A Social Approach to Privacy in Location-Enhanced Computing

Ian Smith, Anthony LaMarca, Sunny Consolvo
Intel Research Seattle
1100 NE 45th St
6th Floor
Seattle, WA, 98105, USA
+1(206)545-2521
{ian.e.smith, anthony.lamarca, sunny.consolvo}@intel.com

Paul Dourish
School of Information and Computer Science
U.C. Irvine
Irvine, CA, 92697, USA
+1(949)824-8127
jpd@ics.uci.edu

ABSTRACT

Place Lab is a system for positioning a user based on passive monitoring of 802.11 access points. Place Lab seeks preserve the user's privacy by preventing disclosures, even to "trusted" systems in the infrastructure. We are pursing two avenues to explore these and other privacy issues in the domain of socially-oriented applications. We are doing fieldwork to understand user needs and preferences as well as developing applications with significant, fundamental privacy concerns in order to expose the strengths and weaknesses in our approach.

KEYWORDS

Privacy, location-based services, location-enhanced computing, ubiquitous computing, context-aware services.

1. INTRODUCTION

Privacy has long been recognized as a central concern for the effective development and deployment of ubiquitous systems [1-5]. As both a technical problem and a social problem, it is difficult to deal with, to design for, and to model coherently.

The traditional frame within which privacy arguments are cast is a trade-off between risk and reward. This is a popular approach in a range of fields from public policy to cryptography. The risk/reward framework, in the pervasive computing context, suggests that individuals make decisions about technology use by balancing perceived risks against anticipated benefits–that is, in a fundamentally economic approach, they trade off costs against benefits and adopt technologies in which the benefits outweigh the costs, while rejecting those in which the costs outweigh the benefits. Therefore, many have argued, creating successful location-enhanced computing requires finding the most effective balance between risks and rewards [7,8]

This approach has a number of problems, though, both as a conceptual framework and, consequently, as a model for design. Studies of actual practice fail to display the sort of rational trade-off that this model would suggest. There are a number of possible reasons.

First, it is likely that the model is over-simplified and neglects a number of related factors that are important for decision-making about technology adoption and use. For example, we have found that naturally-occurring accounts of privacy behaviors depend on *recourse* as much as risk and reward. By recourse, we are referring to the actions that can be taken by users in the event that others misbehave.

Second, recent research in the area of behavioral economics suggests that traditional rational actor approaches fail to adequately account for everyday behavior even within their own fairly limited terms of reference [25]. The notion of stable exchange-values for goods, services, and labor upon which conventional economic modeling is based seems to fare poorly when applied to human actors who are meant to embody these principles. Instead, psychological and social factors seem to interfere with the mathematical principles of neoclassical economics. In a simple example, while you might pay a neighborhood kid $20 to mow your lawn, you would be less likely to mow your neighbor's lawn for $20. Recent approaches that attempt to incorporate psychological elements into economics models, such as prospect theory, revise traditional notions of commodity and value.

Third, and perhaps more fundamentally, studies of technological practice suggest that technology adoption and use should be seen not simply in terms of individual decisions about costs and benefits, but rather in terms of broader patterns of participation in cultural and social life. For example, in Harper's (1992) study [9] of the use of active badges in research laboratories, it is telling that a number of people report participating in the use of the system in order to be seen as team players, in order to provide support to others, etc. In other words, social actions have symbolic value here, and these are frequently the more salient elements of adoption decisions. Ito's studies of mobile messaging amongst Japanese teens [10], or the studies by Grinter and colleagues of the use of SMS and Instant Messaging amongst teens in the US and the UK [11-12] describe the use of messaging technologies as cultural practices, essentially casting the adoption of these technologies as forms of participation in social life. To use the technologies is simply part and parcel of appropriate social practice. As technologies become increasingly integrated into everyday practice, rational decision-making about privacy trade-offs is increasingly irrelevant.

Fourth, studies of privacy management in the everyday world, drawing on studies in social psychology, suggest that privacy management is a much more nuanced and contingent phenomenon. Drawing on the work of Irwin Altman, Palen and Dourish [13] present a model of privacy as a continual and

dialectical process of boundary regulation. These boundaries are not simply barriers to information flow, but are also the boundaries between self and other through which differentiation and affiliation are achieved, and boundaries between past and future that reflect the emergence of genres or conventions for information practice. Some of this can be seen in studies of personal web pages [14] and increasingly, lately, studies of blogs (e.g. Nardi et al [15]) where genres arise that provide both expectations and interpretive norms for understanding information disclosure. For instance, where most personal web pages are unlikely to state the details of where people can be found at particular hours of the day, that is an appropriate and indeed expected form of disclosure for college professors. The dialectic model that Palen and Dourish propose has a number of important implications for design that are quite at odds with traditional rational actor approaches. Principally, they situate information disclosure settings within the immediate circumstances of activity, suggesting that the "costs" and "benefits" of information disclosure are continually subject to negotiation and change.

Finally, one implication of these broader perspectives for traditional approaches to the specification and description of location-based or context-aware computing is that the very notion of "context" may be problematic – it may not be something that can be uniquely defined, but depends on the person to whom the context is being disclosed, or the specific features of the setting in which the formulation is made [16]. We will return to this later, in describing a field study of the ways in which context is formulated [17].

Accordingly, in our work, we have been developing an alternative to traditional formulations of privacy, both as a conceptual framework and a technical approach. Our approach in Place Lab [24] attempts to avoid the abstract formulation of privacy needs and the "disclose and hope" model that requires them (see below.)

Our essential argument, then, is that there are no abstract rules by which privacy is formulated; rather, the information practices that we refer to under the rubric of "privacy" are emergent phenomena of everyday social action.

One common objection to this argument is that, while rules and resources may not be part of our conscious experience of information practice, they must nonetheless be underlying factors, which we have learned and internalized so that they are no longer consciously available to use. We all had to be taught these rules, once upon a time; every one of us, after all, has a story of the moment when, as a young child, we loudly make some remark that was wildly socially inappropriate and embarrassing to our parents. So, the basis for our current practice must be rule-based, even though those rules are no longer part of our conscious experience.

However, this objection is fallacious. It is broadly equivalent to this argument–that when learning to ride a bicycle, we managed to stay upright through the use of training wheels. Once we became competent bicyclists, we no longer used training wheels but, even though the training wheels are no longer visible, they must, nonetheless, be the basis of our balance.

2. CLASSES OF LOCATION APPLICATIONS

Place Lab is a research effort to build a low-cost, widely-available, indoor-outdoor positioning system [18-21]. Devices running Place Lab use radio beacons in the environment (such as 802.11 access points) as known "way points" that can be used to triangulate location. Since an increasing percentage of computation devices are shipping with some radio sensing capability (such as 802.11 or Bluetooth), a map of known beacons allows them to position themselves with no additional hardware. One advantage Place Lab has over many other location technologies is that it is based on passive monitoring of radio signals and local lookups and computation. As such, devices running Place Lab can position themselves completely locally and only need disclose their location when it is desired by the user[1].

Our initial explorations with Place Lab have shown that location-enhanced applications fall broadly into three classes: institutional, social, and personal. These classes of applications differ based on the person or organization to whom location information is disclosed. A *personal* application is one that does not need to disclose location information to anyone to be effective. An example is a pedometer or other personal fitness applications. Another set of personal applications are way finding or route planning applications. These types of applications may need the user's location to function properly, but it is not necessary to communicate that location to anyone given local storage and possibly a cache of content.

Institutional applications are a more common arrangement, requiring that people disclose information to a central authority (normally, an organization) in return for some service. Active Badge systems [23] and related context-based services operate according to this model; information about location is relayed to a central server, while then makes contextualized services available to clients and users. This architectural approach made sense when both client-side computation and network bandwidth were limited, and so has been a common structure in prototype ubicomp systems. However, given the relentless march of time and Moore's Law, alternative technical approaches are now more feasible, and avoid the sorts of privacy commitments being made in this architecture.

It should be noted that it is possible to build the same institutional application with varying degrees of disclosure on the part of users. For example, if Google made their index of web pages publicly available, one could turn Google into a personal application since a user could do their searches while disclosing little to no personal information. In this scenario, one could download the entire medical index and then search locally for a specific condition, revealing the possible interest in a medical condition, but no more beyond that. However, in most cases, institutional applications have substantial commercial, public interest, or intellectual property barriers that prevent them from being organized in this open fashion.

The final class of applications in our taxonomy is *social*. These applications require disclosure to people, rather than institutions to work effectively. Many ubiquitous-computing services, such as friend finder [24] or context-aware chat [29], are examples of social applications. A friend finder is an application that alerts you when one of your "friends" is nearby, facilitating serendipitous social interaction. Clearly, this requires at least that the user and her friend's locations be exchanged in some way.

[1] A number of other technologies including GPS have this same advantage that location is computed locally.

There are risks in social applications, although they are not as clear as some other scenarios. In the friend-finder example, by what mechanism should "friends" be designated? Certainly, it should require some type of mutual acceptance, otherwise the system can and will be abused by anyone with the technology. Avenues for recourse are also unclear. Are the forces of recourse—such as social isolation or embarrassment—strong enough to affect user behavior? With due respect to considerations of risks and recourse, we are more interested in how this technology will be adopted be social actors. It easy to imagine that being on someone's "friends list" in a friend-finder application might be as important as being in someone's cell-phone address book. Studies of the gift-giving practices of teens [26] have revealed the social impact of being "in" the social space of someone's cell-phone address book to be significant.

2.1 Implementation Strategies

There has been a trend in ubiquitous computing research towards systems with a "disclose and hope" feel. These systems require the user to make substantial disclosures to systems like smart houses [27], active work spaces [28, 29] and location-tracking services [23] and then hope that these systems make an effort to keep the information private. We do not believe that this is appropriate, and we feel that these designs should be avoided. At the point that the disclosure is made, the user's privacy is, in reality, lost. Designers of these "hopeful" systems suspect that their algorithms and strategies are correct, that the system managers are not incompetent, and that no subpoenas are issued against some or all of the private data they manage.

We vastly prefer designs that preserve privacy by simply not disclosing unnecessarily, i.e., keeping risk as low as possible. In the case of location information, such as that generated by Place Lab, the risks associated with disclosing information are significant. We believe that Place Lab's passive nature allows us to at least have the option to design systems with minimal disclosure.

3. SOCIAL APPLICATIONS, PRIVACY, AND PLACE LAB

Previously, we argued that simple models that imply that people are rational actors making a narrow choice such as "will I give away this information for that commodity" are insufficient to explain the privacy-related behaviors we observe. If there are areas in which people can be seen as close to making rational choices it is the area of personal applications. Because disclosures to others are not required for personal applications, fewer social forces come to bear and an individual can make decisions "flying solo." This is not say that a simple risk versus reward calculation can be employed to predict user behavior—that would ignore issues such as user-interface concerns that still exist in personal applications. In the case of the pedometer personal application, issues such as size, weight, visibility to others, and battery life are quite significant to ultimate user adoption. Even the designation "personal" is troublesome here; if a pedometer is implemented in a "disclose and hope" fashion, the personal application takes on social dimensions as it can be used to track those that are walking with you.

Institutional applications are also problematic unless situated in their social context. Consider workplace-safety applications of location-tracking technology. Organizations and institutions might view this as a positive development, decreasing accidents or preventing workplace violence. Individuals who work for these organizations are likely to have many complex relationships to the deployment of such a technology [9] and the institution that deploys it. Yet again, the individual user's relationship to the organization and the deployed technology is not a simple matter of a trade-off in risk versus reward.

One of the goals of the Place Lab project is to build location infrastructure that will foster the development of successful applications. Unfortunately, as we have argued in Section 1, the inherent value of an application is a complex and unpredictable metric to predict. Of the three classes of application, the value of those in the social domain is the most unpredictable and often counter-intuitive. For this reason, we have chosen to initially focus our study of privacy and its relationship to location-enhanced computing on this class of application.

In the Place Lab project, we have begun two efforts to better understand the future space of social, location-based applications and how people will formulate the social norms governing their use. The first is a field-study to expose situated user concerns and the second is an application to help us directly experiment with these issues.

3.1 A Field-Study of Privacy Concerns

We are conducting a user study to understand people's perceptions about privacy and how time, place, and other people affect the types of disclosures they might make. In other words, we are trying to understand the social factors that would affect our future application development. Our study design uses the experience sampling method [30] or ESM (often called a "beeper study"). In an ESM study, a participant is given a mobile device such as a PDA that periodically alerts the user and asks a question(s). While incurring more overhead and interruption than techniques such as diary studies, ESM data is typically highly accurate as it is collected in situ and does not require recall.

In our case, this allows the participant to answer questions about location in the actual location, not in a lab or conference room days later. An additional advantage arises from the fact that participants will be carrying a computational device with them during the course of the study. Since we already have the Place Lab positioning infrastructure running on small devices, we can create questions that are customized based on the user's location. For example, through Place Lab, our ESM application might know the user's location and look up that location in a database of business records. We can then discover if the city or county business records, perhaps "Smith's hardware store," matches well with how users self-report their location. We believe that this comparison will shed light on how users' perception of risk varies with time and physical location.

Some examples of the types of questions we are designing into our study are:

- If your boss asked you for your location right now, how would you answer? Your spouse?
- If your mother asked where you were right now, would you answer 'a bank,' 'the corner of 45th and Vine,' or something else?
- Would you tell Alice your location right now in exchange for hers? If so, what would you be

comfortable telling Alice? What would you want to know about Alice's location?

3.2 Ambush: A Dangerous, Yet Privacy-Aware Application

Rather than trying to develop locations-enhanced applications that skirt privacy issues, we have chosen the opposite approach. We have devised an application that we believe offers substantial new functionality while at the same time presents significant privacy risks. In this way, we hope to attack the privacy issue "head on" by experimenting with privacy strategies and mechanisms.

Our application is called "ambush" and is based on the work of Mynatt and Tullio. In [22], Mynatt and Tullio describe an ambush as the use of a shared calendaring system to infer a person's probable location in the future with the intent of "ambushing" them for a quick face-to-face meeting. This process is used frequently in larger organizations, particularly by subordinates, to have brief conversations with senior managers who are between meetings.

We have generalized the notion of ambush to be any location, not just conference rooms visible in a shared calendar system at work. Our ambush application allows a user Alice to define a geographic region—say a public park—and ask to be notified anytime Bob enters that region. If Alice lives near the park and wants to visit with Bob, clearly both can benefit from the possible serendipitous, social encounter in the park. Another use of ambush is micro-coordination. Such tasks are common in urban environments, such as "Let me know when Charles or DeeDee get to the subway station so I can go meet them." Another use of ambush is the creation of social capital [31] through discovery of shared interests that are demarcated by places, such as bookstores, music venues, or civic organizations. It should be noted that current "friend finder" systems offered by cell-phone providers are actually corner-cases of our ambush application in which the only location that can be specified is "near me."

The potential for nefarious activities with ambush are rife, making risk a significant issue. As previously stated, we chose ambush as a test application because it forces to come to grips with the privacy concerns.

As an aside, we are not concerning ourselves right now with the security and authenticity issues of ambush. We are not addressing questions like, "How do I know that no malicious entity modified or hacked the users' devices to steal their location information?" or "How can I be sure that this geographic region is Green Lake Park as Alice purports and is not my home as I suspect?" Although these are interesting questions, we are focusing our initial investigations on the privacy issues.

We have devised several concrete strategies to help us address the privacy concerns in ambush. First, our privacy concerns field study with ESM mentioned above will include questions that are specifically tailored to an ambush-style application. This can help us craft our technical strategies to be sensitive to the social norms and perceptions of our user community.

Without going into tremendous detail, we are considering three significant techniques to blunt the privacy concerns in ambush. All of these are currently be explored through our early efforts.

- Reciprocity: If you get someone else's location you give up your own. Although this strategy is vulnerable to certain types of abuse—notably that people who do more things and go more places have more to lose than those that stay at home constantly—it offers some advantages. It allows those who disclose their location to know who requested the information; if the location offered in reciprocity is of little value ("always at home"), perhaps social norms of recourse can be used to deter abuse.
- Explicit acceptance: This seems central to our strategy of preserving privacy. You have to take explicit action to disclose your location, so it is at least possible for you to be aware of others' attempts to observe you, for good or ill. This has the obvious problem that it does not scale well to large numbers of disclosures of your location. Either you will become irritated with the frequent disturbances or become "numb" to the action and cease to really make a decision about the disclosure. Both this technique and the previous one are situated primarily the social domain for both the user's understanding of the disclosures being made as well as the possibilities for recourse.
- Indirection: Perhaps Alice should "make an argument" to Bob for the release of his information to her. In this model, Place Lab does not disclose Bob's location to Alice, but rather shows Bob Alice's argument (perhaps in text form) when he enters the park. "Bob: We should get our kids together in the park. Call me. –Alice." This technique can easily be combined with either of the first two for additional benefits. This is a similar to many systems that leave information at specific places in the world, but it is focused on the two users rather than leaving information "for anyone."

4. CONCLUSION

Despite being in the early stages of the Place Lab project, we know that accurately recognizing and addressing privacy concerns is critical to the success of our system as a platform for location-enhanced computing. Unfortunately, understanding disclosure of user's information and its relationship to an application's success is difficult to predict. This is especially true in the domain of social applications in which users disclose personal data to other individuals. To increase our understanding of applications in this domain, we are running an ESM study to learn how location, context and place interact with a user's inclination to disclose information to others. To gain experience with a particular application, we are building and plan to deploy "ambush" a request-driven location service. By building and deploying a useful yet dangerous application like ambush, we hope to develop an understanding of how applications interact with social norms.

5. REFERENCES

[1] L. Barkhuus and A. Dey. Location-based services for mobile-telephony: a study of users' privacy concerns. In *Proceedings of INTERACT 2003, 9th IFIP TC13 International Conference on Human-Computer Interaction*, 2003.

[2] V. Bellotti and A. Sellen. Design for Privacy in Ubiquitous Computing Environments. *In Proceedings of The Third European Conference On Computer Supported Cooperative Work (ECSCW '93)*. Milan, Italy: Kluwer Academic Publishers, 1993.

[3] J. Hong, G. Boriello, J. Landay, D. MacDonald, B. Schilit, and J. Tygar. Privacy and Security in the Location-enhanced World Wide Web. In *Workshop on Ubicomp Communities: Privacy as Boundary Negotiation (Ubicomp 2003)*. Seattle, WA, 2003.

[4] E Kaasinen. User Needs for Location-aware Mobile Services. *Personal and Ubiquitous Computing* 7(1): p. 70-79, 2003.

[5] S. Lederer, J. Mankoff, and A. Dey. Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In *Proceedings Of Extended Abstracts of CHI 2003, ACM Coneference on Human Factors In Computing Systems.* Fort Lauderdale, FL, pp 724-725, 2003.

[6] M. Ackerman, T. Darrel, and D. Weitzner. Privacy in Context. Human Computer Interaction 16: 167-176, 2001.

[7] M. Gruteser and D. Grunwald. A methodological assessment of location privacy risks in wireless hotspot networks. In *Proceedings of the First International Conference on Security in Pervasive Computing*, 2003.

[8] B. Schilit, J. Hong, M. Gruteser. Wireless location privacy protection. IEEE Computer 36 (12): 135-137, 2003.

[9] R. Harper. Looking at Ourselves: An Examination of the Social Organization of Two Research Laboratories. In *Proceedings of the ACM Conference Computer-Supported Cooperative Work.* Toronto, Canada, 1992.

[10] M. Ito and O. Daisuke. Mobile Phones, Japanese Youth, and the Re-Placement of Social Contact. Front Stage-Back Stage: Mobile Communication and the Renegotiation of the Social Sphere. Grimstad, Norway, 2003.

[11] R. Grinter and M. Eldridge. Y do tngrs luv 2 txt msg? In *Proceedings of the European Conference On Computer Supported Collaborative Work (ECSCW 2001)*. Bonn, Germany, 2001.

[12] R. Grinter and L. Palen. Instant Messaging In Teen Life. In Proceedings of the ACM Conference On Computer Supported Cooperative Work (CSCW 2002). New Orleans, LA, 2002.

[13] L. Palen and P. Dourish. Unpacking "Privacy" for a Networked World. In *Proceedings of ACM Coneference on Human Factors In Computing Systems.*

[14] K Crowston and M. Williams. Reproduced and Emergent Genres of Communication on the World Wide Web. The Information Society 16, p. 201-205.

[15] B. Nardi, D. Schiano, M. Gumbrecht, and L. Swartz. I'm Blogging This: A Closer Look at Why People Blog, Communications of the ACM, to appear.

[16] P. Dourish. What We Talk About When We Talk About Context. Personal and Ubiquitous Computing 8 (1).

[17] E. Schegloff. Notes on Conversational Practice: Formulating Place. In D. Sudnow (ed.), *Studies in Social Interaction*, 75-119, New York, New York, 1972.

[18] P Bahl and V. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In INFOCOM (2), p. 775-784.

[19] P. Castro, P. Chiu, T. Kremeneck, and R. Muntz. A Probabilistic Room Location Service For Wireless Neworked Environments. *In Proceedings of Ubicomp*, Atlanta GA, 2001.

[20] A. Ladd, K. Bekris, A. Rudys, L. Kavraki, D. Wallach, and G. Marceau. Robotics-based Location Sensing Using Wireless Ethernet. In *Proceedings of MOBICOM*, p. 227-238, 2002.

[21] J. Hightower and G. Boriello. A Survey and Taxonomy of Location Sensing Systems for Ubiquitous Computing. University of Washington Technical Report, CSE 01-08-03, 2001.

[22] E. Mynatt, and J. Tullio. Inferring calendar event attendance. In *Proceedings of the ACM Conference on Intelligent User Interfaces (IUI 2001)*. Santa Fe, New Mexico, p 121-128.

[23] M. Sprietzer and M. Theimer. Providing Location Information In A Ubiquitous Computing Environment. Parel session in *Proceedings Of Fourteenth ACM Symposium on Operating Systems Principles*. Asheville, NC, 1993.

[24] B. Schilit, A. LaMarca, G. Boriello, W. Griswold, D. MacDonald, E. Lazowska, A. Balachandran, J. Hong, and V. Iverson. Ubiquitous Location-Aware Computing and the Place Lab Initiative, First ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH), San Diego, CA, 2003.

[25] M. Rabin. Psychology and Economics. Journal Of Economic Literature, 36, p. 11-46.

[26] A. Taylor and R. Harper. The Gift of the Gab: A Design Oriented Sociology Of Young People's Use Of Mobiles. *Journal of Computer Supported Cooperative Work* (CSCW), 12(3), p. 267-296.

[27] G. Abowd, C. Atkeson, A. Bobick, I. Essa, B. MacIntyre, E. Mynatt, and T. Starner, in Extended Abstracts of ACM Conference on Human Factors in Computing Systems (CHI), 2000.

[28] B. Johnson, A. Fox, T. Winograd. The Interactive Workspaces Project: Experience with Ubiquitous Computing Rooms. IEEE Pervasive Computing Magazine 1(2): April-June 2002.

[29] A. Ranganathan, R. Campbell, A. Ravi, and A. Mahajan. ConChat: A Context-Aware Chat Program. In IEEE Pervasive Computing, p. 52-58, July-Sept 2002.

[30] S. Consolvo and M. Walker. Using the Experience Sample Method to Evaluate Ubicomp Applications. In *IEEE Pervasive Computing Mobile and Ubiquitous Systems: The Human Experience*, 2 (2), p. 24-31, Apr-June 2003.

[31] R. Putnam, *Bowling Alone*, New York, New York, Simon and Schuster, 2000.