Developing Privacy Guidelines for Social Location Disclosure Applications and Services

Giovanni Iachello
College of Computing and GVU Center
Georgia Institute of Technology, USA
giac@cc.gatech.edu

Ian Smith
Sunny Consolvo
Mike Chen
Intel Research
Seattle, WA, USA
{ian.e.smith, sunny.consolvo, mike.y.chen}@intel.com

Gregory D. Abowd
College of Computing and GVU Center
Georgia Institute of Technology, USA
abowd@cc.gatech.edu

ABSTRACT

In this article, we describe the design process of Reno, a location-enhanced, mobile coordination tool and person finder. The design process included three field experiments: a formative Experience Sampling Method (ESM) study, a pilot deployment and an extended user study. These studies were targeted at the significant personal security, privacy and data protection concerns caused by this application. We distil this experience into a small set of guidelines for designers of social mobile applications and show how these guidelines can be applied to a different application, called Boise. These guidelines cover issues pertaining to personal boundary definition, control, deception and denial, and group vs. individual communication. We also report on lessons learned from our evaluation experience, which might help practitioners in designing novel mobile applications, including the choice and characterization of users for testing security and privacy features of designs, the length of learning curves and their effect on evaluation and the impact of peculiar deployment circumstances on the results of these finely tuned user studies.

Categories and Subject Descriptors

D.2.2 [**Design Tools and Techniques**]: Evolutionary Prototyping; D.2.1 [**Software Engineering**]: Requirements/ Specifications—*elicitation methods*; K.4.2 [**Computers and Society**]: Social Issues; K.8.m [**Personal Computing**]: Miscellaneous

General Terms

Design, Human Factors, Security

Symposium On Usable Privacy and Security (SOUPS) 2005, July 6–8, 2005, Pittsburgh, PA, USA
.

Keywords

Social Mobile Applications, People Finder, Iterative Design, Prototyping, Privacy

1. INTRODUCTION

We are interested in designing social mobile applications, *i.e.*, mobile Information Technology (IT) applications which facilitate everyday social interactions. Social mobile applications include text messaging (texting) services, person finders, and availability managers. During the last few years, a slew of specialized applications has emerged, thanks to the availability of powerful computing platforms (the most common being smart phones and networked PDAs), infrastructure software and novel context sensing techniques. In the present article, we will concern ourselves specifically with social location disclosure applications, that is, applications that enable the communication of location among individuals within their social networks.

These applications are widely considered to have a strong commercial potential, especially with consumers in younger age groups. For example, a market survey of US cell phone users conducted in 2004 showed that person finder applications are the second most popular choice among data-intensive applications people would use on their cell phones if they were to spend an additional USD 5–10 on their monthly bill [21].

Despite their promising commercial outlook, the applications that have been launched to date have not performed well in the marketplace. The most widely deployed person finder in the US mobile phone market, AT&T Find People Nearby, has failed to become a large market success, and AT&T Wireless's new parent company (the operator Cingular) may discontinue the application as part of the merger [2]. Other person finders, such as Dodgeball [6], which do not rely on operator support, have a fringe following of dedicated users, but are far from being widespread. Child tracking applications, corporate employee management and similar specialized services, deployed in several countries (*e.g.*, UK, Japan) in collaboration with operators have experienced somewhat better success.

The reasons for the lukewarm acceptance of social location disclosure applications may lay in several interrelated fac-

tors, including privacy concerns, regulatory barriers, steep learning curves, unripe deployment environments, and technical issues (including limitations in the location technology and system usability). We are specifically concerned with those impediments that can be addressed by designers, and for the purpose of this workshop, we will concentrate on issues pertaining to privacy, security and the usability of functions designed to achieve specific requirement goals in these areas.

In addition to design, the evaluation of these applications is also challenging, because it is necessary, when studying their adoption patterns, to take into account a number of interrelated social ramifications. The atypical usage contexts, and unstable use patterns of social location disclosure applications combine with a lack of understanding of how people really use mobile applications which has started to be tackled only recently [8, 9]. This complicates the interpretation of any experimental results that might arise from the observation of usage patterns in a real world deployment. In particular, it has been traditionally difficult to apply acquired knowledge and practices to novel applications, which forces designers to constantly re-invent the wheel when designing new applications.

As a concrete example, consider the aforementioned AT&T Find People Nearby application: although the application was developed several years after the great success of another simple social mobile application, text messaging, many of its compelling features have not been taken into account by the designers of Find People Nearby. Concentrating only on privacy, it is easy to observe how Find People Nearby lacks the ability to fine-tune one's availability (the user is either visible to his/her entire buddy list, or invisible); finely tuned communication is instead one main characteristic of text messaging. Given the lack of published studies on the matter, it is difficult to understand to what degree the inability to carefully manage one's availability has curtailed adoption. Moreover, recent (retrospective) ethnographic literature on the topic highlights the importance of fine-tuning availability in social communications [12, 20]. Newer systems such as Dodgeball, perhaps learning from this, empower users with much more fine-grained management of availability.

In the present article, we build on experience we accumulated over the past year to describe general security and privacy features which impact the design of social location disclosure applications. In Section 2, we discuss related work in the area of design for security and privacy for social location disclosure applications. In Section 3, we describe the design process of an integrated messaging and person finder application, concentrating on the privacy and security questions engendered by this application and how we attempted to solve them. In Sections 4 and 5 we translate our experience into generally-applicable guidelines for driving the design of mobile and ubiquitous computing applications. Finally, in the last section we show how these guidelines can be applied to the development of a new type of social location disclosure application which is map-based rather than text message based.

2. RELATED WORK

This brief overview of related work is by no means intended to be an exhaustive bibliographic treatise. Our intent here is that of pointing out some of the sources which are most closely related to the topic of this article.

Design guidelines for enhancing security and privacy in mobile computing have revolved around various implementations of the Fair Information Practices (FIPS) [19]. Langheinrich discusses guidelines for supporting privacy-related application goals in ubiquitous computing applications, drawing from the FIPS [11]. While certainly worthwhile, these guidelines are also very general and are difficult to translate in practical suggestions for specific problems. With the present article, we try to provide more detailed and hopefully more helpful design guidance.

Lederer *et al.* have investigated the issue of location privacy and reported results from surveys indicating that people decide whether to disclose information about their activities and location based on the identity of the requester more than on the situation in which this happens [13]. Their accounts were quite helpful in the initial phases of this research and were later confirmed by our EMS study (see below).

In a study of college students on a major US campus, Barkhuus and Dey have investigated the balance between security and management burden and have suggested that people are willing to forgo some control over their personal location information if the application is useful to them [1]. We believe that acknowledging and managing this tradeoff is fundamental for successful designs.

The topic of environmental privacy has been discussed in relation to mobile phones and their usage contexts. For example, Ito [9] and Ling [14] report of similar communication practices by teens of very different cultures (resp. Japan and Norway) and how they use mobile messaging as an unobtrusive way of achieving their social communication goals in the face of a strict conduct code imposed by their environments (parents, school, ...). We believe that environmental privacy should be considered an integral component of the impact of communication technologies on the general notion of privacy, as pointed out further below.

3. DESIGN PROCESS

Over the past year, we have been exploring how to design a privacy-observant application that allows people to communicate their location to other members of their social network. Our investigation incorporated three phases of user studies.

We performed a formative Experience Sampling Method (ESM) study, prior to the development of the application. The goal of that study was to inquire to whom people were willing to disclose their location and at what level of detail.

Following that study, we developed an initial application prototype, which we called Reno, and conducted a short pilot study internal to our organization to inquire the usage patterns of the application.

Finally, issues related to privacy, the management of information disclosure and the ability of denying one's availability and location were specifically inquired in a third, longer, deployment with two groups of teenage and adult participants.

Below, we provide an overview of each stage of our work. This cursory overview is intended to provide a framing for our later discussion and is not meant to be comprehensive. Large parts have been published elsewhere and references are provided.

3.1 Understanding the User And the Problem

To begin our investigation, we conducted a two-week long

formative study in July 2004 with 16 adults to help us understand *what* people were willing to disclose about their location to various members of their social networks. This study involved a variety of techniques—both in lab and *in situ*—which included a social network exercise, demographic questionnaires, interviews, two weeks of experience sampling, and a nightly voicemail diary. The two weeks of experience sampling allowed us to explore *in situ* how participants wanted to reply to requests for their location from members of their social networks. That is, several times per day, every day, the participant was interrupted with a hypothetical request for his location from a member of his/her social network and asked to respond. Immediately following the response, we followed up with questions about why they chose to respond as they did. Participants also completed the survey for the Westin/Harris Privacy Segmentation Model [16] which classifies people as being privacy fundamentalists, pragmatists, or unconcerned in the area of consumer privacy. Details of the study are described in [3].

3.1.1 Results

Our most notable findings were that 1) participants want to disclose either the level of detail that they think would be most useful to the requester or that they would deny the request and 2) the participants' privacy classification as determined by the Westin/Harris Privacy Segmentation Model was not a good predictor of how they would respond to requests for their location from social relations. Additionally, we did not see evidence of participants intentionally *blurring* their response, *i.e.*, sharing a level of detail that is accurate but vague, in an effort to protect their privacy. From our results, we reflected on the decision process participants went through to determine what they were willing to disclose about their location to a request from someone they knew; that process is roughly: *who* is requesting, *why* do they need to know, *what* would be most useful to them, and *am I willing* to share that?

There were several findings that directly impacted the prototype we built following this study. One important finding was that participants shared different levels of detail about where they were based on who was asking, what the participant was doing, and/or why they thought the requester wanted to know. For example, participants often shared something vague, like the city they were in, with requesters who lived out of state—not because participants were trying to protect their privacy, but rather because more detail was likely to be meaningless to people from out of state. Additionally, if the participant was out of state for work or vacation, they often felt that sharing the state they were in was the most useful level of detail, as the requester would know that they had arrived where they needed to be. This meant that our prototype would need to support various levels of location granularity.

Despite the fact that willingness to share was strongly tied to the requester's role in the participant's social network (*e.g.*, participants were usually willing to share something with their spouse/significant other), participants expressed a need to be able, on occasion, to "stretch the truth" (or, in blunt terms, to *deceive*[1]). For example, a common scenario that was raised was that if the participant's spouse wanted to know if she was on her way home and she hadn't really left the office/store/friend's house, she wanted to be able to say that she was on her way home already. This meant that the prototype would need to support the ability for the user to disclose a location other than her current/actual location.

Another finding was that most participants expressed that they would only use a location disclosure application with close family and friends. For example, many participants were very bothered by the hypothetical requests that came from their manager and even found requests from some of their (not too close) friends or family members to be awkward, particularly when requests came from people who lived out of town. This finding helped us focus our target user group on close family and friends, most of whom live reasonably near the user, for subsequent phases of this research. We also learned that most participants imagined that a location disclosure application would be most useful for coordinating with others (*e.g.*, trying to meet up someplace, either planned or serendipitously) and 'okayness checking' (*i.e.*, assuring someone that you arrived safely where you were planning to go; that you have not yet made it to where you were supposed to be but that you are okay; or that you want to know that someone else is okay).

While this study gave us a lot of insight about how participants would like to use a location disclosure application and the concerns they have of how it could be abused, the hypothetical nature of the requests meant that the participants did not have to deal with the social ramifications of their responses. The logical next step was to deploy a working prototype with members of a real social network.

3.2 Understanding Usage

In the late summer of 2004 we started developing Reno, a location disclosure application. Later in the Fall, we conducted a pilot study of the Reno application. Both the application design and the study were influenced by the formative ESM study described above. We set out to investigate what people would communicate to whom, in a real setting, and under what circumstances. We focused on a small group of colleagues and their families as suggested by the ESM study results. The pilot study had eight adult participants; all were members of our research group or their family members. The study lasted for five days (Thursday-Monday) and has been more fully discussed elsewhere [18].

This early version of Reno offered basic location functionality. The user could ask for another person's location ("query" another user) as well as disclose his or her own. Location queries and disclosures were sent directly via SMS (*i.e.*, not through a special location "service") and appear in a special message 'inbox' (see Figure 1).

A key feature of Reno is that users teach the mobile phone the definition of their places. When the user is present at a location, like "home," the user tells the phone the name of the current place. Reno then associates the name given by the user with a short, recent history of cell towers to which the phone was connected. This has a positive impact on privacy since users control the name given to the location and presumably will not name locations that they are not comfortable having disclosed. Moreover, users could label the same location with different names, in order to adjust their disclosure to what the recipient would find most useful. Users could also deny their response by simply ignoring the request.

[1]The word deception is used in the present context to indicate the entire spectrum of untrue statements, from stretching the truth to outright lying.

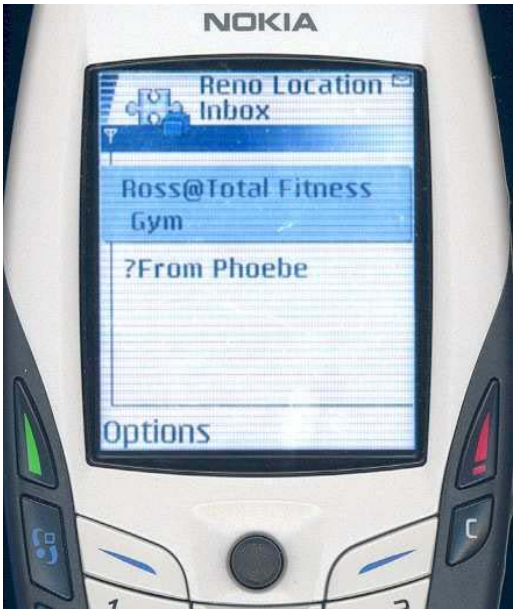

Figure 1: Screenshot of The Reno Application. The location "inbox" shown here contains two messages, a disclosure from Ross and a query for the user's location from Phoebe.

When disclosing location, the user was given a list of "nearby" places to choose from, which speeds up the process of disclosing. The mobile phone sensed its location based on the cellular network, meaning that the list presented to the user was typically quite short and the alternatives were close to the user's current location. The list of nearby places was populated from places previously entered by the user, thus the application never discloses anything that has not been explicitly assigned by the user. In practice, even the simple algorithm used in the pilot typically produced a list with 4–5 elements and with places that were within about 100 meters of the user's actual location.

In addition to manual operation, Reno provided a feature called *triggers*[2] that allowed the user to set up an automatic disclosure. This disclosure was triggered whenever the mobile phone detected that the user was at a specific location, thus allowing semantics like: "disclose that I am at the store to my wife anytime I arrive at the grocery store." Common wisdom in information security suggests that tight control on information is necessary to preserve privacy. On the other hand, the idea of calm technology proposes that users offload their information management to machines. Our goal was that of identifying an acceptable balance between automatic disclosures and the desire for preventing unintended disclosures.

3.2.1 Results

Some key results from our initial investigation of Reno were that people used location as a proxy for many other things and that automatic features, if they are employed at all, have to be designed very carefully.

[2]In the following version, this feature was renamed *waypoints*.

Pilot users found that when the sender and receiver of location disclosure shared context—such as that shared by a husband and wife—the location disclosures were interpreted in a much richer way than at face value. For example, if a sender sent a disclosure that he was at the bus stop, the recipient would use knowledge about the day of the week, time of day, plan for the day, and typical travel times to interpret that message as "I'll be home in 15 minutes." We saw several examples where recipients of disclosures made significant "semantic jumps" to conclusions that were at some distance from the actual text disclosed, including some jumps that were in error!

The feedback about the trigger feature was mixed, although this was clouded by significant numbers of false positives. Oftentimes, trigger disclosures were not accurate because the user had simply transited close to a trigger place, and not actually entered it. Although there were some positive user comments about the use of triggers, it was clear that the application's detection accuracy was a problem.

3.3 Investigating Specific Security Questions

The observations we made in the previous studies informed a third, more comprehensive study involving the deployment of a modified version of Reno with two families, each composed of two parents, teenage children and friends of the latter. This study, executed in the late Fall 2004, set out to investigate two specific issues, prompted by observations we had made during the previous studies:

- The usage of the application within deception and denial practices common to mediated social communication.
- The usefulness of simple automatic functions within social location disclosure applications.

Social psychologists as well as political scientists observe that a small amount of unaccountability actually improves the effectiveness and subjective management of everyday social relations (*e.g.*, when engaging in potentially stigmatizing activities such as visiting a psychologist or more mundanely, when organizing a surprise party for a friend). Designing a support for deception and denial is essential for the acceptability of communication media, as plausible deniability hinges upon such "slack space." In turn, the ability to deceive or deny, and the prevention of abuse in these practices, depend on specific security and privacy features of the medium, including access control on location information, preventing unauthorized data disclosure, and auditing system performance.

In this updated version of Reno, in addition to *Waypoints* (which are equivalent to the triggers in the previous prototype) we introduced the *Instant Reply List*. This feature causes Reno to reply automatically with the current most likely location to any request coming from a person on the Instant Reply List (which is a subset of the Reno contact list). If the location is undetermined, Reno transmits "Unknown Location." In order to increase the usefulness of the automatic features, the location algorithm was overhauled in the second version of the application, in order to provide more reliable and accurate sensing.

The deployment of this version of Reno lasted two weeks, and was performed with two groups, each consisting of a family of four with teenage children, and one to two friends

of the children. Participants were instructed in the use of Reno, and interviewed twice, at the middle and at the end of the study. Every other day, they were sent an email survey which asked about their use of Reno, and the replies were used to drive the dyadic interviews, along with status messages automatically sent by the phones to the researchers. The status messages included samples of sent disclosures and requests and usage statistics.

3.3.1 Results

In the ESM study, participants formulated their replies to other people in order to achieve specific social goals, and this suggested that straight location might not be the best or only choice when disclosing one's location. We were however still unsure to what degree the responses to hypothetical questions were accurate. Our observations in the deployment empirically confirmed in fact this kind of practice. Regarding denial and deception strategies, which are salient to our discussion of security, we observed few cases of straight-out deception, although the social structure and circumstances of the study (time of year) might have skewed these results, as we discuss further below. The low rate of deception notwithstanding, our participants demonstrated in the interviews that they would have been able to use Reno within denial and deception practices, thus supporting our claim that the control provided by Reno is sufficient for achieving plausible deniability.

The second aspect of location disclosure that we addressed in this study was related to the automatic functioning of the application. The participants of this third study by and large did not use automatic functions. As explanation, they cited concerns with potentially misleading their communication partners, and only secondarily privacy concerns. Moreover, we observed strong evidence suggesting that even in larger social networks automatic functions would be unnecessary in the face of loss of control.

4. DESIGN GUIDELINES

From our experience in the various phases of designing a social location disclosure application, we identified a set of design guidelines for social location disclosure applications. Although generally applicable principles are difficult to devise, especially in the case of relatively uncharted domains such as ubiquitous computing, for specific applications, thoughtfully applying narrowly defined design guidelines can help designers. The following guidelines should be viewed in this light, more as suggestions than mandatory rules.

These suggestions derive from observations and reflections on mistakes we incurred during the past year. We hope that these guidelines will help designers of novel social location disclosure applications better address the often conflicting needs of securing users, preserving their privacy, and building appealing and usable applications. A careful reader may observe that some of these guidelines overlap and that more general guidelines (*e.g.*, 4.4) should subsume the more specific ones (*e.g.,* 4.5). However, we have elected to point out explicitly these issues because we believe that they have the greatest impact on design.

4.1 Don't Start With Automation

Automatic functions that communicate on behalf of the user should not be introduced by default, but only when a real need arises.

In its latest form, Reno provided two automatic functions: an option to automatically reply to location requests coming from a predefined set of people; and sending the location of the user to someone else automatically whenever the user arrived in a pre-defined location. These features have a strong impact on user privacy, as they take control away from the user. Nevertheless we introduced them to see how people would manage them and how they would reconcile their privacy requirements with the benefit of having Reno work in the background. These two features were not used by the study participants, who cited several reasons for the lack of use. First, participants felt that the loss of intentionality occurring with an automatic location disclosure defied the purpose of the communication (*i.e.*, they attached meaning to the fact hat they were sending a message, meaning which was muddied by the application acting on its own). Second, they did not feel the need to have the application automatically reply to incoming location requests, because the number of these requests was manageable.

4.2 Flexible Replies

Users should be able to choose what the system discloses as a reply to a location request.

As hinted at above, participants used the content and the mere absence of replies for communicating. For example, not replying (both intentionally and not) was often used to signal unavailability (and in most cases it was correctly interpreted by the other party).

Control of what information is disclosed was also used to achieve communication goals. For example, in the third study, participants gave different names to the same location to provide their communicating partner with what they thought would be the most useful information (*e.g.*, "Making Lunch" and "Home" were both associated with the home of the user).

Social psychologists define the process of choosing what to tell a *selection problem*, among different, semantically correct alternatives [17]. This selection process involves, also, a judgment of what the person desires to disclose to his/her communicating partner, and this implies an impact on user privacy. This point is different from 4.4 below as flexibility does not necessarily imply deception.

4.3 Support Denial

Communication media should support the ability to ignore requests.

We observed, both in the ESM study and in the second deployment, that participants valued the ability to deny requests for their location. Among the various ways this can be achieved, simply ignoring incoming location requests was the strategy that participants most often adopted. Denial is important for availability management, as shown by common practice in telephone conversations, and effectively functions as a secondary communication channel.

This secondary role of denial demands that the judgment about denial be made on a case by case basis. This observation further questions the usefulness of the application automatically replying to requests.

4.4 Support Deception

Communication media should support the ability to deceive in the reply.

Cases of outright deception about location occurred relatively rarely, both in the ESM study and in the actual deployment; however, interviews with our participants indicated that in those rare instances, having the option to do so would be important. Our participants affirmed that they would lie about their location in order to preserve their individual privacy, or as a way of achieving certain long term, positive social effects.

While supporting deception may appear as an unethical proposition for designers to follow, we are convinced, by overwhelming literature, and by our conversations with the participants in all three phases of our investigations, that people do want to deceive, from time to time, for purposes that are subjectively relevant. Technologies that curtail this ability run the risk of not being adopted or in being used in unintended ways.

4.5 Support Simple Evasion

Designs should include the ability of signaling "busy" as a baseline evasive reply.

The formative study showed that blurring was not the deceptive strategy preferred by participants. More useful responses were generic messages such as "I am busy." In two cases of deception in the third study, the participant used "Running Errands" as a generic way of signaling being busy (the option "I am busy" was not available as response).

4.6 Start With Person-to-Person Communication

Social mobile applications should support person-to-person communication before attempting group communication.

Our experience suggests that it is important to achieve well-functioning designs for one-to-one communication before attempting to support more complex interaction. Tools designed to support group, as opposed to individual, communication involve more complex security requirements and policies; for example, managing group access control is more complicated than managing individual principals. This suggest to introduce groups as principals in the application only if really necessary.

Participants in the third field study did not feel the need for group communication features, even if they engaged in group rendezvous. It follows that the complex policies required by group access control systems may be overkill for simple applications. Accommodating the needs of 'power-users' should be left for later refinement stage.

4.7 Status/Away Messages

Provide a way of signaling availability status.

Several participants suggested to provide support for availability cues, similarly to "away" messages in email or Instant Messaging. This requirement is supported by theoretical arguments which view the management of personal privacy as a boundary definition exercise [15]. Providing an automatic "away" notification does not contrast with point 4.1 "Don't start with automation" above, as it is not as misleading as an erroneous or unintended location disclosure, because the user has to activate the away message. Also, this is different from point 4.5 "Support Simple Evasion," as away messages are communicated automatically, whereas that form of evasion requires manual action by the user.

4.8 Operators: Avoid Handling User Data

Social location disclosure applications should not be provided by centralized services.

In major world markets (*e.g.*, the EU), location information used for call routing enjoys a lighter regulation than location information used for so-called *value-added services* (which is subject to standard informed consent requirements). In the case of person finder applications, which are value-added services, when location is computed or stored on the infrastructure's side, the operator is required to comply with stricter regulation.

Calculating and storing location information on the phone simplifies information management on behalf of the operators because the operator only delivers messages to and from users. Revenue from the application is not necessarily impacted by this choice, as use can be charged by data traffic generated as opposed to on a per-use basis of the service. Furthermore, performing the calculation and storage of location information on the user's terminal increases the perception of control over their personal information.

5. HARD LESSONS LEARNED FROM EVALUATION

Below we provide some lessons learned, salient to the study of privacy and security with user studies, and, more specifically, to studies of hard-to-observe practices such as deceptive communication with technology. Moreover, we discuss how some of our process choices have impacted our observations, and how these must be taken into account when devising design guidelines and applying them to the development of novel products.

The message that we would like to get across is that in order to understand the unique challenges of the design for privacy and security, it is necessary to carefully sieve observations which pertain to different interrelated concerns (*e.g.*, effort of use, perceived reliability, social effects of mediated action, *etc.*) which can have an often subtly polluting effect.

5.1 Power Relationships

Select user groups which are likely to generate need for the privacy features.

In order to evaluate security features of an application, it is important to select user groups that will likely have a need to use them, such as groups with imbalanced power relationships. This was manifested in our deployment of the revised version of Reno with families.

The parent-teenager groups were specifically chosen to expose the tensions in parent-child relationships, which would presumably cause participants to use the denial and or deception options provided by the phones. We did not observe the amount of deception we had expected, and this led us to reconsider our assumptions about how the participants would have behaved.

While we still believe that it is necessary to choose appropriate power relationships, we now recognize that it may be difficult to characterize power relationships correctly (see below, 5.2). We suggest to make sure to expose desired dynamics, by carefully screening participants during recruitment. This can be done, for example, by employing interpersonal trust evaluation instruments.

5.2 User Characterization

Carefully characterize users and their use of security features and privacy requirements.

Attributing needs for specific security features to certain user groups based on common sense can be very misleading. With the support of social psychology literature and common sense, we assumed that deceptive practices among teens and parents would be quite prominent. When we failed to observe the amount of deception we had hypothesized, we were forced to look back and reconsider our assumptions. First, we observed large variability among teens in their use of the technology. Each user appropriated the system in different ways, and each teenager had different privacy needs which required different security strategies. For example, one participant formed a very clear mental model of the application and demonstrated the ability to actively deceive his parents about his location in order to achieve personal privacy with regard to the places he desired to go. This participant admitted that he would have never voluntarily used the Instant Reply feature, because he worried that his parents would have used it as a surveillance tool to prevent him visiting certain friends.

Some participants relied on their parents providing transportation to achieve their own social goals and thus were quite insensitive to privacy concerns. Other teenagers thought that socially induced self-restraint would have prevented abuse of the automatic reply function (*e.g.*, stalking), because the requesting party's identity would have be visible in their Reno inbox.

Thus, characterizing use of privacy and security features based on broad social groups like 'parents' and 'teens' demonstrated to be exceedingly blunt. Characterization must be more fine-grained to provide high quality results.

5.3 Account for Long Learning Curve

Plan for long deployments: application features which deal with security and privacy are appropriated late.

User studies which target the security and privacy-related features of applications are hampered by the fact that these features, deriving from non-functional requirements, tend to remain invisible until users really need them. This mundane observation has the consequence that it is difficult to define the length of a user study that reliably produces observations on the use of such features.

We found that the application we designed and tested was not fully appropriated even after 14 days of deployment. (The application was running, in the average across participants, 48% of the total study time, thus for the majority of the wake hours). Reno was arguably a simple application composed of 40–50 interaction steps (screens) accessible to the user (approx. 15,600 lines of Java code). Considerable effort went into fine-tuning the interface, which had been reviewed after the pilot study, and subject to cognitive walkthrough by two experienced HCI professionals.

Notwithstanding training and access to detailed documentation, most participants took one full week to become acquainted with the application's basic functionality, and a majority of participants never used advanced features with security implications such as the automatic features and the auditing functions.

5.4 Account for Specific Circumstances

In planning and evaluation account for circumstances of deployment.

As people's social activity and practices vary not only over the course of a lifespan but in yearly, weekly and shorter cycles, the results of user studies can be strongly influenced by the specific circumstances of the deployment. This demands careful selection of appropriate times of year for performing specific studies. In the experience with the second deployment study, for example, we observed that the tight schedules of most teens, and the time of year, during school, just before major holidays, reduced their independent mobility. This impacted our observation of deceptive practices both between peers, and with their parents.

Participants spontaneously noted that repeating the study during summer vacations could have produced very different usage patterns. Again, this issue is emphasized by the latent character of security and privacy requirements.

6. BOISE: A MAP-BASED SOCIAL LOCATION DISCLOSURE APPLICATION

After our experience with Reno, we embarked on a new design effort, focused on another social location disclosure application, Boise. Boise has been informed by the positive lessons we learned from Reno's design and deployment as well as some of our missteps. Boise, which at present is at the paper prototype stage, is motivated by four primary considerations.

First, Reno users found it to be a compelling coordination tool. Frequently, users needed to have a rendezvous of one or more people. In these scenarios, location was often exchanged via Reno with significant context overloading. For example, participant A might use Reno to send participant B a disclosure like "School." This disclosure was not meant to indicate that A was at school, but rather that A was at school waiting for B to come and pick A up in a car. With Boise, we are trying to support this type of rendezvous activity more explicitly.

Second, in addition to simply communicating location for meeting up, users desired to communicate Estimated Time of Arrival (ETA). Understanding the various people's ETA is almost always needed in rendezvous scenarios. Several participants in both deployment studies expressed that they had tried to use Reno's location functionality to assess others' ETA or disclose their own. The success of these efforts varied, and we wanted to try to make ETA a more integral part of the new application. This related interestingly to deception, as ETA provides less information about one's activity than location. One participant suggested that if he had to deceive his communication partner by making up a location, he would be careful to choose a location which would induce a credible expectation for ETA (see points 4.2, 4.4).

A third motivation in Boise is the desire to generalize the rendezvous notions in Reno from one to many. In Reno, we made an explicit choice that the communication would be peer-to-peer. This type of architecture offers significant benefits for privacy, since only the participants are involved in the communication and the control of each disclosure is made with respect to a single recipient, as synthesized in point 4.6 above.

Some participants, though, faced complex communication tasks, such as getting an entire family together for an event. In both families we observed, the mother was in charge of

coordinating family activity and was the primary means by which the children without cars were ferried around. Often these plans required several stops to pick up and drop off people who wanted to go to different places. Reno offered little support for this multi-person problem, and even if it were to be extended to include a notion of groups of other users, problems would remain.

While participants of the third study did not ask for supporting group communication, we recognize that in larger groups (more than 4 people) coordination may become overly problematic. Given that Reno uses SMS as message transport, some users may be deterred by transmission costs, since multiple messages would need to be sent in a peer-to-peer fashion. Furthermore, recipients might face a high cognitive barrier in trying to coordinate multiple text messages to achieve a complex coordination task. Thus, in Boise, we decided to support group communication, and to explicitly test for potential privacy issues arising in group communication.

The fourth and last consideration which influenced Boise related to the presentation of information and interaction style. Several users during the third study had suggested the use of maps instead of text messages as a more natural way of communicating ETA, route plans, *etc.* Boise allows users to visualize the location of the people they are coordinating with on a map. By visualizing the location of other people, we felt that users could use their knowledge of local geography to approximate the ETA of persons involved in a coordination activity. Our hope is that by putting special care in the design of the application we will prevent the somewhat justified "big brother" reaction some users might have when seeing their location on a map.

6.1 Design Proposal

Boise runs in one of three modes: normal mode, tracking mode, and away mode. These modes are described below, with reference to the design guidelines provided above.

6.1.1 Normal Mode

In normal mode, Boise offers users a way to query other users for their location, and to respond to the queries, similarly to Reno. Unlike Reno, however, Boise does not have a location inbox but simply gives users a notification that a query has been received, from whom, and at what time. This decision was based on the relatively low volume of queries observed in the Reno studies and the observations that Reno's participants tended to either attend to those messages immediately or ignore them if they were too old (and thus, not salient anymore). The notification disappears after a while, and then, the location of people is displayed by positioning an icon on the map (in Figure 2, a smiling face). The icon slowly fades away over time, so that older location disclosures appear thinner, and eventually disappear. We have not yet decided the time needed to completely fade the icons.

Like Reno, Boise allows users to disclose their location at a time and place of their choosing. Disclosures are addressed to a single recipient, similar to Reno, as suggested by point 4.6. The design of Dodgeball and the need for managing rendezvous with several people, have induced us to try to design in a buddy list that allows the user to manage who will receive disclosures. We hope to test buddy lists in upcoming field studies to see if there are cases in which users would prefer to disclose to multiple buddies with one action.

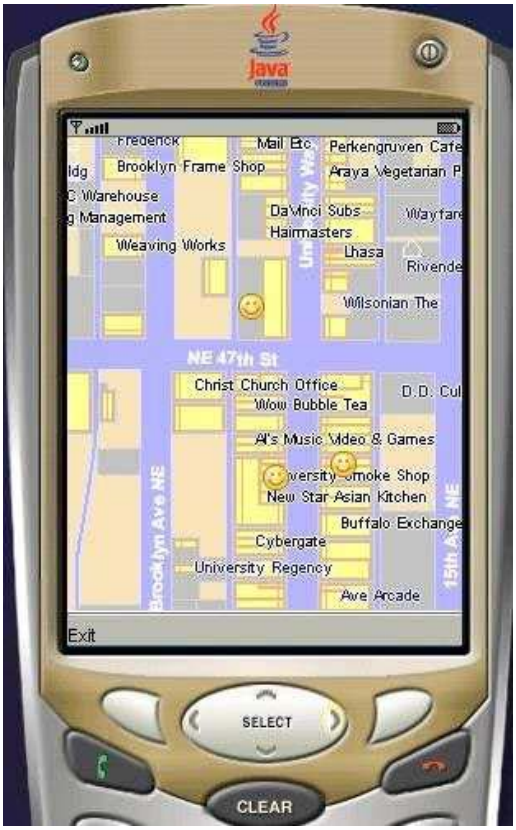

Figure 2: Prototype implementation of Boise. The owner of the device is located at the center of the image (at the *Wow Bubble Tea Shop* in Seattle's university district). Three other people (the smiling faces) have disclosed their location to the user. The background map was taken from Lost In Seattle. (www.lostinseattle.com)

Boise's disclosure options are more subtle than Reno's simple text messaging model. One way to disclose your location in Boise is to move the name of a place to the center of the display, like *Wow Bubble Tea* in Figure 2. These place names could be defined by the user, like in Reno, or more simply seeded with Yellow Pages or directory listings. However, place names are not associated with cell towers, but with the geographical location of the places (see below for a discussion of the location technology employed). The selection of the location to disclose is up to the user and independent of his/her actual location.

Once at the center, the business name will highlight and it can be selected. A set of crosshairs is used to show the user the center of the screen and these are not shown in Figure 2. The disclosure is then sent to the intended recipient and the user's icon is placed on the recipients' map. At present, we have not yet designed the interaction steps for choosing disclosure recipients.

Another way to disclose location is for the sender to scroll the map so that desired location appears roughly at the center of the screen under a crosshair, not shown in Figure 2. This causes that location to be disclosed (in the form

of geographical latitude/longitude coordinates), again independently of the user's actual location.

6.1.2 "I'm Running Late" or Tracking Mode

In some cases, users of Reno expressed a desire to have other users know and track their current location. This was commonly expressed in situations where one party was running late for a rendezvous, or when they were engaged in a series of activities throughout the day and wanted to provide a status indication, and is confirmed by other user studies on similar applications (*e.g.*, [7]). When the user enables tracking mode, the recipients of the disclosures get frequent updates about the sender's physical location, allowing them to track the sender, *e.g.*, coming down the freeway toward their location.

Giving up their true location information may be viewed by the discloser as a reasonable loss of privacy for avoiding the other person unnecessary waits. Further, in Boise the sender must express when the tracking mode is to be terminated. By default, the termination condition occurs when a specified amount of time has elapsed. Another termination condition is when the sender comes in close proximity to a designated recipient—the rendezvous has been accomplished. The former condition is intended to support social availability, such as with a group of friends out for the evening who might want to rendezvous. Dodgeball supports the opposite scheme, by allowing users to *block* communication for the remainder of the evening [6].

A key innovation of Boise's tracking mode is that the user's location is snapped to locations that they have disclosed in the past. For example, when a user in tracking mode transits an area where the user had once disclosed a place called "Pabla Indian Cuisine," the location disclosed to other recipients might be "Pabla Indian Cuisine" for the entire period that the sender is in the proximity of that place. We plan to test the effectiveness of this strategy, since it negatively impacts some rendezvous activities by giving out less than the software's best estimate of the user's location. On the other hand, it protects the user from disclosing places that may carry a negative connotation. This behavior was inspired by point 4.5, as the system only discloses place names which are supposedly harmless because they had been already previously disclosed.

The ability of tracking the actual location of users places a higher burden on the location system used by Boise. In Reno, since users were compelled to name their own places, the application could use a class of algorithms referred to as *fingerprinting* algorithms, where a fingerprint of recent cell towers is associated with a place name. Since tracking mode requires the application to calculate the actual location of the user, Boise must use a different algorithm, which can estimate the user's location without the need for fingerprints. Boise uses Place Lab to accomplish this [10]. This algorithm relies on maps being created a priori of the radio environment of the entire extent of the usable area of the software.

6.1.3 Away Mode

The third operation mode in Boise is the away mode. In this mode, Boise provides an automatic away message, when someone asks for the user's location. Similar to what happens with Instant Messaging applications, away mode allows users to signal their unavailability to location requests, without specifying any reason; this provides a secondary communication channel (the person is not available for initiating communication or meeting up) without sacrificing privacy by telling the requester where the person is or what he/she is doing.

6.2 Critique

In this section we review each of the design guidelines provided in section 4 and assess how design choices made in Boise might be affected by the guidelines.

6.2.1 Don't Start With Automation

The only automatic feature in Boise is the tracking mode, which arises from a very specific and loudly voiced need of supporting ETA and rendezvous activities. As mentioned above, Dodgeball allows users to disable the automatic notifications for a specified amount of time, whereas our design is more conservative. As suggested by point 4.1, Boise automatically terminates tracking mode after a default timeout to prevent accidental disclosures if the user forgets to turn it off.

6.2.2 Flexible Replies

In the current version of Boise, users can disclose a named place or a geographic location, which arguably does not provide great flexibility, and could overly constrain how people would use the tool.

In response to this guideline, we are currently evaluating alternatives for enriching communication, such as allowing the party who is disclosing their location to select a specific icon to indicate his/her activity or availability in addition to or instead of the location. This icon would then appear on the recipients' screen and support secondary-channel communication.

Another design option would be that of allowing a user to 'tap' other people by hitting their icon on the map. The other person's phone would then vibrate shortly, thus providing a subtle and tactile communication mechanism.

6.2.3 Support Denial

In Boise, denial strategies are arguably limited. Basically, the user can only choose not to reply to a message.

6.2.4 Support Deception

In Boise, the user can select what location to disclose to others, independently of his/her actual location, both when disclosing a place or geographical coordinates. We feel that this support for deception is sufficient, and that it would be unbalanced to enable deception during tracking mode. Further, it would defeat the point of determining (presumably accurately) when the other party will actually arrive at a rendezvous location.

6.2.5 Support Simple Evasion

Currently Boise does not fully comply with this guideline. The user can elect not to reply, but not to provide an evasive answer. However, we are considering to enable Boise users to reply with a busy message instead of a location or a place name, *e.g.*, by providing a specific icon—see point 6.2.2 above. In light of the ESM study results we do not think that supporting evasion by changing the scale of the location (*i.e.*, telling the city as opposed to the street name) is necessary.

6.2.6 Start With Person-to-Person Communication

In Boise, disclosures are addressed to a single recipient, similar to Reno, as suggested by point 4.6. As mentioned above, however, we are planning to try a buddy list that allows the user to send at once disclosures to everyone on the buddy list.

The need for simplifying interaction by providing easy group support does compete with the privacy requirements embodied in point 4.6. In fact, experience in the field of usability of security functions suggests that in most cases, users might not be willing to manage complex buddy lists. The simplest form of group support would thus be to have only one buddy list, which presents on the other hand clear problems if people want to partition their social milieu.

In observance of the guideline, if we implement buddy lists, this feature will be subject to specific tests during the summative evaluation of the application, in order to understand if it causes acceptance problems and if it is really needed.

6.2.7 Status/Away Messages

The away mode of Boise was designed specifically to account for point 4.7 above.

6.2.8 Operators: Avoid Handling User Data

To calculate user location, Boise uses Place Lab, which computes the location without assistance from the operator network. Communication occurs through regular SMS messages.

To reduce the traffic of SMS messages in the case of large buddy lists, one could imagine a centralized server storing the buddy lists, and forwarding Boise traffic to the entire buddy list, similar to how Dodgeball works. We are still evaluating techniques to protect user data while incurring the least burdensome regulatory constraints in this case. A simple approach in this case would be for all the buddies to share a common secret key and some cryptographic technique. This makes the server less interesting to attackers since it would not keep trace of users' locations.

6.3 Evaluation

We are acutely aware of the fact that usability of Boise is a key factor in the success of the application—in addition to the privacy issues discussed in this paper. In the specific case of Boise, one can easily imagine that family and friends might be spread over a large geographic area, such as the entire Puget Sound region, the entire San Francisco Bay area, or greater London. Understanding the displays presented by Boise on a screen limited to 160 by 200 pixels is a significant challenge. We are working with visualization design experts to design effective-for-small-screen map-based displays that preserve our privacy design features. These will be tested in controlled, laboratory settings before being deployed. These visualization experiments are outside the scope of this work and will be reported elsewhere. Here, we ignore the issue of scaling maps.

Once the visualizations have been fully designed and validated by user testing, we plan to do a field study to study Boise in action. Our current design for the field study is to employ groups that need to rendezvous, particularly for social reasons. To this end, we are seeking a few (2–5) groups of people who have reasonably large social circles (5–7 people) which require significant amounts of coordination. We plan to deploy Boise using the visualizations obtained from the lab tests as key parts of its user interface.

We are interested in two different kinds of coordination activities. In order to obtain high quality data, we plan to employ different test groups for Boise, and will screen them to expose some of the social dynamics we are interested to study. This procedure is suggested by point 5.2 above. In the first type of coordination, a large family will use Boise to solve the pragmatic coordination problems often observed in modern families [4]. In this type of coordination, a tool like Boise will be judged primarily on its effectiveness and speed. We hope to find that using Boise offers a significant advantage to a family, perhaps in terms of less waiting time, fewer late/missed appointments, fewer phone calls, or the ability to successfully attend to even more coordination activities.

The other type of coordination activity that we would like to study is more social, more opportunistic, and less clearly useful. These types of activities are those performed just for social benefit—such as going out to dinner with friends. We hope to recruit one or more sets of participants who would use Boise as part of their social activities, preferably with all or most of the social group using Boise. In such a deployment, Boise would be deemed a success if it opened up new (or better) opportunities to socialize with friends, rather than the more utilitarian metrics offered by the family deployment.

Given the potentially long learning period for this application, we plan to deploy Boise for a time longer than 2 weeks (the length of the third study reported above)—see point 5.3 above. Cost concerns may limit the duration of deployment studies. However, in our experience, most of the cost of a deployment is incurred during the planning and development phases, and not during the actual deployment. On the other hand, long deployments may be problematic for recruitment, requiring higher compensations and incurring in higher recruitment costs. Long deployments also require to hedge against usability issues or malfunctions in the software. Segmented deployments may provide the best results. In these deployments, participants use the application for a certain amount of time. The application is then fine-tuned, and showstopper bugs are fixed. The same set of participants then use the application again for the remainder of the study. This reduces deployment costs related to learning.

With regards to the family deployment, we would choose, if possible, to deploy Boise during the summer vacations, in order to leverage the higher mobility of the participants, and potential deceptive practices related to their specific activities outside of the home (see point 5.4). In relation to the deployment with a group of friends, we plan to enroll a group of college students or young professionals during the normal working season (*e.g.*, early university semester or quarter).

7. CONCLUSION

Over the past year, we have been investigating the use of social location disclosure applications and have iteratively developed a design by employing a host of formative and summative methods. This experience has allowed us to identify a number of guidelines which help designers in building location disclosure services which are more sensitive to user privacy, and especially support the essential space for plausible deniability within social interaction. These guidelines

are divided into two groups: design guidelines and methodological comments related to the evaluation of social location disclosure tools. These guidelines are specifically targeted at pointing out privacy concerns in this kind of applications, and should be viewed as necessary but not sufficient for achieving a successful design or deployment.

Applying design guidelines has been traditionally problematic in the mobile and ubiquitous computing field, which suffers from the lack of an established design practice. We tried to show how to apply these guidelines to a novel design and to demonstrate how they can inform the development of a new application. We indicated where the guidelines came into play in the new design, both as justifications for design choices, or as warnings that design choices made for satisfying other requirements (such as group communication support) might cause acceptance problems.

The guidelines and the demonstration of how they are applied represents in our opinion the major contribution of this work. We hope that by providing these guidelines to designers, we can improve the design and quality of future social location disclosure applications and services.

8. ACKNOWLEDGEMENTS

This extensive work would have not been possible without the collaboration and help of a very large group of people, at Intel Research, Georgia Tech, the Intel PAPR group and elsewhere.

We especially thank Jeffery Hughes, James Howard, Jeffrey Hightower, Anthony LaMarca, Tara Matthews, Wendy March, Fred Potter, Pauline Powledge, Tony Salvador, James Scott, Tim Sohn, John Sherry, Jason Tabert (Intel) and our cognitive walkers Gillian Hayes (Georgia Tech) and Jehan Moghazy (IBM) for their collaboration on this project and thoughtful comments. We also thank our anonymous reviewers and all the participants of our user studies.

This work was funded in part by Intel Corp., the NSF through the Graduate Research Fellowship Program and Georgia Institute of Technology. Human subjects research is in part covered by Georgia Tech IRB protocol H04232.

9. REFERENCES

[1] Barkhuus, L., Dey, A. Location-Based Services for Mobile Telephony: a study of users' privacy concerns. In *Proc. Interact 2003,* 709–712, IOS Press, Amsterdam, 2003.

[2] Brown, K. On The Trail Of Location Services. *Wireless Week*, p. 18, March 1, 2004, Reed Business Information.

[3] Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J., Powledge, P. Location Disclosure to Social Relations: Why, When, & What People Want to Share. In *Proc. of the Conference on Human Factors and Computing Systems: CHI 2005,* 81–90, ACM Press, 2005.

[4] Darrah, C.N. Family Models, Model Families. In *Revisiting the American Dream: How U.S. Families Cope with Work and Family Life*, American Anthropological Association Annual Conference, Chicago, IL, November 23, 2003.

[5] DePaulo B.M., Kashy D.A. Everyday Lies in Close and Casual Relationships. *Journal of Personality and Social Psychology*, 74 (1), 63–79, American Psychological Association, January 1998.

[6] Ubiquity Labs, Dodgeball. Available online at http://www.dodgeball.com. Last accessed: Feb 17, 2005.

[7] Fithian, R., Iachello, G., Moghazy, J., Pousman, Z., Stasko, J. The Design and Evaluation of a Mobile Location-aware Handheld Event Planner. In *Proc. of the Fifth International Symposium on Human Computer Interaction with Mobile Devices and Services* (Mobile HCI 2003), Udine, Italy, Sept. 8–11, 2003, LNCS 2795, 145–160, Springer Verlag.

[8] Grinter, R. E., Eldridge, M. 'y do tngrs luv 2 txt msg?'. In: *Proc. Seventh European Conference on Computer-Supported Cooperative Work* ECSCW '01, 219–238, Kluwer Academic Publishers, 2001.

[9] Ito, M. and Daisuke, O. Mobile Phones, Japanese Youth and the Replacement of Social Contact. In Ling, R. and Pedersen, P. (eds.) *Front Stage/Back Stage: Mobile communication and the Renegotiation of the Social Sphere*, Conference Proceedings, 22–24 June 2003, Grimstad, Norway.

[10] LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., Schilit, B. Place Lab: Device Positioning Using Radio Beacons in the Wild. In *Proc. Pervasive 2005*, Springer Verlag, 2005.

[11] Langheinrich, M. Privacy by Design: Principles of Privacy-Aware Ubiquitous Systems. In *Proc. Ubicomp 2001*, LNCS 2201, 273–291, Springer Verlag, 2001.

[12] Laurier, E. Why People Say Where They Are During Mobile Phone Calls. *Environment and Planning D: Society and Space*, 19, 485–504, Pion, London, UK, 2001.

[13] Lederer, S., Mankoff, J., Dey, A. Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In *Proc. of the Conference on Human Factors and Computing Systems: CHI 2003*, ACM Press, 2003.

[14] Ling, R. *Norwegian teens, mobile telephony and SMS use in school.* Technical Report Telenor R&D N 7/2000, Telenor, 2000.

[15] Palen, L., Dourish, P. Unpacking Privacy for a Networked World. In *Proc. of the Conference on Human Factors and Computing Systems: CHI 2003*, 129–136, ACM Press, 2003.

[16] P&AB: Consumer Privacy Attitudes: A Major Shift Since 2000 and Why. In *Privacy & American Business Newsletter*, 10 (6), Sep 2003.

[17] Schegloff E.A. Notes on a Conversational Practice: Formulating Place. In Sudnow, D. N. (ed.), *Studies in Social Interaction,* 75–119, MacMillan, The Free Press, New York, 1972.

[18] Smith, I., Consolvo, S., Hightower, J., Hughes, J., Iachello, G., LaMarca, A., Abowd, G.D., Scott, J., Sohn, T. Social Disclosure Of Place: From Location Technology to Communication Practice. In *Proc. Pervasive 2005*, Springer Verlag, 2005.

[19] United States Department of Health, Education and Welfare. *Records, Computers and the Rights of Citizens.* Report of the Secretary's Advisory

Committee on Automated Personal Data Systems, 1973.

[20] Weilenmann A. "I Can't Talk Now: I'm In A Fitting Room:" Formulating Availability And Location In Mobile Phone Conversations. *Environment and Planning A*, 35, 1589–1605, Pion, London, UK, 2003.

[21] The Yankee Group. *2004 Mobile Users Survey*. Synopsis available at http://www.yankeegroup.com. Last accessed on Feb. 17, 2005.