location privacy policy is often inadequate, presenting users with a small number of privacy personas they can choose from can help them take advantage of expressive settings without requiring them to explicitly select from a large number of options.

Default privacy personas can help

Presenting users with just two location privacy personas for configuring their privacy preferences when it comes to sharing their location with family members is sufficient (accuracy of 95%) and adding a third or fourth persona does not add any value. For close friends and members of the university community, a third persona helps increase accuracy over a situation where the user can only select between two personas. All in all, these results show that a small number of privacy personas is often sufficient to capture people's location sharing preferences, even if these preferences themselves are possibly fairly complex – additional details on how to learn canonical personas that are understandable by users can be found in [4].

A key to empowering users to effectively manage their location privacy preferences involves giving them access to simple auditing functionality. Using this functionality, users are able to review when their location has been shared and with whom. This in turn helps them better appreciate the behaviors their current privacy policies give rise to and empowers them to more effectively refine their preferences over time. Studies have shown that users report being much more comfortable accessing location-based

applications when given auditing functionality [5]. In addition, experiments conducted with location sharing applications show that, when given this functionality, users tend (on average) to selectively relax their privacy preferences over time, eventually leading to more sharing [1,5] (see Figure 3 below). In social networking contexts, where the value of an application derives from the amount of sharing it gives rise to, auditing functionality makes applications more valuable. The same has been measured for location-based advertising applications [2].

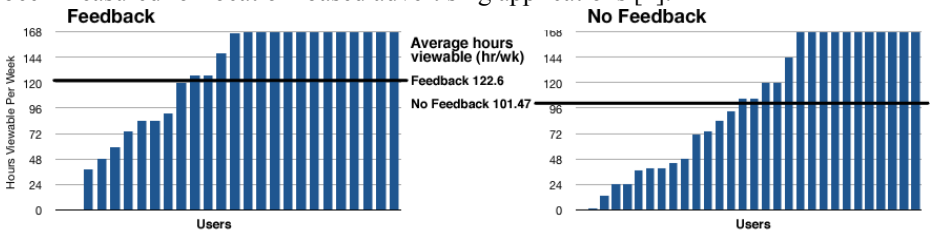

Fig. 3. Providing users with access to auditing functionality increases user comfort and also empowers users to more selectively relax their privacy preferences. Results from a month-long Facebook location sharing pilot involving nearly 60 users split into two conditions, one with auditing and one without.

Supplementing auditing functionality with dialogues and user-oriented machine learning technology to suggest to users how they could possibly improve their current privacy settings can help further improve user comfort by enabling users to converge towards privacy settings that better capture their true preferences [6].

References

1. N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao, "Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application," Journal of Personal and Ubiquitous Computing, Vol. 13, No. 6, August 2009. http://www.normsadeh.com/file_download/8/People+ Finder+PUC.pdf
2. M. Benisch, Patrick Gage Kelley, Norman Sadeh, Lorrie Faith Cranor, "Capturing Location Privacy Preferences: Quantifying Accuracy and User Burden Tradeoffs," Carnegie Mellon University Technical Report CMU-ISR-10-105, March 2010.
3. Norman Sadeh, Fabien Gandon and Oh Buyng Kwon, "Ambient Intelligence: The MyCampus Experience," Chapter 3 in "Ambient Intelligence and Pervasive Computing," Eds. T. Vasilakos and W. Pedrycz, ArTech House, 2006. (Also available as Technical Report CMU-ISRI-05-123, School of Computer Science, Carnegie Mellon University)
4. R. Ravichandran, M. Benisch, P.G. Kelley, and N. Sadeh, "Capturing Social Networking Privacy Preferences: Can Default Policies Help Alleviate Tradeoffs between Expressiveness and User Burden?" *PETS*, August 2009. http://www.cs.cmu.edu/~CompThink/probes/papers/captsocnetwork.pdf
5. J. Tsai, P.G. Kelley, P. Hankes Drielsma, L. Cranor, J. Hong, N. Sadeh "Who's Viewed You? The Impact of Feedback in a Mobile Location Applications", in Proceedings of the 27th annual SIGCHI Conference on Human Factors in Computing Systems (CHI 2009), April 2009. http://www.andrew.cmu.edu/user/jytsai/papers/paper0691-tsai.pdf

6. P.G. Kelley, P. Hankes Drielsma, N. Sadeh, and L.F. Cranor, "User-Controllable Learning of Security and Privacy Policies," First ACM Workshop on AISec, CCS 2008 Conference. Oct. 2008. http://normsadeh.com/file_download/61/CSS+AISec2008+camera+ready.pdf
7. J. Tsai, P.G. Kelley, L.F. Cranor and N.M. Sadeh, "Location Sharing Technologies: Privacy Risks and Controls," to appear in "I/S: A Journal of Law and Policy for the Information Society" A shorter version was presented at TPRC 2009. http://normsadeh.com/file_download/37/TsaiKelleyCranorSadeh_2009.pdf
8. Gandon, F. and Sadeh, N., "Semantic Web Technologies to Reconcile Privacy and Context Awareness" *Journal of Web Semantics*. Vol. 1, No. 3, 2004. http://reports-archive.adm.cs.cmu.edu/anon/2003/CMU-CS-03-211.pdf
9. E. Toch, J. Cranshaw, P. Hankes Drielsma, J.Y. Tsai, P.G. Kelley, L.F. Cranor, J. Hong, N. Sadeh, "Empirical Models of Privacy in Location Sharing." Proceedings of the 12th ACM International Conference on Ubiquitous Computing. Copenhagen, Sept 26-29, 2010