Social Disclosure of Place: From Location Technology to Communication Practices

Ian Smith, Sunny Consolvo, Anthony Lamarca, Jeffrey Hightower, James Scott, Timothy Sohn, Jeff Hughes, Giovanni Iachello, and Gregory D. Abowd

Intel Research Seattle,Intel Research Cambridge,
University of California, San Diego,
College of Computing & GVU Center, Georgia Institute of Technology, Atlanta, GA
{ian.e.smith, sunny.consolvo, anthony.lamarca, jeffrey.r.hightower, james.w.scott}@intel.com
tsohn@cs.ucsd.edu
{giac, abowd}@cc.gatech.edu
jeffdh@cs.washington.edu

Abstract. Communication of one's location as part of a social discourse is common practice, and we use a variety of technologies to satisfy this need. This practice suggests a potentially useful capability that technology may support more directly. We present such a social location disclosure service, Reno, designed for use on a common mobile phone platform. We describe the guiding principles that dictate parameters for creating a usable, useful and ubiquitous service and we report on a pilot study of use of Reno for a realistic social network. Our preliminary results reveal the competing factors for a system that facilitates both manual and automatic location disclosure, and the role social context plays in making such a lightweight communication solution work.

1 Introduction

In many situations we use location, more specifically a notion of place, to communicate with friends, family and colleagues. People send postcards, teasing friends with catchy phrases like "Life's a beach in Biarritz" or "Thinking of you at the Grand Canyon." Even more often people exchange location information with people in their social network, using a variety of technologies, such as phone conversations, SMS, instant messaging, and email. User studies of SMS by teenagers in England [15] and Germany [18] agree that the top three uses are 1) to keep in touch with friends and acquaintances; 2) to coordinate and schedule physical encounters and phone calls; and 3) to chat. At least the first two of these uses suggest the use of location within the content of SMS messages.

Our focus is on an emerging class of pervasive computing application, the *social location disclosure application*, motivated by the explicit sharing of location information in social communication. One example of an existing social location disclosure application is Dodgeball [4]. Dodgeball helps cell phone users meet up with other members of their social network in the physical world, with a particular focus on the

H.W. Gellersen et al. (Eds.): PERVASIVE 2005, LNCS 3468, pp. 134 – 151, 2005.

serendipitous encounter. Although Dodgeball has demonstrated value—it has approximately 1500 daily users—the space of social location disclosure applications is larger than just rendezvous scenarios.

Due to the large amount of shared social state between friends, co-workers and family members, the small amount of information contained in a place name is often sufficient to stand proxy for a wide variety of communications. We want to understand how people leverage social context using lightweight communication tools to share location. Furthermore, the emergence of technologies that allow mobile devices to learn and recognize their location reduces the overhead of sharing such information. This creates the opportunity to develop an application that allows users to exchange location information, manually or automatically, with other members of a social network. The design challenge is to create such an application with low interaction cost and high availability, without sacrificing the privacy or social control individuals want.

In this paper, we introduce Reno, a social location disclosure application that allows users to send their current location, manually and automatically, to other people in their social network as well as request the location of others. Reno runs on Nokia Series 60 mobile phones (such as the Nokia 6600) and uses cell-tower based location estimation to aid the user in defining and recalling places as well as triggering automatic disclosures. We describe our design, as well as the accompanying principles that helped us create a service that is both useful and respectful of people's privacy. We present the primarily qualitative results of a pilot user study of seven early adopters using Reno over the course of five days.

2 The Intended User Experience

The following scenario, adapted from a reported use in our pilot study, helps explain the basic capabilities of Reno. Figures 1 and 2 show screenshots of the user interface (UI) on a Nokia 6600 running the Reno application.

It's near the end of the workday, and as Phoebe prepares to go home, she wonders what she and her husband, Ross, will do for the evening. Ross should be heading home now too, but Phoebe thinks he may be in a meeting, so she doesn't want to phone him to ask. Instead she pulls out her mobile phone and, with a few quick clicks, sends a location query to Ross. Ross, in fact, is busy and working later than normal today. He sees Phoebe's request (Figure 1A), but chooses not to answer it.

His meeting ends about 45 minutes later, and while he is waiting at the bus stop near his office, Ross accesses Reno on his phone to send a location update to Phoebe (Figure 1B). The screen on his phone lists 4 place names that correspond roughly to his current location. His bus stop is outside Merchant Mick, a favorite store where he often buys dessert on his way home from work. When Ross sees that place name at the top of the list, he selects it, figuring that Phoebe will understand that he has left work and will be getting on the bus soon. Phoebe receives the message at home (Figure 1C) and is now eagerly anticipating Ross' arrival home.

When Ross gets off the bus, just outside the Easton Hotel, a location update is automatically sent to Phoebe, (Figure 2A) and she knows that he is only 10 minutes from home, so she sets out dinner just in time so that it is a warm welcome home for her husband. When Ross gets home, he smiles at the set table and says, "You are *wonderful! I am late, and you already have dinner waiting for us!" Phoebe replies, "I made your favorite dish to go with my favorite dessert you just bought from Merchant Mick's." Ross, realizing that Phoebe misinterpreted the Reno message he had sent from the bus stop and feeling a bit embarrassed, admits that he didn't buy dessert tonight, because he was already running late. Seeing the disappointment in Phoebe's eyes, he decides to take her to Merchant Mick's for dessert after dinner. On the way home from dessert, standing at the same bus stop, Ross pulls out his phone and defines a new place in Reno that connects that place with the label "Bus stop by work," (Figure 2B), hoping to avoid the same confusion in the future.*

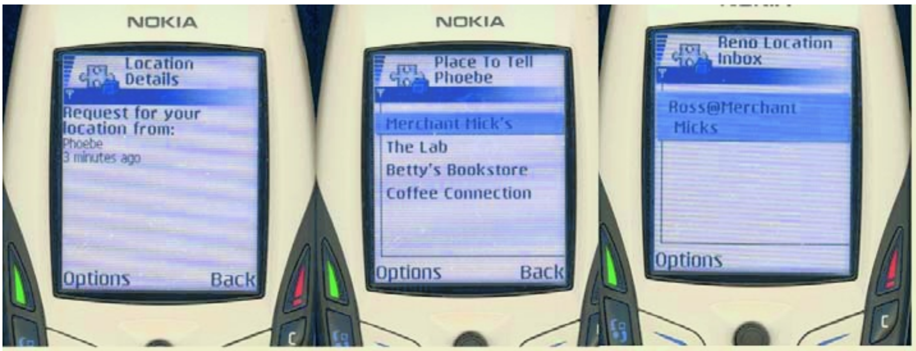

Fig. 1. A, Left) Phoebe's request for location received by Ross. **B, Center)** Ross' list of nearby places he can send. **C, Right)** View of Phoebe's inbox after she receives Ross' location

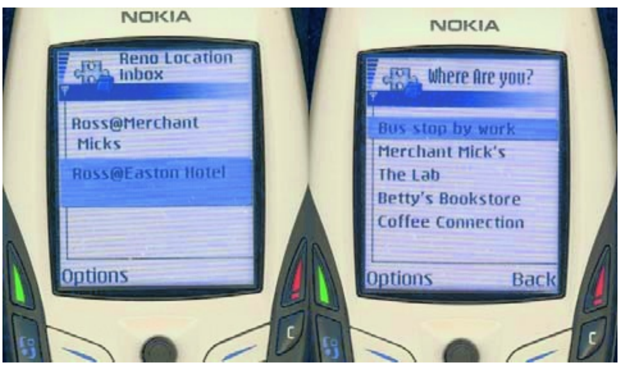

Fig. 2. A, Left) Phoebe's location inbox after receiving Ross' automatic disclosure. **B, Right)** Ross' list of nearby places after adding the bus stop

This scenario demonstrates four of the basic capabilities we built into Reno:

- explicitly and manually "push" a location disclosure to someone;
- automatic triggering of a location disclosure to someone else based upon entry or exit from a pre-defined and agreed upon location;

- explicit request of "Where are you?" to someone, with a subsequent, but not required location disclosure from that person; and
- ability to define new places that can be used for future location disclosures.

To allow the user to see queries for disclosures of their location, we use a common concept: an "inbox." The location inbox (see Figures 1C and 2A) contains a list of recent *queries*, marked with a '?,' and *disclosures*, sent manually or automatically, marked with an '@' character. These items can be "opened" with the phone's soft keys or joystick to see details; a detail view of a query is shown in Figure 1A. Query details include the time it was sent and the name of the requester; disclosure details include the time the disclosure was made, the full name of the place, and whether it was a manual or automatic disclosure.

In addition to clarifying how Reno works, this scenario also demonstrates how the communication of a place name is interpreted, sometimes incorrectly, within the social context of human discourse. It is precisely this leveraging of social context that makes Reno an intriguing research platform. While we are unable to provide a complete analysis of the implications of adopted Reno use over an extended period of time, our initial deployment study does reveal an interesting balance between the four functional capabilities and the social context of its use.

In the following section, we will review the guiding principles for designing, implementing and deploying Reno. We will then report on a pilot user study of its use over a 5-day period within a specific social network.

3 The Reno Application: Guidelines, Design and Implementation

Reno was designed with three guiding principles in mind: always-on service, avoid real or perceived privacy threats, and minimize deployment or acceptance barriers. By *always-on service*, we mean that Reno must be available 100% of the time, since otherwise users would not consider it robust enough to employ it for their day-to-day tasks. *Avoid real or perceived privacy threats* refers to the notion that users must be able to trust that Reno is not disclosing their location without their explicit consent, and it should avoid reliance on third parties who may not be perceived by users as trustworthy. To *minimize deployment or acceptance barriers*, Reno must be convenient for users, it must be low-cost and it must avoid relying on users to carry around hardware which they would otherwise not carry. It also should not interfere with the other hardware and services users already rely on. This also means that there should be minimal reliance on external infrastructure. Even though this infrastructure might improve the functionality of Reno, it presents obstacles to deployment that had to be avoided.

Certain principles were *not* prioritized in Reno. Location accuracy is useful, but wide coverage and availability is more important. Complex features (e.g., automatic detection of nearby "friends") might prove useful in some situations, but without a clear understanding of how basic social location disclosure is used in practice, it is unwise to design for those features at this point.

The principles above were used to make four key design decisions about Reno's platform, location sensing, place representation, and message transport. Regarding the message transport, Reno communications are entirely peer-to-peer, using the Short Message Service (SMS). This transport method was chosen because it is widely available, interoperates across mobile phone networks worldwide, has a clear cost model, and is already familiar to many users. The infrastructure that implements SMS is under the control of service providers and regulated by various national and international bodies. More importantly, the casual user of SMS understands it to be a person-to-person and private communication. Leveraging SMS for message transport, therefore, means that no real or perceived privacy threats are *introduced* by Reno, respecting the second and third of our guiding principles.

3.1 The Mobile Phone Platform

The mobile phone platform was chosen primarily because it is the most powerful computing platform which is already continually carried by many people. Requiring that a Reno user carries a PDA or other hardware would create deployment barriers, and threaten to make the service unavailable when users forget or do not bother to carry the extra hardware with them.

Reno is targeted at the Nokia Series 60 phone platform and is written primarily using Java 2 Micro-Edition (J2ME), with the Connected Limited Device Configuration (CLDC) and Mobile Information Device Profile (MIDP) APIs. All development and deployment was done using the Nokia 6600 phone, chosen for its large screen and good developer support. Disadvantages of this platform include the difficulty of debugging and the limited and slow persistent storage support.

The mobile phone platform satisfies some of the guiding principles mentioned earlier, but brings to the forefront a deployment barrier. Mobile telephony is a crucial function for many people—particularly for some of our pilot users with school-age children—so we were very cautious to avoid degrading the phone functionality of the device, e.g., by draining the battery or crashing the phone. Note that, in order for Reno to be able to send automatic location disclosures, it must always be running as a background task on the phone.

3.2 Location Sensing

There were a number of possible sources for determining the phone's location. The most widely commercialized location sensor, GPS, was rejected because it is not always available (indoors, in buses, etc), and the cost and extra hardware introduce deployment barriers. Another possibility would be to leverage location technology which is being implemented by some telecom operators to comply with "Enhanced 911" requirements [12]. Although network-based sensing is a compelling business case for operators, we have chosen to compute location on the handset because of the lack of standard methods for querying location from the network, and because we feel that this solution provides for simpler regulatory compliance in jurisdictions with strong location privacy legislation (*i.e.*, private use of others' personal information is exempt from most data protection provisions). Computing location on the phone also respects the user's control for disclosure.

A more promising option is Place Lab [2,21], which uses the known locations of visible radio beacons such as GSM cell towers and Bluetooth beacons to provide a current location estimate. However, Place Lab relies upon a user community having generated detailed "maps" of the radio beacons in a given area before location sensing works in that area. This goes against the principle of always-on service: even if large maps were gathered, users may move outside the mapped area. If, in the future, such mapping databases are well populated[1], a client-based location service such as Place Lab would be suitable for Reno.

Reno keeps track of the phone's currently associated cell tower, and each place defined in the phone includes a list of cell towers which have been seen while at that place. Location is determined entirely on the phone and does not rely on any special hardware or external service. This method has very high availability, introduces no new privacy threats, and does not hinder deployment, thus making it very suitable for a quick and realistic study. The biggest disadvantage is the relatively low accuracy of this technique, since cell tower footprints range from hundreds of meters to several kilometers.

Bluetooth beacons can be used in a similar way to GSM cell towers, potentially providing more accurate location information. However, we decided not to use Bluetooth in this study, for two reasons. Many Bluetooth nodes are mobile, and associating mobile Bluetooth beacons with a place will lead to false detections of that place. However, there is a deeper problem because there are adverse privacy implications in the (even unwitting) recording of the location and movements of other users carrying Bluetooth devices. While there are ways around this (e.g., allowing the users to specify particular Bluetooth beacons as being "stationary" and therefore usable), in the interests of simplicity and usability by novices, we did not do so in our initial study.

3.3 Representing Place

There are many ways to represent places. For use by humans, expressing places with mathematical coordinates or beacon IDs is not acceptable, regardless of how convenient it would be to implement. Furthermore, previous studies have shown that even symbolic names like exact street addresses are infrequently a user's first choice when communicating her location to other people [7]. As a result, we decided not to use pre-defined place lists (e.g., extracted from the web) in Reno, and instead required users to define their own places. This allows users to mark places with very personal labels (e.g. "10 minutes from home", "Phoebe's office"). Places in Reno consist of a user-defined name and a list of associated unique cell tower IDs, as shown in Table 1.

Table 1. Example of Place Definitions

Place Name	Cell Tower List
Ralph's Market	2561, 2221
Home	923
Bank	12087, 8921, 12071

[1] http://www.wigle.net is an example of a community-generated database which currently maps 1.8 million WiFi access points; phones with WiFi functionality have been announced by a number of manufacturers and may become commonplace.

A place is created when the user activates the “Record this place” feature. While platform restrictions mean that Reno is only capable of monitoring the currently visible cell tower, a list of all recently seen cell towers is maintained. When the user records a new place, this list is associated with that place. The user is also able to “Record this place” for already defined places; in such cases, the list of currently seen cell towers is merged into the place’s existing list by taking the intersection of the two sets. Empirically, we have found a time limit of about two minutes for recently seen cell towers works well. This allows some smoothing of the GSM beacon environment seen by the mobile phone, while avoiding the accumulation of old GSM beacons.

4 Evaluation and Results

4.1 Management Burden for the List of Nearby Places

When a user selects “Tell someone I am here”, a minimal list of nearby places is shown. We aimed to show the most probable nearby places in order of relevance to avoid requiring the user to scroll through a long list of places that have no current relevance. To validate whether this happens in practice, an experimental “walk” was conducted by one of our pilot users around a local shopping district. Using our nearby places algorithm, we sampled the returned list of results at each of the pre-defined places (triangles) and measurement (circles) points shown in Fig. 3.

In every case where the user was actually present at a place, the place received a non-zero score meaning it would have been on the user’s list of nearby places. When the user was at a measurement point, the closest place was given as either the highest or jointly highest in each case. The highest or jointly highest rank place was given by the algorithm when at 6 of the 8 places. In the two cases where the algorithm did not put the user’s true place as the highest scoring choice, it made significant errors in selecting the best choice; an error of about 350m at place 5 and an error of 550m when the user was truly at place 6. Further, the algorithm clearly offers no discrimination power between places 2, 3, and 4. These places are roughly all on the same block (distance north to south between place 2 and 4 is 205 meters). A check of the user’s phone revealed that all these places were defined with the same, single cell.

During real use, we have found that our cell-tower-based algorithm works very well and has performance characteristics very similar to those shown in this section. One confounding factor in the experiment above involves Reno’s cache of recently seen beacons. The current Reno system has no model of the speed of travel of the device it is running on. This means that the two minute decay of beacons has a very different effect when the device is in a car and when the user is walking, affecting both the definition of the place and the measurement of the environment to compute nearby places.

4.2 Pilot User Study

A pilot study of Reno was conducted in Fall 2004 with eight participants over five days. The participants were members of the research team and their families. We

Fig. 3. Beacon Based Place Experiment. This figure shows an area 1170 meters wide by 1520 meters tall

chose this set of participants because they were part of an actual social group, yet would be forgiving of problems with both the technology and unintended and potentially awkward location disclosures to members of their social network. This group of participants provided a reasonable first pass at uncovering some of the social ramifications and privacy concerns that a social location disclosure application like Reno could create among co-workers, friends and spouses. A primary goal of this pilot study was to refine the study design and Reno for a longer term study with participants who are representative of the target user base.

We begin by describing the profiles of the eight participants and follow with the methodology used in the three phases of the pilot study. We then discuss key results.

4.2.1 Study Design

Participant Profiles. Eight participants (four male), aged 26-40, participated in the pilot study. Participants were a mixture of Reno project members and their families (see Fig. 4 for a representation of the network). Five were employed full-time, worked together, and were members of the Reno project team (participants a, d, e, g, and h in Fig. 4). Four of those five worked in the same office; the other (h) worked overseas but communicated regularly with the four via phone, instant messaging, and email. The remaining participants were family members of participants a and g: two spouses and one sister-in-law. One of the spouses was a part-time designer (f); the other spouse (b) and the sister-in-law (c) were homemakers.

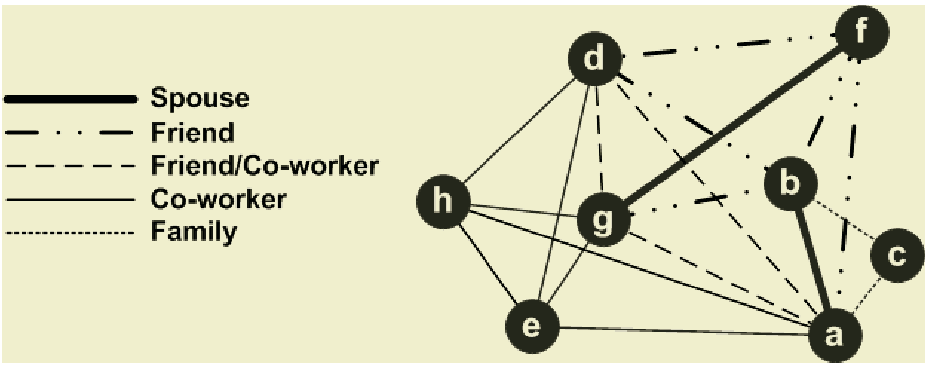

Fig. 4. Representation of the participants' network. Participants are represented by circles; letters are used to identify individuals. Lines show social relations

We used the P&AB-Harris Interactive privacy classification survey [28] to determine if participants were privacy *fundamentalists*, *pragmatists*, or *unconcerned*. Four participants were fundamentalists, three were pragmatists, and one was unconcerned. The trend we saw is slightly different than those reported by P&AB-Harris Interactive, with fundamentalists being the largest group in our pilot study, followed by pragmatists, then unconcerned (in 2003, P&AB's largest group was pragmatists, followed by fundamentalists, then unconcerned).

Pilot Study Methodology. The pilot study consisted of three phases: background, Reno deployment, and post-study feedback. In Phase 1, participants signed a release form, completed a couple of questionnaires, and worked with us to customize their Reno application. The questionnaires included basic demographic information, experience with a variety of technologies, the P&AB privacy classification survey, and exercises to explain how they knew and communicated with the other participants in the pilot study. To customize Reno, they each provided at least two triggers they wanted another participant to receive from them when they arrived at or departed a particular location and at least two notifications they wanted to receive about another participant. (They also provided a set of names for places they thought would be important during the five day deployment of Reno; the combined list of names from all participants was used to create a list of "Word Cheats." These cheats were simply a list of places that could be relied on to avoid typing in text on the phone, a task that several of our pilot users were unfamiliar with.) Once we received the trigger and notification requests, we asked the appropriate participants to accept/reject the requests.

For the Reno deployment (Phase 2), we supplied participants with Nokia 6600 cell phones with the Reno application pre-installed. Based on the customization exercises from Phase 1, their phones were already set up with the requested triggers that had been accepted. Participants were instructed on how to use the application, including how to "teach" their phone about places. Despite the triggers being set up in advance, each participant was responsible for teaching his phone about the relevant places (*i.e.,* even if participant *h* had a trigger set to notify participant *g* whenever *h* arrived at work, if *h* never taught his phone where work was, the notification would never be sent to *g*). Though we provided participants with cell phones, they used their personal SIM cards[2] and existing service.

During the Reno deployment which ran for five days from Wednesday through Sunday, participants filled out a daily email log that included several questions about their experiences with Reno and the other participants for that day. This set of questions was emailed to them every morning. They were asked to fill it out and send it in before going to bed that day. Participants received a reminder every evening. If they had not returned their daily email log by the next morning, they received a reminder to please fill it out and return it as soon as possible. The study concluded by returning the Nokia 6600 cell phone and filling out a post-study questionnaire which asked about their general reaction to Reno and their experiences during the study. Usage logs from the phones were extracted and analyzed.

4.2.2 Results: Basic Usage Results

The pilot study lasted for 5 days and we expected that the device would be with the person and on the vast majority of the time. Despite giving the participants fully configured phones and letting them keep their existing phone number, we met with mixed success on this point. We had an average of 53% for the amount of time the phone was on as a percentage of the 5 day (140 hour) study time, with a minimum and maximum of 22% and 85% respectively.

Fig. 5 shows the number of automatic and manual disclosures that were made by each participant over the course of the study. The total number of disclosures was 306 and an average of 38.25, or about 7-8 per day of the study. Fig. 5 shows that there is a significant difference between the systems heaviest and lightest users; although the two heaviest users shown, Ross and Joey, are developers of the system, their spouses, Rachel and Phoebe, used Reno a great deal as well, including a signifycant number of manual disclosures.

We expected that the amount of location disclosure would be correlated with the strength of a social relationship—the strongest of these being the spousal relationship. 45% of all requests for location (21 total requests) were sent to a spouse/partner. Similar requests to a colleague, a friend, and a "friend & colleague" were 23% (11), 17% (8), and 15% (7), respectively. Considering all location disclosures (automatic and manual) the percentage distribution between spouses/partners, colleagues, friends, and "friends & colleagues" was 44% (124 total disclosures), 34% (98), 4% (11), and

[2] One participant did not have a cell phone, so we provided her with a phone and service.

16% (46), respectively. The spousal relationship did have the highest amount of location disclosures, but our data, both on real strength of relationship and amount of disclosures, do not allow us to draw any further conclusions.

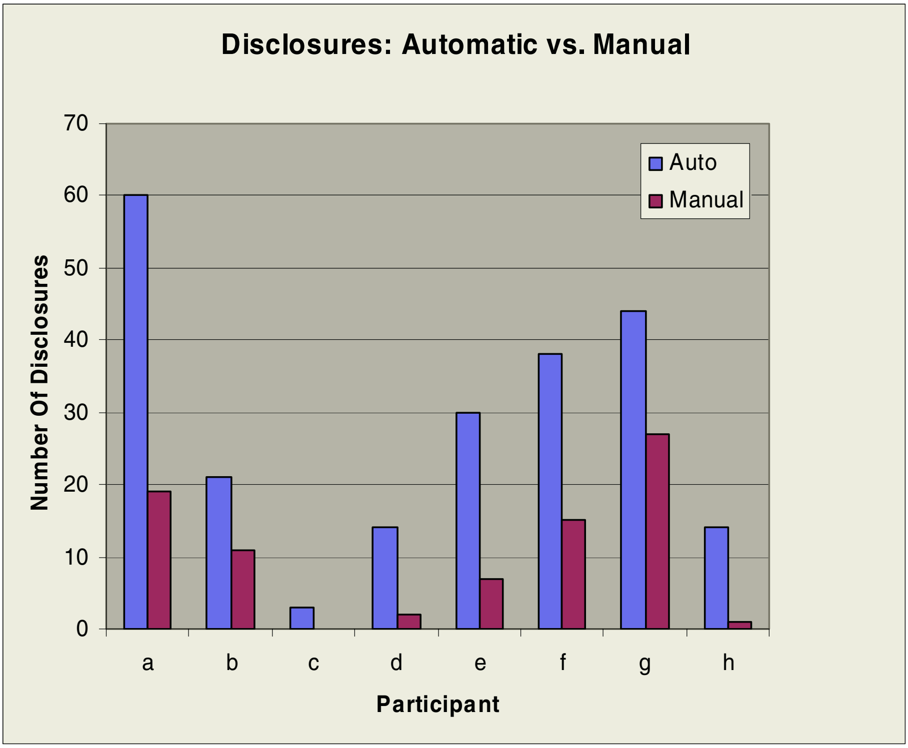

Fig. 5. Number of automatic and manual disclosures, by participant

4.2.3 Results: People's Context Makes Place/Location Communication Effective

Location information as we have defined it in Reno, is a relatively low data bandwidth channel. Reno makes it easy to share the user's location, either automatically or through user interaction, but what is shared is basically a few characters of text. Since we supplied users with a list of place names to avoid typing, very few places had names that were "customized" for the particular user, e.g. "Merchant Mick's" was chosen although better names might have "the Merchant Mick's on the 33rd street at 4th avenue" or "Merchant Mick's near my house." Despite this relatively impoverished channel, a great deal of information flowed between the senders and receivers. Our pilot users had social relationships and these are relationships that are quite rich; the richness allowed a simple string to take on a great deal of meaning.

Quote 1: Day 2 (Thu), participant a: *Late tonight, [f] pushed her location to me from <anonymized> airport. That reminded me that some of her friends (also* *acquaintances of mine) were coming into town today. I wanted get together with them socially when they arrived. I had forgotten it was this weekend.*

Quote 1 is completely dependent on the fact that *a* knew that *f* had friends coming into town and approximately when; the sender and receiver had significant shared context allowing what would otherwise be a seemingly random disclosure at a strange time to be interpreted correctly. It should also be clear from this example that the sender clearly *knew* that the receiver would interpret the message through the shared context.

Quote 2: Day 5 (Sun), participant f: *I knew [g] was coming home from work and when he was at McDonald's.*

Participant *f* in **Quote 2** is referring to two disclosures received from *g*, both of these involved *g*'s path. In the first case, an automatic disclosure revealed to *f* that *g* was nearly home from the office since the location was a place that would be passed on the route home. The latter was a manual disclosure that told the recipient that the sender would be home shortly with a fast-food dinner. In both cases, the context of the two participants allowed a simple place name to be taken for both a path as well as an approximate estimated time of arrival. Without shared context information, either of these could easily be misinterpreted by the receiver; for example, a different take on the latter disclosure is "He's telling me his diet is out the window."

We did see a number of instances where misinterpretation did occur. The types of misunderstandings of sender or receiver's context varied as did the consequences of the confusion. The scenario discussed in section 2 was based on an actual misunderstanding that occurred between *f* and *g*.

Quote 3: Day 1 (Wed), participant f: *"I saw that [g] was at <Merchant Mick's> but he was really waiting for the bus, and I was disappointed."*

As was discussed earlier the sender and recipient did not agree on the significance and meaning of the place in question. However, it is interesting to note that *neither* of the two participants took the place at face value. One was expressing the activity "waiting for the bus" and the other thought "buying dessert." We also note that the exact same message being exchanged between *a* and *d* of the same social group at lunch time between Monday – Thursday would mean, "If you want me to pick up a salad for you, call me" and on Friday during summer between any of *a, d, e,* or *g* would mean, "If you want me to pick something up for you for the barbeque, call me."

Quote 4: Day 1 (Wed), participant e:.*I knew [a] was on the east side and I wasn't sure how close he was to being back. He replied with 'Home Depot' … I wasn't sure which Home Depot.*

After following up with *a* and *e* about the situation, we learned that *e* wanted to know which Home Depot because he thought *a* was trying to give him an indication of how close to the office he was, *i.e.,* that he was driving past a Home Depot, and *e* didn't know if he was at a nearby or east side Home Depot. However, what *a* was trying to communicate was "I'm done with work, so I'm taking a trip to Home Depot to run a personal errand. I'm finished with work for the day." Again, although this pair got their signals crossed, *neither* participant thought that this name would communicate the simple idea of *a*'s location!

We found that the location disclosures, both automatic and manual, were intepreted through the lens of the context shared between the sender and receiver. Since manual disclosures were sent *intentionally* by the sender, they were often able to achieve a higher level of abstraction, such as in **Quote 1**. The automatic disclosures, since they had no significance to the sender tended to be less significant to the receiver. Overall, we found that there were three types of interpretations commonly seen. These were, in increasing order of the use of the shared context, location, path, and activity. The first of these occurred most frequently when the sender and receiver were trying to rendezvous and the location of the sender could be interpreted literally. More frequently, the receiver used knowledge of the place transmitted and the local geography to understand the path of the sender, as in **Quote 2**. The last class of interpretations, activity, were quite amazing and varied dramatically in their level of understanding (or misunderstanding) as was shown in **Quote 1** and **Quote 4**.

4.2.4 Results: Automatic Disclosure is a Double-Edged Sword

There is a cost and a value to automatic disclosure. In our pilot study opinions varied significantly about the value of the automatic disclosures. **Quote 5** and **Quote 6** show that even within a single participant, opinions varied!

Quote 5: Day 2 (Thurs), participant a: *I had an 'anti-meeting' situation today in which Reno was VERY helpful. When [b] is out for the night I always let the kids stay up later than I should and she always gets peeved when she come home and finds them not in bed. I got my Reno alert that she was home and since she had to get stuff out of the car and walk up the flight and a half of stairs, I had the kids in bed before she opened the door. So Reno saved me some irritated spouse today.*

Quote 6: Day 3 (Fri), participant a: *When [b] and I went to our house together, we get cross-notifications, neither of which is useful at the time.*

The important point is that both were situations where *a* and *b* were in the same place—in the negative example, he wants Reno to not give the notification because it is useless. In the positive example, *a* is glad to receive the notification because he has a few moments to "clean up" his mess. Though *a* alludes to the fact that it might be when they arrive together being the problem, it's easy to imagine a situation in which *a* is not doing something wrong but is home before *b* is, receives the notification, and finds it useless. The key thing was that *a* was doing something he was not "supposed" to be doing. If *a* had been behaving properly so to speak, the notification would have been irritating, as in the negative example. The double-edged sword of automatic disclosures would have been even more pronounced had we been able to explain **Quote 5** to *b* and gauge her reaction!

Another example of the mixed feelings engendered by the participants can be seen in **Quote 7**. This quote shows that the participant did not fully understand the context in which this trigger might fire and when confronted with the unpleasant, privacy-invasive use of the trigger felt violated.

Quote 7: Day 2 (thurs), participant g: *My phone disclosed my location to [participant a] last night when I was out running an errand and was returning to my house ([a] has a trigger for my house). I'm now reconsidering that trigger because I felt* *weird about that one. I didn't feel weird when he got notified I was "home from work" but did about the late-night errand running."*

5 Related Work

Location-based services (LBS) have been hailed by telecom operators as the next killer app following the unexpected success of SMS [29]. However, marked success for LBS has lagged, in part due to lack of a clear regulatory framework and privacy concerns[3]. Japan has witnessed some success in this area: person-to-person LBS have been marketed by DoCoMo in the form of location-augmented iMode websites (used for, e.g., dating) [26]. Ubiquitous cell-phone tracking was commercially launched by KDDI with GPS MAP, a person finder service targeted at corporate customers, in 2002 [19]. This system works with GPS-enabled handsets that can be tracked through client software on PCs. Location information is stored on a server system managed by the operator.

A more general cell-phone based person-finder application has been developed by Kivera Inc. for AT&T Wireless (now part of Cingular). The system, called Find People Nearby (formerly known as Find Friends) [6] allows a user to build a buddy list, and to locate other service subscribers in any area covered by AT&T. After locating a friend, the user can call the person, send a message or invite him/her to some meeting point chosen from businesses in the AT&T Yellow Pages. The success of this application has been arguably limited by the lack of interoperability with other providers; Reno overcomes this limitation by leveraging the universal interoperability afforded by SMS. Another difference is that the Find People Nearby application is web-based: the location of the phone is sensed by the infrastructure and the application is accessed through a web browser on the phone, along with other AT&T-provided services. Reno computes user location on the device; given the different legal status of location information used for call routing as opposed to value added services in major world markets (e.g. the EU [1]), this simplifies information management on behalf of the operators and increases user control on their personal information. Finally, users of Find People Nearby cannot label locations: location names are predefined and typically refer to urban coordinates (e.g. street intersections).

Recently consumer-oriented person-finder applications (e.g. dating and child-tracking applications) have started to appear amidst mounting privacy concerns and a clamp-down on websites linked to high-profile 'horror' stories of abuse [3]. Several child tracking applications are currently available in the United Kingdom. For example, Mapamobile allows to track a phone with the precision allowed by cell-tower coverage. [24] As opposed to Reno, these services are marketed as single-purpose applications, are centrally managed, and subject to strict organizational oversight.

Apart from dating and child tracking, applications of location technology include mobile guides and games. In the Lighthouse location-aware museum guides by Galani and Chalmers [14], mobility is a resource for interaction as participants use motion to communicate between each other, to signal presence and to support contextual awareness. We observed that our participants also used location information as a resource for interaction, e.g. for planning, communication and coordination activities.

Mobile location-aware devices serve not only explicit interaction, but also more subtle purposes. In the WatchMe prototype [25] GPS location and activity information are used to provide users with awareness of each other, and to initiate explicit communication. Our observations indicate that participants use Reno for similar purposes, especially when location is disclosed automatically in "push mode" (through triggers). When a disclosure occurs, people reconnect the location information with prior knowledge of the other's activity; location disclosure thus acts as a reminder and awareness mechanism of the whereabouts of the other person. Finally, like WatchMe, Reno is build on a ubiquitous, always-available device.

In 'Uncle Roy All Around You' users self-report their location (by pin-pointing it on a map on a touchscreen. Interestingly, Benford et al. noted that, similarly to how Reno users use location to express activity and share context, players in their game used self-reported location to express intent (where they will be going) and history (places they had visited) [10].

In location-based games 'Can You See Me Now?' [13] and George Square [11], mobility also represents a tool for interaction. These systems use both GPS and WLAN beaconing to enrich the gaming experience and the imprecise nature of this kind of location sensing is exploited by the designers to enrich the game by creating uncertainty. Imprecision and ambiguity are very important in Reno, as they afford an essential space for privacy. Thus, the imprecise nature of cell-phone tower-based localization can be viewed as both a problem and an advantage.

Laasonen *et al.* show how cell-phone tower localization can be used in combination with user-based labeling schemes [20]. They also use predictive algorithms to recognize when a user is in a certain *base* (i.e., a recurring location such as workplace and home), or in *areas* (which are clusters of bases, *e.g.*, a neighborhood) and *routes* between bases. A different approach to movement prediction is followed by Ashbrook and Starner, who propose associating place labels to GPS information and Markov models to infer the user subsequent movements [5]. Reno uses cell phone tower location because such information is more available in buildings and dense urban environments than GPS readings. Despite the name, GPS is less ubiquitous than cell-phone technologies in places of human significance, because where there are humans there is an interest in providing cell phone coverage (e.g. subways, buildings, factories).

Harrison and Dourish have brought to the attention of the ubicomp community the long-standing distinction between space and place (i.e., a location associated with a meaningful semantic) [17]. Adopting their observations, Reno hides the spatial coordinates from the user. Rather, it provides tools for creating and using meaningful place names, instead of hard-to-understand geographical or urban coordinates.

The view of interpersonal privacy as a continuous negotiation process and a tool of social action dates back several decades. Goffman, for example, characterizes privacy as one of the fundamental instruments that enable people to create their social presentation [16]. Palen and Dourish have adapted this concept of privacy to information technology, and suggest that IT must provide the tools for finely managing this negotiation in order to preserve our ability to function as social beings [27]. This kind of negotiation clearly emerges from a web-based survey performed by Lederer *et al.* [23] in 2003, which shows that people tend 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. This is confirmed by a precursor study to the present [7], which, in addition, indicates that the supposed reason for the request is a co-determinant for deciding whether to disclose, whereas current activity and mood have only secondary influence on the decision. This study also highlighted that based on that decision, users provide either the information that they think will be most useful to the requester, or none. The present work is directly influenced by these findings: Reno is designed to support fine-grained control on the projection of their social image.

Fine control over information disclosure has been the topic of much work in the security community; more security implies increased administrative burden. This balance has been inquired specifically for context-aware technology by Barkhuus and Dey, who have suggested that people are willing to forgive some control over their personal location information [8]. Our experience with Reno shows that users display similar feelings towards automatic disclosure of location information — what these authors term “active context-awareness”. Feedback and control of information disclosure are central to the ubicomp community from the very beginnings of the field [9]. Recently Lederer *et al.* have stressed that successful designs must make information flows visible, provide coarse-grain control, enable social nuance, and emphasize action over configuration [22]. Reno’s design confirms these guidelines; when Reno does not perform as expected by its users, it can be attributed to one of these pitfalls.

8 Conclusions

The social communication of location or place information is common, and we often co-opt a variety of technologies to facilitate this practice. In this work, we explored what it would mean to provide a purpose-built social location disclosure service that met certain criteria for rapid deployment and adoption. We introduced Reno as a prototype location disclosure service for a mobile phone platform, and we presented some significant results from a pilot user study of its use within a realistic social network.

While our implementation of Reno was not as available and transparent on the mobile phone as we would have liked, it did provide a simple means of communicating user-defined place information, both automatically and manually. When combined with the rich, shared social context among family, friends or colleagues, this simplified disclosure of place facilitates effective communication. The automatic disclosure of location, while at times valuable, suffers because the explicit communication act by the sender, and its accompanying knowledge of intended context for interpretation, is lost.

The advantages and disadvantages hinted at by our pilot study suggest deeper research questions which we intend to explore next. Specifically, we want to better understand whether unauthorized or unintended location disclosure will be restrained to an acceptable rate by a combination of algorithmic techniques, user interaction and pressure from social norms. Another important feature for any technology supporting social interaction is the extent to which it supports the human need for plausible deniability. Adoption of Reno, or of any similar location disclosure application, will engender denial practices, attributable to user activity, the supporting infrastructure, or even the operational or failure modes of the service. Understanding these denial and restraint strategies, and supporting them explicitly in new technologies, will be essential condition for the acceptance and, ultimately, the success of pervasive technologies.

Acknowledgments

The authors would like to acknowlege the support of the National Science Foundation's Graduate Research Fellowship Program and the MacArthur Foundation for their support (Iachello). We would also like to acknowledge the Intel corporation's support of academic collaborations (Iachello, Abowd). This work would not have been possible without the excellent technical work provided by James Howard, Jeff Hughes, and Fred Potter.

References

1. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) Official Journal of the European Communities (2002) 37-47
2. Placelab.[Online]. Available: http://www.placelab.org
3. What the Operators Are Doing. in Mobile Location Analyst, Oct 2003.
4. Dodgeball. Ubiquity Labs LLC, [Online]. Available: http://www.dodgeball.com/. Downloaded January 26, 2005.
5. Ashbrook, D., Starner, T.: Learning Significant Locations and Predicting User Movement with GPS. In: Proc. Int. Symp. on Wearable Computers (2002) 101-108
6. AT&T Wireless. Find People Nearby.[Online]. Available: http://www.attwireless.com/personal/features/organization/findfriends.jhtml
7. Consolvo, S., Smith I., Mathews T., LaMarca A..: Location Disclosure To Social Relations: Why, When, & What People Want To Share. To appear: ACM Conference On Human Factors In Computing Systems (2005), Portland, Oregon.
8. Barkhuus, L., Dey, A.: Location-Based Services for Mobile Telephony: a study of users' privacy concerns. In: Proc. Interact 2003 (2003) 709 -712
9. Bellotti, V., Sellen, A.: Design for Privacy in Ubiqui-tous Computing Environments. In: Proc. ECSCW '93 (1993)
10. Benford, S., Seager, W., Flintham, M., Anastasi, R., Rowland, D., Humble, J., Stanton, D., Bowers, J., Tanadavanitj, N., Adams, M., Farr, J. R., Oldroyd, A., Sutton, J.: The Error of Our Ways: The Experience of Self-Reported Position in a Location-Based Game. In: Proc. Ubicomp 2004 (2004) 70-87
11. Chalmers, M.: George Square -- Demo Game at Mobile HCI 2004 Conference. (2004)
12. Federal Communications Commission. Enhanced 911.[Online]. Available: http://www.fcc.gov/911/enhanced/
13. Flintham, M., Anastasi, R., Benford, S. D., Hemmings, T., Crabtree, A., Greenhalgh, C. M., Rodden, T. A., Tandavanitj, N., Adams, M., Row-Farr, J.: Where on-line meets on-the-streets: experiences with mobile mixed reality games. In: Proc. CHI 2003 (2003)
14. Galani, A., Chalmers, M., et al.: Developing a mixed reality co-visiting experience for local and remote museum companion. In: Proc. HCI international, HCII2003 (2003)

15. Grinter, R. E., Eldridge, M.: 'y do tngrs luv 2 txt msg?' In: Proc. Seventh European Conference on Computer-Supported Cooperative Work ECSCW '01 (2001) 219–238
16. Goffman, E.: The presentation of self in everyday life. Anchor Books, New York (1959)
17. Harrison, S., Dourish, P.: Re-place-ing space: The roles of space and place in collaborative systems. In: Proc. CSCW (1996) 67-76
18. Höflich, J. R., Rössler, P.: Mobile schriftliche Kommunikation – oder: E-Mail für das Handy. Die Bedeutung elektronischer Kurznachrichten (Short Message Service) am Beispiel jugendlicher Handynutzer. Medien & Kommunikationswissenschaft 49 (2001) 437
19. KDDI. GPS MAP, a Location Service For Mobile Phones.[Online]. Available: http://www.kddi.com/english/corporate/news_release/archive/2002/0718/
20. Laasonen, K., Raento, M., Toivonen, H.: Adaptive On-Device Location Recognition. In: Proc. Pervasive 2004 (2004) 287-304
21. 21 LaMarca A., Chawathe Y., Consolvo S., Hightower J., Smith I., Scott J., Sohn T., Howard J., Hughes J., Potter F., Tabert F., Powledge P., Borriello G., Schilit B., Place Lab: Device Positioning Using Radio Beacons in the Wild, To appear: 3rd International Conference On Pervasive Computing (2005), Munich, Germany..
22. Lederer, S., Hong, J., Dey, A., Landay, J.: Personal Privacy through Understanding and Action: Five Pitfalls for Designers. Intel, IRB-TR-03-035 (2003)
23. Lederer, S., Mankoff, J., Dey, A. K.: Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In: Proc. CHI 2003 (2003)
24. Mapamobile. [Online]. Available: [http://www.mapamobile.com/].
25. Marmasse, N., Schmandt, C., Spectre, D.: WatchMe: communication and awareness between members of a closely-knit group. In: Proc. CHI 2004 (2004)
26. NTT DoCoMo. iArea: Location Based Services.[Online]. Available: http://www.nttdocomo.com/corebiz/imode/services/iarea.html
27. Palen, L., Dourish, P.: Unpacking "Privacy" for a Networked World. In: Proc. CHI 2003 (2003) 129-136
28. Privacy & American Business: Consumer Privacy Attitudes: A Major Shift Since 2000 and Why. 10 (2003)
29. UMTS Forum: Enabling UMTS / Third Generation Services and Applications, UMTS Forum Report 11. (2000)