

Context-aware Telephony: Privacy Preferences and Sharing Patterns

Ashraf Khalil

College of Computer Science and
Information Technology
Abu Dhabi University
Al Ain, United Arab Emirates
khalil.ashraf@gmail.com

Kay Connelly

Computer Science Department
Indiana University
150 S. Woodlawn, Ave
Bloomington, Indiana, 47401
connelly@indiana.edu

ABSTRACT

The proliferation of cell phones has led to an ever increasing number of inappropriate interruptions. Context-aware telephony applications, in which callers are provided with context information about the receivers, has been proposed as a solution for this problem. This approach, however, raises many privacy issues that may render it infeasible. In this paper, we report on an in-situ study of user privacy preferences and patterns of sharing different types of context information with different social relations. We found that participants disclosed their context information generously, suggesting that context-aware telephony is not only feasible, but also desirable. Our data shows a distinct sharing pattern across social relations and different types of context information. We discuss the implications of the results for designers of context-aware telephony in particular and context-aware applications in general.

Categories and Subject Descriptors

H.5.2 User Interfaces- user-centered design, H.5.3 Group and Organization Interfaces – evaluation/methodology, collaborative computing; K.4.1 Public Policy Issues – Privacy.

General Terms

Human Factors.

Keywords

Context-awareness, mobile phones, interruptions, experience sampling, social relations, information disclosure, privacy.

1. INTRODUCTION

The explosion in the number of cell phones has made them the most ubiquitous communication devices [27]. Consequently, cell phone interruptions have sparked problems that affect not only the owner, but people in the surrounding environment as well [12]. The problems, which have already been studied extensively, are many in number and can be serious in nature [6, 20, 31].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW'06, November 4–8, 2006, Banff, Alberta, Canada.

Copyright 2006 ACM 1-59593-249-6/06/0011...\$5.00.

Context-aware telephony has been proposed as a way to minimize cell phone interruptions that are mostly caused by the mismatch between the user's context and the cell phone settings [18, 26, 28, 32]. This mismatch is largely due to the static nature of cell phone configuration, which depends on the user's memory to change the configuration every time the context changes. One approach to solving this problem is to empower the caller to make better, more informed decisions about the appropriateness of making a call by providing him information about the receiver's context. Context information can be any kind of information that helps in conveying the receiver's availability and circumstances, such as location, activity, ambient sound and social cues such as company. A recent study showed that providing the caller with contextual information about the receiver indeed decreases the frequency of the mismatch and enhances the level of agreement between receivers' desires and callers' decisions [1]. The study also showed that different contextual information generates different levels of improvements.

This approach, however, raises many privacy issues that need to be answered before judging its feasibility and practicality. For one, are people willing to disclose context information in exchange for less inappropriate interruptions? In other words, are the incentives and services offered by the context-aware telephony application good enough to overcome the privacy concerns? What kind of context information are they willing to share? How does the type of context information affect sharing patterns? What is the effect of the caller-receiver relationship on context disclosure? How do people differ, and on what do they agree?

By addressing these issues, we are not only trying to examine the feasibility of context-aware telephony but also aim to provide valuable insight to designers of context-aware telephony applications. Only by understanding the sharing patterns of different contextual information with different social relations and user privacy preferences can designers formulate efficient, user-friendly and privacy-aware systems with improved tools for privacy management. The above issues are not unique to context-aware telephony that deals with interruptions, but are common to all context-aware telephony applications that require context sharing, such as those that support social interactions [8, 17].

We report a formative study with 20 participants in order to address the above questions. This study was conducted in-situ over a period of 10 days where participants received inquiries asking them what kind of context information they were willing

to disclose to potential callers from 6 different types of social relations. Interviews were conducted at the end of the study as part of an overall evaluation.

The remainder of the paper is divided as follows: We first review the various context-aware telephony solutions that have been proposed in the literature as well as related privacy studies. Then we describe the study design including details about the different types of context information we consider. We detail the main findings, followed by a discussion of the design implications. Finally, we discuss the relation between context disclosure and privacy and then conclude.

2. RELATED WORK

2.1 Context-Aware Telephony

Context-aware telephony solutions that aim to minimize interruption can be divided into two categories. The first includes approaches that empower cell phone owners by improving the capabilities and awareness of cell phones. Quiet Calls enables the caller to send prerecorded audio messages thus enabling him or her to negotiate with the caller a more appropriate time for the call [19]. Such a system may decrease interruptions for the surrounding environment but does not affect personal interruption since the user is still expected to receive the call and act upon it. SenSay is a system that uses input from different sensors such as accelerometers, light detectors, and microphones to capture the context of the user [30]. The context is then used to adjust the modality of cell phone notification (i.e. vibration, ringer). Schmidt et al. have introduced an adaptive cell phone that changes its configuration automatically based on the recognized context [28]. Khalil and Connelly investigated the use of calendar information to automatically adjust the configuration according to the user's scheduled activities [13].

The second category is to empower the caller to make better decisions about the appropriateness of the call before making it by providing her information about the receiver's context [18, 26, 32]. This approach leverages human judgment in determining the appropriateness of making a call given the current social environment of the receiver. The Calls.calm system uses the web to activate an interaction webpage that provides the caller with a set of available communication channels as well as information about the receiver's current situation, and leaves it up to the caller to make an educated choice [26]. Milewski and Smith have used the address book to display dynamic information about the recipient's availability and location [18]. Their solution applies the same concepts of "Buddy List", used in instant messengers. Marmasse et al. have developed a wristwatch prototype that captures the user's context and communicates this information to the members of a closely-knit group [17]. This information can be used by the caller to assess the availability of the receiver before attempting the call. Our study investigates the feasibility of the second category of context-aware telephony applications through studying privacy preferences and sharing patterns of users.

2.2 Privacy Studies

Extensive research has been conducted with the aim of studying privacy issues related to context-aware computing and CSCW. The study most relevant to ours is the one conducted by Consolvo et al. which explored user's location disclosure to social relations [3]. They found that location disclosure depends mainly on the

recipient, the degree of detail, and the reason for the disclosure request. The study also examined the decision process that determines whether and what to disclose to social relations. Our study goes beyond location information and examines the disclosure of other kinds of context information as well as the relationship between different types of context information.

Olson et al. studied the sharing of private information with the purpose of identifying clusters of information and recipients in order to create a simple and efficient privacy management system [22]. The study focused on personal information rather than dynamic context information. Patil and Lai investigated sharing preferences for location, availability, calendar information and messaging activities for a collaborative application called MySpace [25]. The study explored how people control their privacy permission settings for social relations. They found that grouping mechanisms offer a balance between privacy control and configuration burden. Our work is not aimed at exploring privacy permission but rather at sharing patterns for a different set of contextual information that is more relevant to mobile users. Lederer et al. presented a mechanism to allow people to control the disclosure of their context information [16]. They introduced the "face" concept as a metaphor to what the user is willing to show others in a certain situation. Their user study found that the inquirer plays a larger factor than the situation when people decide what to disclose. Danezis et al. examined the extent to which location information is valued by people [7]. They used economical incentive to study how much people are willing to receive before they release their location information. In the case of context-aware telephony, the incentive for releasing contextual information is the benefit of having fewer inappropriate interruptions. One main purpose of our study is to explore whether such an incentive is valued enough by users for them to disclose their contextual information.

Several studies have aimed at exploring the privacy requirements in context-aware applications. Lederer et. al have identified five pitfalls that designers should avoid when designing interactive systems [15]. These pitfalls include obscuring potential information flow and lacking coarse-grained control. Reciprocity was introduced by Bellotti et al as an important privacy control feature in context-aware applications [2].

We share with Palen and Dourish the same conceptual framework for privacy, which states that privacy management is not about setting static rules and enforcing them but rather about managing boundaries that are continually in tension with each other [24]. Palen and Dourish identified three interrelated boundaries that characterize privacy management: the disclosure boundary, the identity boundary, and the temporal boundary. The goal of our study is not to set rules regarding sharing of contextual information but rather to examine the viability of the caller-based approach in enhancing the accuracy of cell phone interruptions by investigating the extent and the patterns of shared contextual information.

3. STUDY

We performed a study that focuses on exploring the sharing patterns and privacy preferences for different types of context information and for different social relations in the domain of context-aware telephony. In order to avoid measuring privacy in isolation of the natural setting, we used the Experience Sampling

Method (ESM). First introduced in the mid-70s, this method is simply an automated version of time diaries [5]. It is now widely used in ubiquitous computing to simultaneously and systematically study people in natural environments [4]. Our choice to use the in-situ approach rather a survey was further driven by many recent studies that have found significant discrepancy between privacy attitudes and behavior. For example, Jensen et al examined privacy practices of internet users and found that users' behaviors did not match their survey responses [11]. They concluded that surveys are best suited to study attitudes and opinion rather than behavior or experience and called for a reevaluation of the role of surveys in studying privacy behavior.

3.1 Context Types

We identified four different types of contextual information relevant to mobile telephony: Location, Activity, Company, and Conversation. These four types were specifically chosen for the balance they provide between privacy concerns that are commonly considered high risk factors (namely, Location and Activity) and low risk factors (Company and Conversation), as well as for their relevance to context-aware telephony. In addition, capturing the contextual information of these four types of context is technically feasible, given the recent advances in context-aware computing and sensor technology. Below we describe the four contexts in more detail.

- **Location:** Location is one of the most widely used types of context, as demonstrated by the numerous location-aware applications on the market. Advances in location tracking systems such as GPS and GSM have made it feasible to obtain location information. According to the location, people can often infer the activity and thus the availability of the receiver. Location is especially relevant to mobile telephony, considering that many people start their phone conversation with the question "Where are you?" [14].
- **Activity:** Information about a person's activity, like location, provides a direct clue as to his or her availability. Many systems have been developed in order to infer the user's current activity and to predict future ones. These systems use cues such as calendar information, computer activities, and real-time analyses of audio and video streams to infer the current activity [9, 21]. For example, the SEER system uses models to identify a set of human activities through computer activity, ambient audio, and a video stream [21].
- **Company:** This type of context information provides knowledge about the presence of surrounding people but not about the number of people or their identity. The mere fact that the receiver is surrounded by people or alone has been found to convey valuable cues about the user's availability [1]. Combining the company information with other cues, such as prior knowledge about the receiver, can only magnify the value provided by that context. For example, learning that a friend, whom you know to have his own office, has company can help one infer that he is busy, and that one should try to call at a later time.
- **Conversation:** Whether somebody is speaking or engaged in a conversation can be a very reliable indicator of one's availability. Fogarty has found that a speaking sensor can predict the interruptibility of the user with high accuracy (76%)

[9]. The speaking sensor was found to correlate with many other activities that directly affect our availability, such as being in a meeting or talking on the phone. Moreover, that it is technically simple to sense talking makes it a good candidate for useful contextual information.

3.2 Social Relations

Based on Olson et al., we chose 6 distinct categories of social relations between caller and receiver [22]. These include Significant other, Family member, Friend, Colleague, Boss and Unknown. Even though some of the categories could be further refined (e.g. Friends could be broken down into "close friends" and "best friend"), we opted to keep these six groups in order to keep the task manageable for participants. The Unknown category is not expected to include only strangers because, in the context of cell phone communication, unknown numbers do not necessarily imply that a stranger is calling.

3.3 Methodology

At the beginning of the study, participants were familiarized with the purpose of the study and answered a P&AB-Harris Interactive privacy classification survey so they could be grouped according to their level of privacy concern [23]. Demographic data was also collected. Participants were each given a Palm PDA that ran a general purpose ESM application partially developed at Intel research labs named iESP [4]. They were given a brief introduction on how to use the PDA with a walk-through scenario similar to the ones they were expected to encounter in the study.

During the study, each participant carried a PDA. Throughout the day, every participant received inquiries prompting her to choose what context she would like to disclose to a potential caller



Figure 1. ESM question asking about the kind of context information the participant would be willing to disclose to the caller (a friend in this case).

(Figure 1). Participants were asked to assume the role of the receivers of a cell phone call. The caller assumed one randomly chosen role out of the 6 social relations. In addition to inquiring about participants' willingness to disclose different types of context information, every questionnaire included a list of questions about the current location, activity, number of surrounding people, and the social relationship to the surrounding people, as well as the participant's availability under those circumstances to receive a phone call from that particular caller. These questions allowed us to examine the effect of the environmental conditions on participants' pattern of context disclosure and availability preferences.

The questionnaires were triggered randomly throughout the day in order to maximize the chances of capturing a variety of different situations and also to minimize their predictability and thus decrease any potential bias in the answers. The questionnaires were triggered during the day from 10 am to 10 pm, unless the participant requested different times at the beginning of the study. The questionnaires were also equally distributed as to the roles of the callers. The study lasted for 10 days during which participants were prompted to answer the questionnaire 13 times a day. The 10-day period was chosen so as to elicit a representative sample from participants without overwhelming them with too many questionnaires per day. In addition, by conducting the in-situ study over 10 days, we were likely to capture participants' responses throughout a wide range of usual and unusual daily activities. Each questionnaire took 1-2 minutes to complete.

At the end of the study, semi-structured interviews were conducted with participants in order to obtain their overall feedback about the usage, acceptability and concerns of context-aware telephony. The interviews lasted approximately 20 minutes each.

3.4 Participants

We recruited 20 participants equally divided between males and females. Participants were ages 18-51 (average 24), were mostly students and all had either a full-time or part-time job. All participants had owned cell phones for more than a year (4.4 years on average) and regularly used their cell phones (daily average: made 5 calls and received 4).

Using the P&AB-Harris Interactive privacy survey, we found that twelve participants were privacy fundamentalists, six were pragmatists, and two were unconcerned. The P&AB-Harris study reported similar rates [23].

3.5 Design Tradeoffs

Since the study involved hypothetical context disclosure, and no real context was actually disclosed, this may have caused the participants to be less concerned about their privacy and thus disclose more context information than they would disclose in a real application. However, employing the in-situ technique to capture people's responses in real life is expected to minimize such bias.

The goal of the study is not to provide a comprehensive account of the details of privacy preferences for context-aware telephony, but rather to provide designers with high-level guidelines about user privacy preferences and sharing patterns.

4. FINDINGS

The overall participation rate was around 80%, which means participants answered 80% of the total questionnaires and missed 20%. After removing the partially answered questionnaires, a total of 2422 questionnaires were considered for the analysis.

In this section, we report the main findings of the study. We begin by examining overall context disclosure rates, the correlations between the examined contexts that emerge, and how we classify the contexts according to their disclosure pattern. We then analyze the effect of the social relationship between the caller and receiver on context disclosure rates and caller availability, identifying distinct patterns. We also look at how the current physical and social contexts (as reported by the participant) affect both context disclosure rates and availability. We present a discussion of how we can classify the participants in terms of their privacy concerns as gleaned from their in-situ answers, and how that compares to the P&AB-Harris survey. We end with a discussion of the effect of gender.

4.1 Context Classification

For every questionnaire, participants selected any combination of the four different types of context information they were willing to disclose, including "none". In most cases, participants opted to either reveal all context information or none at all. Participants revealed all the information (location, activity, company, and conversation) 41% (992 responses) of the time whereas they chose to reveal nothing 19.1% (463 responses) of the time. The third most frequent selection was Company and Conversation (18.7%, 452 responses), followed by Company in the fourth place (8.1%, 195 responses) and then Conversation in the fifth (3.6%, 88 responses). These data show that users tend to share as much information as possible, removing contexts from their disclosure lists selectively based on their specific situation.

4.1.1 Context Disclosure Rates

Table 1 shows the frequency with which participants chose to release each type of context information. Company was the most frequently released context type, closely followed by Conversation, then Location, and then Activity. Overall, this high rate of context disclosure indicates that solutions which aim to enhance cell phone interruptions based on the receiver's context are both feasible and practical.

Table 1. The disclosure frequency and percentage for each type of context information for all participants out of a total of 2422 questionnaires

Contextual info	Frequency	Percentage
Company	1799	74.3
Conversation	1681	69.4
Location	1148	47.4
Activity	1123	46.4

The disclosure rates for Location and Activity are very close to each other across all participants. However, participants appear to share Company slightly higher than Conversation. This may be because Conversation indirectly suggests the existence of Company either physically, as in the case of a face-to-face

conversation, or virtually, as in the case of a phone conversation. The opposite is not necessarily true, however. Location and Activity information were disclosed at virtually identical rates with Activity disclosed slightly less than Location. Even though the Activity and Location can convey the same information in many cases, Activity may convey more information than Location in some cases.

4.1.2 Grouping of Context Information

After examining the four different context disclosure rates for each participant, we found a strong correlation between disclosing Location and disclosing Activity of 0.92 (significant at 0.001 level). Another statistically significant correlation is found between Company and Conversation (the correlation value is 0.78, significant at 0.001 level). The high correlation between Company and Conversation on one side and Location and Activity on the other, together with the similar values for their disclosure rates across the participants, means that people tend to mentally categorize them into two different groups. The different context types within the same group are almost equivalent in terms of their pattern of disclosure and the rates of disclosure, allowing designers to choose the best context for their context-aware application from the desired group.

We also found a strong and statistically significant negative correlation between disclosing no context information and disclosing company and conversation contexts (correlation with company context is -0.91 and correlation with conversation context is -0.81, both significant at 0.001 level). There was no such correlation between choosing not to disclose any context information and disclosing Location and Activity. Thus, for certain situations, participants appeared to make a choice between disclosing Company/Conversation information and disclosing no information at all. This indirectly implies that the Company and Conversation group is perceived to have a higher privacy comfort level than the Location and Activity group.

4.2 Social Relationship of Caller and Receiver

The number of questionnaires per social relation or caller role is shown in Table 2. Ideally, the number of questionnaires would be equal across the types of social relations in order for any comparison across social relations to be precise and unbiased. However, this was not possible in our study since the social relation associated with each questionnaire was randomly selected and 20% of the questionnaires were not answered.

Table 2. The # and % of questionnaires per social relation

	SO	Friend	Boss	Colleague	Family	Unknown
#	374	427	380	395	416	430
%	15.4	17.6	15.7	16.3	17.2	17.8

4.2.1 Context Disclosure and Social Relations

Figure 2 shows the average disclosure rates across all participants for the four different types of contextual information for different social relations. The rate of context disclosure decreases as the strength of social relationship between the caller and the receiver becomes weaker. For example, participants disclose their locations around 76% of the time to their significant others, 60%

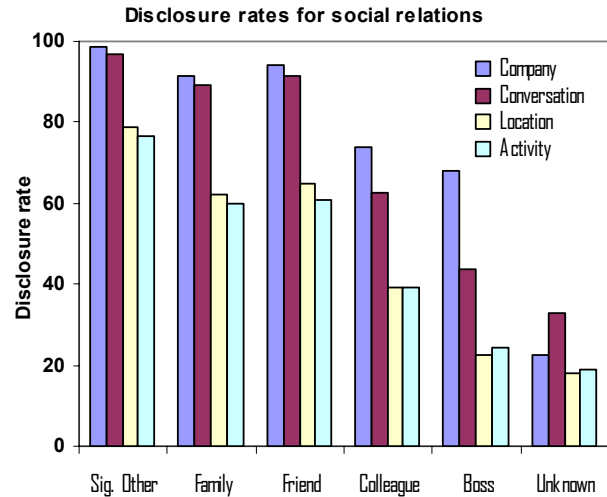


Figure 2. Average disclosure rate for different types of context information across different social relations

of the time to their family members, 61% of the time to their friends, 39% for their colleagues, 24% for their boss and 19% for anonymous or unknown callers. Interestingly, the rate of context disclosure is slightly higher for friends than family members. Even though the difference is not significant, the trend is consistent across all four context types, which suggests that people are in general more open to sharing contextual information with their friends (whom they choose) than with their family members.

The rate of sharing decreases drastically when moving from Location/Activity group to Company/Conversation group. This pattern is consistent across all social relationships. Participants shared Company more with Boss and Colleague than Conversation, (68%, 74%) and (44%, 63%), respectively, pointing to a higher level of privacy concern for Conversation than for Company. This may be because both Boss and Colleague already know the company information when participants are at work, but this is not the case for the conversation information. However, participants disclosed Conversation more than Company (33% and 22.6% respectively) when the caller was Unknown. The change in the rank order of disclosure rates for Company and Conversation moving from Boss and Colleague on one side to Unknown on the other side indicates that the ranking of different types of context information according to their rate of disclosure is dynamic and can change from one role to another.

4.2.2 Grouping of Social Relations

In order to study the sharing pattern of different types of contextual information amongst different social relations, we employed a hierarchical cluster analysis technique. The dendrograms shown in Figure 3 depict the different clusters for different social relationships. Each dendrogram shows the clusters for one type of contextual information using the disclosure rate by each one of the 20 participants for that particular type of context. We used a between-group linkage as a clustering method and Squared Euclidean Distance as a clustering measure. From the dendrograms we can see that there are 3 different main clusters.

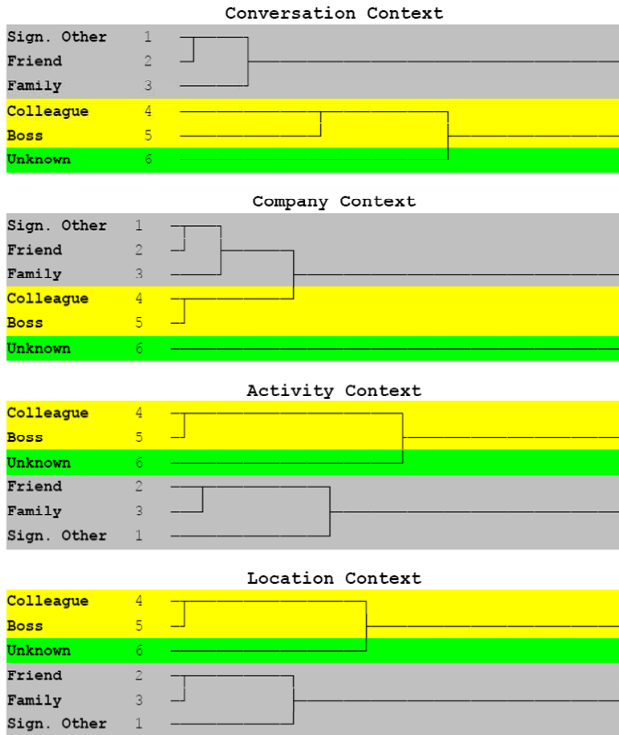


Figure 3. Dendrograms showing the clustering pattern of caller roles for each type of context information.

The first cluster contains the Significant other, Friend and Family member relationships while the second cluster contains Colleague and Boss, and the third cluster contains the Unknown category. The same clustering pattern was obtained when using the disclosure rates for the four different types of contextual information. However, the distances between clusters when they are joined vary from one type of context to another.

The three different clusters correspond to three levels of sharing: high, medium, and low. The high sharing group corresponds to the first cluster that contains social relations with whom participants are less concerned about sharing context information. On the opposite extreme, the low sharing group corresponds to the third cluster which only contains the Unknown social category. The medium sharing group corresponds to the second cluster and so contains the Boss and Colleague relationships in which participants share information only on an as-needed basis.

Having the same clustering pattern of social relations across different types of context suggests a pattern of sharing contextual information among different social relations. This common pattern of sharing can help designers of context-aware telephony applications in particular and privacy-aware systems in general to build simpler systems either by using an access control approach or a learning and recommending approach. These results are in agreement with the work of Olson, who obtained similar results investigating sharing patterns of 40 different types of personal data across different social relations [22].

4.2.3 Availability and Social Relations

During the in-situ study, participants were asked at the end of each questionnaire whether it was an appropriate time for the caller to call. Our results show that the availability rate (the percentage of the appropriate calls) does not only depend on the context of the receiver but also on the relationship between the caller and the receiver. In fact, the social relation was found to have a main effect ($F(5,114)=10.7, p< 0.001$). The availability rate for Significant Others is the highest (75%), followed by Friends (68%), Family members (63%), Boss (50%), Colleagues (47%), and Unknown (39%). This shows that designers of context-aware telephony applications that aim to minimize unwanted interruptions should take into consideration the social relation between the parties in addition to the receiver's context.

The availability rates across different social relations for any given participant seem to be independent of each other, as demonstrated by the weak correlations between availability and social relation. An exception to this is the strong correlation between family and significant others (0.77, significant at 0.025 level). This overall lack of correlation means that a high level of availability for a certain social relation does not necessarily mean a high level of availability for other social relations, and likewise with low levels.

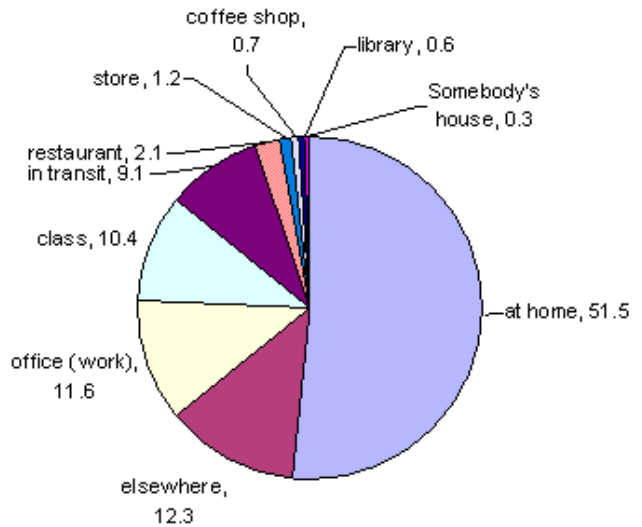


Figure 4. Distribution of locations as reported by the participants during the in-situ study.

The average availability over all different social relations and for all participants is around 57%. This means that a little less than half of the calls are not received at appropriate times. This high rate of unwanted incoming calls stresses the importance and the need for solutions to minimize cell phone interruptions. It is worth mentioning that the rate for inappropriately received calls may be lower in real life than the one obtained from our study due to the familiarity of friends, family members, and significant others with the work pattern of the receiver.

4.3 Effects of Physical/Social Context of Receiver

4.3.1 Location and Sharing Rates

Figure 4 shows the distribution of location information, or where participants spent most of their time, as reported by the participants during the in-situ study. When participants optionally chose to specify their exact meaning of “elsewhere”, the most popular entries were ‘at gym’, ‘office hour’, ‘park’ and ‘concert’. To explore the effect of location on disclosure rates, we examine the two most frequent locations: home and work.

Figure 5 compares the sharing rate for location information when participants reported being ‘at work’ and ‘at home’. As expected, more privacy is desired when at home than at work. One can see a steady decrease in sharing rates going from Significant Other to Unknown. One interesting observation is that participants disclosed location information at a considerably higher rate for Significant Other when at work than when at home. However, this is not the case for Friend and Family member. Colleague and Boss social relations appeared to be affected the most where participants disclosed a significantly higher rate of location information when at work than when at home. Similar patterns were observed for the other types of context information.

4.3.2 Availability and Surrounding People

Participants were asked to report the number of surrounding people during each questionnaire of the in-situ study. We examined whether the number of surrounded people they reported have any effect on their availability for the same situation to receive a phone call. We found that participants were alone for 38% of the times and were surrounded by more than 4 people

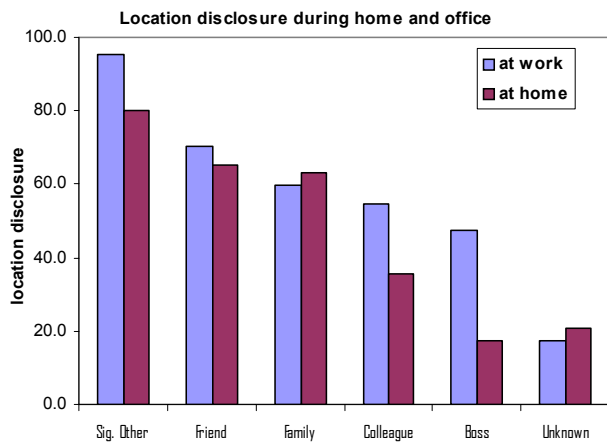


Figure 5. Comparison of sharing rate for location information across different social relations during home and work times

23% of the time. Figure 6 shows the average availability rate for different numbers of surrounding people. The availability rate decreases linearly as the number of surrounding people increases. This suggests that the number of surrounding people is a good indicator of the user’s availability. Number of surrounding people

provides information about user availability that could not be achieved by company information (whether a participant has a company or not) by itself. Company information only provides a binary value of whether the user has company or not regardless the number of people. This does not mean, however, that the number of surrounding people is the only factor at play when determining one’s availability. The pattern merely reflects average behavior. In many situations, as in waiting at a bus stop, one remains available despite a large number of surrounding people. Laurier has concluded that availability is a negotiable process between the caller and the receiver [14]. Availability prediction is thus a complex process dealing many interrelated factors.

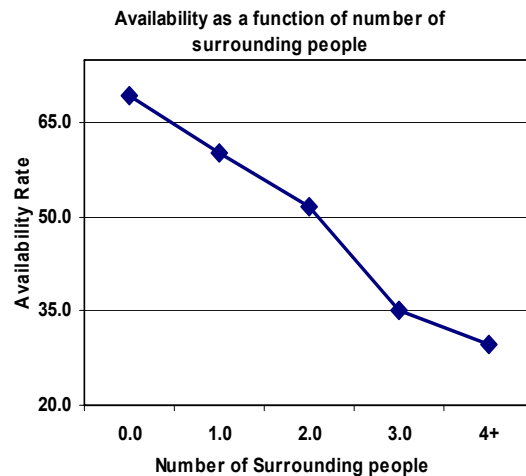


Figure 6. The average availability rate for different number of surrounding people

4.4 Participants’ Privacy Classifications

Previous research in consumer privacy [22, 23, 29] has shown that one fourth of consumers are “privacy unconcerned”, having little to no concern about their privacy, one fourth are highly concerned “privacy fundamentalists”, and around half fit in between those extremes as “privacy pragmatists”. To test whether our results reflect that breakdown and to examine the differences in participants’ privacy classifications, we categorized the participants according to their willingness to share context information with others. We chose the average disclosure rate for location information across all social relations to represent the privacy index for each user because Location and Activity were responsible for most of the variation in sharing rates among participants. The privacy index ranged from 0 to 100.

Figure 7 shows a histogram of participants’ average rate of sharing location information. From the figure we can see that 4 participants can be considered “privacy fundamentalists”, having sharing rates between 0-20%, 12 can be considered “privacy pragmatists” with sharing rates between 20% and 80%, and another 4 participants can be considered “unconcerned” due to their sharing rate between 80% and 100%. The overall average location sharing was 6.3% for privacy fundamentalists, 94% for privacy unconcerned, and 54.5% for privacy pragmatists. Clearly, “privacy pragmatics” make up the largest group, a fact which

agrees to a large extent with the results obtained in consumer research [23, 29]. However, participants were located in different categories from those in the P&AB. The same observation was cited by Consolvo et al.: the P&AB privacy index is not a suitable measure for privacy concerns in context-aware computing [3].

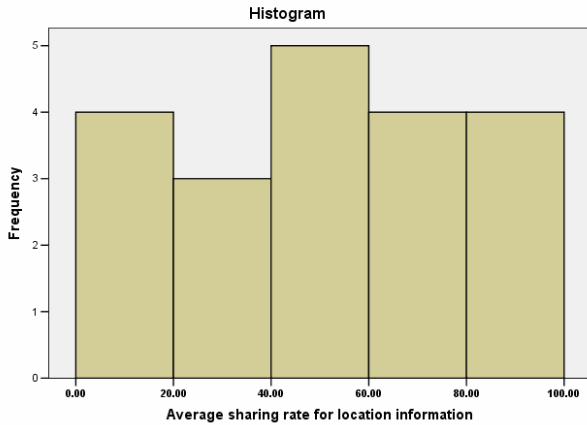


Figure 7. Histogram of average location disclosure rate across all social relations. There are 4 “privacy fundamentalists” (at left) and 4 “privacy unconcerned” (at the right).

4.5 Gender Effect

Gender was found to have a significant effect on the sharing rate across the four context types. Figure 8 shows the difference in location sharing between the male and female groups across different social relations. One can see that males shared significantly more location information than females for each one of the social relations. We performed an independent sample t-test (single-tailed) for every female-male pair’s rate of disclosure for each social relation. We obtained a statistically significant difference (at 0.05 level) for Friend, Family, Colleague, Boss, and Unknown (values for $t(18)$ are 2.27, 1.80, 2.75, 2.07, 1.90 respectively), and near statistical significance for Significant Other ($t(18)=1.71, p=0.052$).

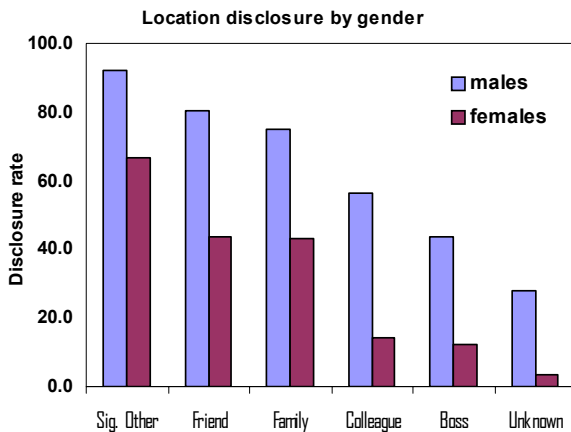


Figure 8. Comparison between males and females in location sharing across different social relations

The same behavior was found for the disclosure of Activity. However, even though males shared Company and Conversation at a higher rate, the differences in averages were not significant. Looking back at the ‘privacy fundamentalists’ group, we found that 3 out of the 4 members were females and only one female was found to be a member of the ‘privacy unconcerned’ group. Such results imply that females, in general, are more sensitive or cautious in sharing their contextual information than males.

4.6 Qualitative Results

During the end-of-study interview, 70% of participants reported they were willing to use a service that publishes their context information comparable to the one used in our study if their cell phones were equipped with it and if they were provided with a tool to manage their privacy preferences. Twenty percent of the participants answered that they would possibly use the service, while the remaining 10% would avoid the service due to privacy concerns. Participants were also asked to rate the usefulness of the service on a scale of 1 to 5, with 5 being the most useful. Fifty-five percent of the participants rated it 4, 30% rated it 3, 10% rated it 2, and 5% rated it 1.

We also found that participants frequently changed their cell phones’ profiles, an average of 4 to 5 times per day. One participant reported keeping his cell phone in vibrate mode all the time, despite the risk of missing some calls, rather than changing the profile back and forth. Another participant reported that she stopped carrying her cell phone around because of the inconvenience of untimely interruptions.

5. IMPLICATIONS FOR DESIGNERS

Our findings suggest that people are willing to share personal information in exchange for useful services. Enhancing the agreement between the caller and the receiver in the context of cell phone communication such that inappropriate calls are minimized was found to offer a good incentive for participants to share their context information.

The results showed significant differences in the disclosure rates for different types of context information. For the four contexts we examined, two distinct groups emerged which we classified as low sharing rate and high sharing rate. Designers of context-aware telephony applications can use this classification to decide to use any contextual type within the same group depending on design requirements other than the rate of sharing and privacy comfort level. For example, a designer can choose to enable sharing of conversation or company information depending on the presence of sufficient audio or video infrastructure, respectively.

Focusing too much on location information, as is common practice, prevents systems from making use of other types of context information that could convey valuable details about the situation while simultaneously being sensitive to users’ privacy concerns. For example, information about the number of surrounding people was found to be a good indicator of participants’ availability.

By classifying participants’ privacy sensitivity in terms of the frequency in which they disclosed Location, we found that users showed varying degrees of willingness to share information. Designers should take this into consideration when designing context-aware telephony and avoid rigid rules, allowing for

flexibility to accommodate individual differences. Indeed, it would behoove designers to provide users with options for sharing context from both low and high privacy comfort groups in order to accommodate the privacy fundamentalists and the privacy unconcerned.

Our findings demonstrate the important role of social relations on the rate of disclosure for different contexts. For the 6 social relations we examined, we found 3 clusters with distinct sharing patterns. By using both the clustering of context information and the clustering of social relations, designers can create efficient and user-friendly privacy management tools. The inclusion of an efficient privacy management tool that provides users with full control over their privacy preferences without the burden of complex configuration may be critical to the acceptance and adoption of such a service.

Finally, even though our study predicted similar ratios of the three different groups of privacy classifications to those found by P&AB, some participants were categorized differently. This suggests that the P&AB privacy index does not serve as a global measure of privacy concerns, and different privacy indexes are needed for different areas such as context-aware telephony.

6. DISCUSSION

People perceive different kinds of personal information with varying degrees of sensitivity or privacy comfort [10]. Patil and Lai attribute the difference in the rate of context disclosure to the level of privacy comfort associated with a particular type of context information [25]. The higher the disclosure rate for a certain type of context information, the less sensitive the person is and the higher the level of privacy comfort regarding this context. Looking back at our results, we can see that both Location and Activity information were disclosed less frequently than Company and Conversation. This suggests that Location and Activity triggered lower privacy comfort levels while Company and Conversation triggered higher levels of privacy comfort. This is not surprising given that in many cases learning the activity or the location of a close friend automatically conveys substantial information about the friend's company and conversation status. For example, finding out that a friend is at a movie theater indirectly suggests with high probability that she has company and is not engaged in a conversation. However, this ability to infer information does not work reciprocally, i.e., knowing Company or Conversation status does not generally lend information about Location and Activity.

The association of the different types of context information to varying levels of privacy comfort is not only evident from the disclosure rate but also from the negative correlation between disclosing Company and Conversation and disclosing nothing. This means that, in many cases, participants choose to disclose Company and Conversation over not disclosing anything, suggesting that people disclose less sensitive context information when they are concerned about their privacy.

7. CONCLUSION

We have presented a study investigating context disclosure and sharing patterns for context-aware telephony with the aim of decreasing interruptions and enhancing agreement between callers and receivers. However, the results of the study, such as

clustering of context information and social relations, are relevant to context-aware applications in general.

Our most prominent findings are:

- Context-aware telephony is not only feasible but also desirable, as is reflected by the high level of acceptance and the high rate of context disclosure.
- The clustering of context information into two groups of high sharing rate and low sharing rate, as well as the clustering of social relations, indicates a pattern of privacy preferences. These patterns can be used to guide designers of context-aware applications.
- Contextual information can provide valuable cues as to people's availability. Different pieces of contextual information offer varying levels of clarity or ambiguity about the user's availability. However, combining different sources of context information along with background and social information drastically decreases ambiguity in availability cues.
- When people are concerned about privacy, they will selectively remove contexts from their disclosure list instead of disclosing no context at all. This indicates that people want to share as much information as possible without compromising their privacy in return for useful services.

The low availability rate of our participants to receive cell phone calls (only 53% of the time) highlights the usefulness of such a service and may explain the high acceptability rate. Privacy preferences and disclosure patterns discussed in the paper should provide guidance for the designers of context-aware telephony, as it is essential that designers provide users with an easy and efficient way to control their information disclosure.

8. ACKNOWLEDGMENTS

This work was partially funded by an NSF Research grant.

9. REFERENCES

- [1] Avrahami, D., D. Gergle, S.E. Hudson, and S. Kiesler, Improving the match between callers and receivers: A study on the effect of contextual information on cell phone interruptions. To appear in *Behaviour & Information Technology*.
- [2] Bellotti, V. and A. Sellen. Design for privacy in ubiquitous computing environments. In *Proceedings of Third European Conference on Computer Supported Cooperative Work (ECSCW'93)*. 1993.
- [3] Consolvo, S., Smith I., Mathews T., and L. A. Location Disclosure To Social Relations: Why, When, & What People Want To Share. In *Proceedings of CHI 2005*. 2005. Portland, Oregon.
- [4] Consolvo, S. and M. Walker, Using the Experience Sampling Method to Evaluate Ubicomp Applications: The Human Experience. *IEEE Pervasive Computing Mobile and Ubiquitous Systems*, 2003. 2(2): p. 24-31.
- [5] Csikszentmihalyi, M., R. Larson, and S. Prescott, The ecology of adolescent activity and experience. *Journal of Youth and Adolescence*, 1977. 6: p. 281-94.

- [6] Cutrell, E., M. Czerwinski, and E. Horvitz. Notification, Disruption, and Memory: Effects of Messaging Interruptions on Memory and Performance. In Proceedings of Interact 2001. 2001.
- [7] Danezis, G., S. Lewis, and R. Anderson. How much is location privacy worth? In Proceedings of Workshop on Economics of Information Security (WEIS 05). 2005.
- [8] Eagle, N. and A. Pentland, Social Serendipity: Mobilizing Social Software. IEEE Pervasive Computing, 2005. 4(2).
- [9] Fogarty, J., S. Hudson, C. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. Lee, and J. Yang, Predicting Human Interruptibility with Sensors. ACM Transactions on CHI, 2004.
- [10] Goffman, E., The Presentation of Self in Everyday Life. 1959, Garden City, NY.
- [11] Jensen, C. and C. Potts, Privacy Practices of Internet Users: Self-report versus Observed Behavior. To appear in International Journal of Human Computer Studies special issue on HCI Research in Privacy and Security.
- [12] Kern, N., S. Antifakos, B. Schiele, and A. Schwaninger. A model for human interruptability: experimental evaluation and automatic estimation from wearable sensors. In Proceedings of IEEE International Symposium on Wearable Computers. 2004.
- [13] Khalil, A. and K. Connelly. Improving cell phone awareness by using calendar information. In Proceedings of INTERACT 05. 2005. Rome, Italy.
- [14] Laurier, E., Why people say where they are during mobile phone calls. Environment and Planning D: Society & Space, 2001: p. 485-504.
- [15] Lederer, S., J.I. Hong, A. Dey, and J.A. Landay, Personal Privacy through Understanding and Action: Five Pitfalls for Designers. Personal and Ubiquitous Computing, 2004. 8(6).
- [16] Lederer, S., J. Mankoff, and A.K. Dey. Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In Proceedings of Extended Abstracts of CHI 2003. 2003.
- [17] Marmasse, N., C. Schmandt, and D. Spectre. Watchme: Communication and Awareness between Members of a Closely-Knit Group. In Proceedings of UbiComp. 2004.
- [18] Milewshi, A.E. and T.M. Smith. Providing Presence Cues to Telephone Users. In Proceedings of CSCW 2000. 2000.
- [19] Nelson, L., S. Bly, and T. Sokoler. Quiet Calls: Talking Silently on Mobile Phones. In Proceedings of CHI 2001.
- [20] O'Connell, B. and D. Frohlich. Timespace in the workplace: Dealing with interruptions. In Proceedings of CHI '95. Extended Abstracts. 1995: ACM Press.
- [21] Oliver, N., E. Horvitz, and A. Garg. Layered Representations for Recognizing Office Activity. In Proceedings of International Conference on Multimodal Interaction (ICMI 2002). 2002.
- [22] Olson, J., J. Grudin, and E. Horvitz. A study of preferences for sharing and privacy. In Proceedings of CHI 05. 2005.
- [23] P&AB, Consumer Privacy Attitudes: A Major Shift since 2000 and Why. Privacy & American Business Newsletter, 2003. 10(6).
- [24] Palen, L. and P. Dourish. Unpacking "Privacy" for a Networked World. In Proceedings of ACM Conference on Human Factors in Computing Systems CHI 03. 2003.
- [25] Patil, S. and J. Lai. Who Gets to Know What When: Configuring Privacy Preferences in an Awareness Application. In Proceedings of CHI 2005. 2005.
- [26] Pedersen, E.R. Calls.calm: Enabling Caller and Callee to Collaborate. In Proceedings of CHI 2001. 2001.
- [27] REUTERS, Mobile phone users double since 2000, <http://www.computerworld.com/mobiletopics/mobile/story/0%2C10801%2C98142%2C00.html>
- [28] Schmidt, A., A. Takaluoma, and J. Mantyjarvi, Context-Aware Telephony Over WAP. Personal and Ubiquitous Computing, 2000. 4(4): p. 225-229.
- [29] Sheehan, K., Toward a Typology of Internet Users and Online Privacy Concerns. The Information Society, 2002. 18(1): p. 21-32.
- [30] Siewiorek, D., A. Smailagic, J. Furukawa, N. Moraveji, K. Reiger, and J. Shaffer. SenSay: A Context-Aware Mobile Phone. In Proceedings of IEEE International Symposium on Wearable Computers (ISWC). 2003.
- [31] Speier, C., J.S. Valacich, and I. Vessey. The effects of task interruption and information presentation on individual decision making. In Proceedings of Eighteenth International Conference on Information Systems. 1997.
- [32] Tang, C., N. Yankelovich, J. Begole, M. Van kleek, F. Li, and J. Bhalodia. ConNexus to awarenex: extending awareness to mobile users. In Proceedings of CHI 2001.