The Impact of Expressiveness on the Effectiveness of Privacy Mechanisms for Location Sharing

**Michael Benisch, Patrick Gage Kelley,
Norman Sadeh, Tuomas Sandholm,
Lorrie Faith Cranor, Paul Hankes Drielsma,
Janice Tsai**

December 2008
CMU-ISR-08-141

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This work has been supported by NSF grants CNS-0627513 and IIS-0427858. Additional support has been provided by Nokia, France Telecom, Nortel, the CMU/Microsoft Center for Computational Thinking, ARO research grant DAAD19-02-1-0389 to Carnegie Mellon University's CyLab, and the CMU/Portugal Information and Communication Technologies Institute. The authors would also like to thank Lucian Cesca, Jialiu Lin, Tony Poor, Eran Toch, and Kami Vaniea for their assistance with our study.

Keywords: Expressiveness, Usable privacy, Location sharing, Web services, Social networking, Mechanism design

Abstract

A recent trend on the Web is a demand for higher levels of expressiveness in the mechanisms that mediate interactions such as the allocation of resources, matching of peers, or elicitation of opinions. In this paper, we demonstrate the need for greater expressiveness in *privacy mechanisms*, which control the conditions under which private information is shared on the Web. We begin by adapting our recent theoretical framework for characterizing expressiveness to this domain. By leveraging prior results, we are able to prove that any increase in allowed expressiveness for privacy mechanisms leads to a *strict* improvement in their efficiency (i.e., the ability of individuals to share information without violating their privacy constraints). We validate these theoretical results with a week-long human subject experiment, where we tracked the locations of 30 subjects. Each day we collected their stated ground truth privacy preferences regarding sharing their locations with different groups of people. Our results confirm that i) most subjects had relatively complex privacy preferences, and ii) that privacy mechanisms with higher levels of expressiveness are significantly more efficient in this domain.

1 Introduction

Over the past few years we have seen an explosion in the number and different types of websites that allow individuals to exchange personal information and content that they have created. These sites include online social networks, photo and video-sharing sites, and location-sharing services on the Internet. While there is clearly a demand for users to share this information with each other; recently, we have started to see a change in attitude, with users demanding greater control over the conditions under which their information is shared. This change has led to expanded privacy controls on sites such as Facebook and Flickr.

In this paper, we conduct a user study where we track 30 participants over a one week period. Based on their location trails, we ask them to rate when, where, and to whom they would be comfortable sharing their locations. We then apply our recent theoretical framework [5] for studying expressiveness to the domain of privacy for Web-based information sharing. We focus on a class of mechanisms that we call *privacy mechanisms*, or mechanisms that allow individuals to control the circumstances under which certain pieces of private information are shared. Our notion of "expressiveness" refers to the level of detail or granularity that users are able to use to control the sharing of their personal information. By applying this theoretical framework to the use of mobile location-sharing technologies in a user study, we find that providing users with a greater amount of expressiveness in the creation of rules governing the sharing of their location can lead to the design more efficient privacy mechanisms – or mechanisms that allow individuals to share more of the information they want to share, without violating their privacy constraints.

More than 40 different location-sharing applications exist on the Web today, many of which emerged over the last year.[1] These applications allow users to share their location (frequently, their exact location on a map) and other types of information, but have extremely limited privacy controls. Typically, these mechanisms only allow users to specify a *black list*, or a listing of the individuals with whom they would never share their locations. Despite the number of location-sharing applications available, there does not seem to be a specific service that has captured a large proportion of market share, perhaps indicating that the existing levels of control are not adequate enough to allay users' privacy concerns.

Recent work has suggested that individuals have a difficult time expressing their privacy preferences in the sharing of their location information [7, 8, 20, 24, 30, 36]. One reason for this difficulty is that these systems may lack the expressiveness to capture users' true preferences. To determine the level of expressiveness needed in the context of a location-sharing application, we conducted a user study to measure how often and under which conditions users would share their information. The goal of our experiment was to better understand the complexity of real-world privacy preferences, and to determine the most appropriate forms of expressiveness for privacy mechanisms that control access to location information. We tracked 30 subjects for one week, and analyzed more than 3,800 hours of location information with corresponding subject-stated ground truth privacy preferences.

[1]This rapid increase Web-based location sharing services is largely due to the introduction of Yahoo!'s easy-to-use location sharing FireEagle API.

Among our findings are the following:

- Most subjects have complex privacy preferences regarding when, where, and with whom their locations can be shared.
- The privacy settings offered by today's location sharing applications (i.e., black lists) appear to be unsuitable to the wide array of privacy preferences revealed by our study. This finding may help explain the lack of broad adoption encountered by these applications so far.
- Mechanisms that allow subjects to hide locations based only on time of day, or based only on location, are roughly equivalent in terms of their performance. However, for individuals in the university community, location appears to be significantly more important than time.
- Expressions about time and location do not appear redundant. Allowing subjects to block certain individuals from seeing their locations based on time of day *and* location leads to significantly better performance than either time or location on its own.

While our results suggest that expressive privacy mechanisms are necessary to capture users true preferences, added expressiveness does not come without cost. It generally implies collecting more preference information from people or businesses, which in turn may incur additional cost or user burden. It can also lead to confusion (e.g., Herb Simon's concept of "bounded rationality" [34]), if not outright misery (e.g., Barry Schwartz's "tyranny of choice" [33]). What we provide is a methodology to inform the design of expressive privacy mechanisms by identifying the most relevant privacy dimensions for a particular user population, and to quantify the cost of limiting users to less expressive mechanisms.

2 Theoretical background

In prior work, Benisch, Sadeh and Sandholm [5] introduced the first domain-independent formal framework for studying expressiveness in mechanisms. This framework allows us to meaningfully characterize the expressiveness of different mechanisms, and demonstrates the strong ties between a mechanism's expressiveness and its efficiency. In this section, we describe how we can adapt this theory to study privacy mechanisms.

One key difference between the formal model of expressiveness in this paper, and that of earlier work is a move to a single agent setting. In this paper, we assume that the behaviors of agents other than the one making an expression are stochastic, rather than strategic (e.g., requests for one's private information are assumed to come from some probability distribution, rather than the behavior of other rational agents). Despite this difference, we will show that our theoretical framework for studying expressiveness can be naturally applied to this domain.

2.1 A general privacy mechanism model

The formal setting we study in this paper is that of a single request for a piece of private information, such as an individual's geographical location. We assume that a request can be described by a vector of m attributes, $\vec{a} = \{a_1, a_2, \ldots a_m\}$, such as the individual behind the request, or the time the request was placed. In general, each of these attributes can be discrete valued or real valued (however, in practice we discretize real-valued attributes, such as time). We assume that the attribute vector, $\vec{a}$, of a request is stochastically drawn from the set of all possible requests, $\vec{A}$, according to a joint probability distribution, which we denote as $P(\vec{a})$.

In our model, an agent interacting with the mechanism has a type, t, which is unknown to the mechanism. The agent's type is drawn according to some probability distribution, $P(t)$, from the set of all possible types, T, and represents the agent's attitude towards releasing any piece of private information under any circumstance (the set of all types can be finite or infinite). For example, an agent may have a type that is highly secretive about releasing its location during certain times of day, or its type may be more concerned about releasing certain locations.

The agent interacts with the mechanism by making an expression about its privacy preferences, which we denote as θ, from the space of all possible expressions, Θ. Based on the privacy preferences that the agent expresses and the attributes of a request, the mechanism computes the value of a binary outcome function, $f(\Theta, \vec{A}) \rightarrow \{0, 1\}$. The outcome function determines whether the request is granted (i.e., when $f(\theta, \vec{a}) = 1$) or denied (i.e., when $f(\theta, \vec{a}) = 0$). [2]

We assume that the agent has a utility function, u, which depends on the agent's type, the attributes of a request, and the outcome chosen by the mechanism. The utility function maps these inputs to a real-valued utility indicating how happy or unhappy the agent is with the outcome chosen by the mechanism, $u(T, \vec{A}, \{0, 1\}) \rightarrow \mathbb{R}$. We will also define an agent's strategy, $h(T) \rightarrow \Theta$, as a mapping from each possible type to an expression. A strategy dictates how the agent will interact with the mechanism depending on its type. Typically we assume that the agent will choose a strategy, h^*, that maximizes its expected utility.

$$h^*(t) = \arg\max_{\theta} \int_{\vec{a}} P(\vec{a}) u(t, \vec{a}, f(\theta, \vec{a}))$$

Using this model we can describe the expected efficiency of a particular privacy mechanism with the following equation (where expectation is taken over the possible types of the agent and the different possible request attributes, when attributes and types are considered to be discrete the integrals in the following equation would be summations instead):

$$E[\mathcal{E}(f)] = \int_t P(t) \int_{\vec{a}} P(\vec{a})\ u(t, \vec{a}, f(h^*(t), \vec{a})) \tag{1}$$

[2] In this paper we assume that the outcome function is binary: it either grants or denies a request. However, it is possible to generalize our notion of binary outcomes to include cases where a request can be granted to differing degrees, such as releasing an individual's city, rather than exact location.

2.2 Policy-based utility functions

In our empirical analysis we focus on one simple class of utility functions, which we call *policy-based utility functions*. An agent always has some underlying privacy preference function, $\pi(T, \vec{A}) \rightarrow \{0, 1\}$, which indicates the outcome that the agent prefers for any possible request. With a policy-based utility function we assume that the agent suffers a cost c whenever the mechanism inappropriately grants a request, the agent suffers a cost of c' whenever the mechanism denies a request that should have been granted, and the agent receives reward r whenever the mechanism correctly releases information. Typically we assume that the cost for mistakenly revealing a piece of private information is much greater than the reward for correctly sharing it, (i.e., $c >> r$). Table 1 illustrates this class of utility functions under each of the four possible scenarios: i) the mechanism correctly grants, ii) correctly denies, iii) inappropriately grants or iv) inappropriately denies.

	Mechanism denies ($f(\theta, \vec{a}) = 0$)	Mechanism allows ($f(\theta, \vec{a}) = 1$)
Agent denies ($\pi(t, \vec{a}) = 0$)	$u(t, \vec{a}, f(\theta, \vec{a})) = 0$	$u(t, \vec{a}, f(\theta, \vec{a})) = -c$
Agent allows ($\pi(t, \vec{a}) = 1$)	$u(t, \vec{a}, f(\theta, \vec{a})) = -c'$	$u(t, \vec{a}, f(\theta, \vec{a})) = r$

Table 1: An illustration of the policy-based utility function class under each of the four possible scenarios: i) the mechanism correctly grants, ii) correctly denies, iii) inappropriately grants or iv) inappropriately denies.

2.3 Expressiveness in privacy mechanisms

In our prior work on expressiveness, we introduced a measure called *impact dimension* as a measure of the expressiveness of mechanisms [5]. Impact dimension measures the extent to which an agent can impact the outcome that is chosen by a mechanism, by counting the number of different *impact vectors* that an agent can distinguish among. In a privacy mechanism, an impact vector describes the impact of a particular expression by an agent under all possible requests that could be placed for the agent's information.

Intuitively, more expressive privacy mechanisms allow an agent to distinguish among larger sets of impact vectors. The adaptation of the impact dimension measure for the privacy mechanism setting captures this intuition; it measures the number of different impact vectors that an agent can distinguish among.

By extension, the results in our earlier work imply that when designing a privacy mechanism, any increase in allowed expressiveness can be used to achieve strictly higher expected efficiency.[3] In addition, they imply that even a small increase in allowed expressiveness can be used to achieve an arbitrarily large increase in a mechanism's expected efficiency. These

[3]Proof of all theoretical claims can be found in the Appendix. The results described in this section have been adapted to this domain from our prior work [5]. The primary departure from our prior work is the move to a stochastic setting, rather than a strategic setting.

results taken together suggest that privacy mechanisms can be made significantly more efficient by designing them with greater levels of expressiveness. In the next section, we will describe an extensive human subject experiment designed to test these findings in practice.

3 An empirical study of location sharing privacy mechanisms

In the previous section, we demonstrated how greater levels of expressiveness can be used to design more efficient privacy mechanisms in theory. We now discuss a user study that we performed to validate this theory with real-world data. Our findings confirm that, under certain reasonable assumptions about the cost associated with revealing sensitive information, more expressive privacy mechanisms will indeed be significantly more efficient in the context of an actual location-sharing application.

3.1 Experiment overview

Our experiment was conducted over the course of two weeks in early October 2008. We supplied 30 participants with Nokia N95 cell phones[4] for one week at a time (15 subjects were run at once). The subjects were required to transfer their SIM cards to the phones we provided and use them as their primary phones for an entire week. This requirement ensured that the subjects kept their phones on their person and charged as much as possible. Each of the phones was equipped with our location-tracking program, which recorded the phone's location at all times using a combination of GPS and Wi-Fi-based positioning.

Each day, subjects were required to visit our web site and upload a file containing their location information from their phone. They were then asked to audit the location information by answering a set of questions about each location that they visited since their last login. For each location a subject visited, we asked whether or not he or she would have been comfortable sharing that location with different groups of individuals. These groups consisted of close friends, immediate relatives, people within the university community, and strangers. While no location-sharing to others actually occurred, we solicited the names of people from the friends and relatives groups so that the questions the users answered were more meaningful to the participant (i.e., are you comfortable with sharing your location with your mom?).

Subjects were paid a total of \$35, corresponding to \$5 per day, for their participation in the study. We also administered surveys before and after the study to measure the level of concern about their privacy that people had about sharing their location information, to collect relevant demographics, and to determine qualitative measures of the subjects' privacy attitudes.

[4]These phones were generously provided by Nokia.

3.2 Materials

The primary materials we used in our experiment included location-tracking software written for the Nokia N95 phones, a web application that allowed subjects to audit their location information each day, a pre-study survey to collect demographics and qualitative measures of privacy attitudes, and an exit survey. We will now describe each of these components in detail.

3.2.1 Location tracking software

Our location tracking software was written in C++ for Nokia's Symbian operating system. It runs continuously in the background, and starts automatically when the phone is turned on. During normal operation, the software is completely transparent – it does not require any input or interaction.

When designing our software, we faced three primary challenges: i) managing its energy consumption to ensure acceptable battery life during normal usage, ii) determining the phone's location when indoors or out of view of a GPS signal, and iii) communicating a significant amount of location information back to our server without relying on expensive data channels.

To address these challenges, our software is broken down into three different modules: a *positioning module* that tracks the phone's location using a combination of GPS and Wi-Fi-based positioning, an *output module* that writes a minimal amount of location information to a file, and a *management module* that turns the positioning module on and off to save energy.

Management module. Our initial tests revealed that leaving the GPS unit on at all times resulted in an unacceptable battery life of 5-7 hours on average. The management module depends on the N95's built in accelerometer to address the issue of energy consumption. It constantly monitors this low energy sensor, and only activates the positioning module when the accelerometer reports substantial motion. When substantial motion is sensed, the positioning module is activated for a period of at least five minutes, which is typically the amount of time needed by the GPS unit to determine its position. After this time, the positioning module is deactivated unless additional motion is sensed. Any time new motion is sensed while the positioning module is active the deactivation is delayed by one minute.

The phone's accelerometer sensor records acceleration in three dimensions at a rate of about 40 readings per second. In our software, the output of this sensor is smoothed by maintaining a moving average of the total acceleration sensed in all directions. The duration of the moving average (2 minutes) and the threshold for determining whether or not the phone has undergone substantial motion during that period (0.1 g's after accounting for gravity) were determined empirically. In practice we found that this technique improved the phone's battery life to 10-15 hours on average.

Positioning module. To estimate the position of the phone, our positioning module makes use of the Nokia N95's built in GPS unit, and Wi-Fi unit. When activated, the positioning module registers itself to receive updates from the GPS unit at a regular interval (15 seconds). When the GPS unit is able to determine the phone's position, the positioning module records its latitude and longitude readings.

In our initial tests we found that the GPS signal was unreliable when the phone was indoors, and even when the phone was outdoors on cloudy days. For that reason, whenever the positioning module is active it also records the MAC addresses and signal strengths of all nearby Wi-Fi access points at a regular interval (3 minutes). Our server is able to use this information to determine the physical address of the phone using Skyhook Wireless.[5]

The subscription interval for the GPS unit and the scan interval for the Wi-Fi unit were chosen based on energy considerations. The GPS unit consumes a substantial amount of energy when initially acquiring a lock on the phone's position. However, subsequent readings are relatively inexpensive, allowing us to subscribe at a fine granularity for a small marginal cost. Wi-Fi scans are performed less frequently because each scan consumes a substantial amount of energy (roughly equivalent to running the GPS for 3 minutes).

Output module. While the position module is active, the output module appends all location information (i.e., latitude and longitude readings from the GPS unit, or MAC addresses and signal strengths from Wi-Fi scans) to a file on the phone's built in memory. It also appends a heart beat to the file at a regular interval (3 minutes) to record exactly when the software is running. To transfer the file to our server, subjects connected their phone to a PC via USB cable and uploaded the file directly from the phone to our web application.

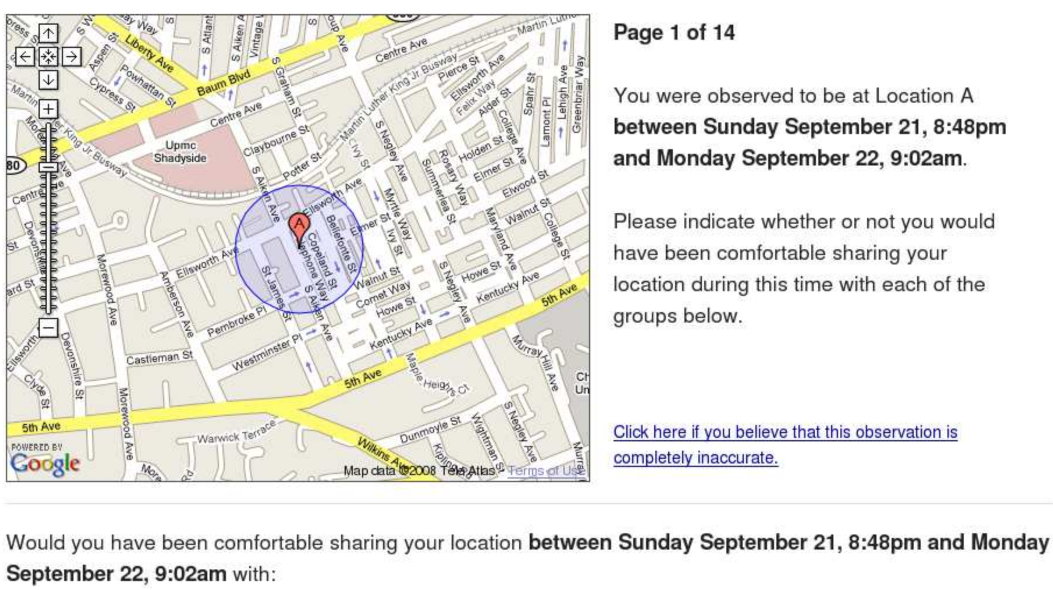

Figure 1: A screen shot of the web application displaying an example location between 8:48pm and 9:02am.

[5]Details about the Skyhook API are available at http://skyhookwireless.com/.

3.2.2 Web application

Each day subjects were required to visit our web site to upload their current location file and audit the location they visited that day.

Location file processing. When a subject uploads his or her location file to our web application, it iterates through each of the GPS and Wi-Fi readings that have been recorded since the last time the file was uploaded. Each of these readings is either associated with a location observation or a path observation between two locations. An observation was considered to be a new location whenever a subject moved more than 200 meters and remained stationary for at least 15 minutes.

Audit administration. After a subject's location file is processed, our web application takes the subject through a series of pages that trace his or her location since the last time the file was uploaded, in chronological order. Each page displays a location on a map inside a 200 meter ring indicating the subject's estimated location during a particular time period.[6] The times when the subject arrived and departed from the location are indicated next to the map. Each page also includes a link which allowed subjects to indicate that an observation was completely inaccurate (inaccurate observations accounted for less than 1% of the time, and were removed from our data set). A screen shot of the user interface for this part of the web application is shown in Figure 1.

Underneath the map on each page, our web application presents a collection of four questions, each of which corresponds to a different group of individuals. Each question asks whether or not the subject would have been comfortable sharing his or her location with the individuals in one of the groups. The groups we asked about in our study were: i) close friends, ii) immediate family[7], iii) anyone associated with our university, and iv) the general population. Subjects were given the option of indicating that they would have shared their location during the entire time span indicated on the page, none of the time span, or part of the time span (when part of the time is chosen, a drop down menu appears allowing the subjects to specify which part of the time they would have allowed). In addition, questions about the friends and family groups included a fourth option allowing subjects to indicate that they would have shared their location with some of the individuals in the group, but not all of them. This option was chosen less than 1% of the time and is treated as denying the entire group in our analysis. Figure 2 shows an example screen shot of a question for the close friends group.

[6] Path observations between locations were also depicted on some pages. However, we do not address those observations in this paper since they accounted for less than 1% of the observed time.

[7] For close friends and immediate family, subjects were required to provide three or four names to give them context while auditing.

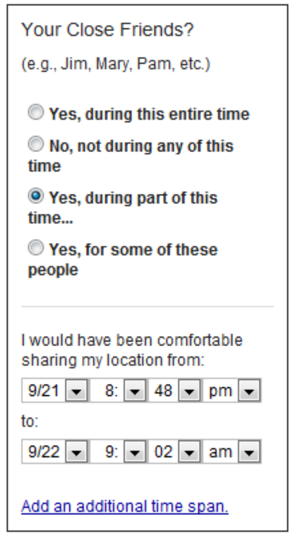

Figure 2: A screen shot of an audit question asking whether or not a subject would have been comfortable sharing his or her location between 8:48pm and 9:08am. Drop down menus are only displayed because "Yes, during part of this time..." is selected.

3.3 Mechanisms we compared

In this study, we focused on evaluating the expected efficiency of the following four different privacy mechanisms. We will illustrate the differences between these mechanisms by considering a hypothetical user named "Alice," who wishes to share her location only with her friends when she is at home between the hours of 9am and 5pm. The default setting for each mechanism or rule is to deny the sharing of location information.

- **Black list (BL).** The black list mechanism is the least expressive mechanism we consider; it only allows users to express whether or not they would be comfortable sharing their locations with each group at all times.

  Alice will need to define *who* (individually or by group) is allowed to see her at all times and at all locations. Similarly, she may also create a rule that everyone is allowed to see her at all times with a list of exceptions.

- **Location-based (LOC).** The location-based mechanism allows users to express specific locations at which they would be comfortable sharing their locations with each group. This mechanism has a higher impact dimension, and is thus more expressive, than the BL mechanism. The LOC mechanism allows the same expressions as the BL mechanism (black listing a group can be simulated in the LOC mechanism by not sharing any locations with that group), as well as some additional expressions about specific locations.

  Alice will need to create a rule allowing friends to view of her location when she is at home. Friends will be able to see when she is home regardless of the time of day.

- **Time-based (TIME).** The time-based mechanism allow users to express time intervals (discretized into 15 minute blocks) during which they would be comfortable sharing their locations with each group (it does not consider the day of the week). Similar to the LOC mechanism, this mechanism is more expressive than the BL mechanism because it allows a larger set of possible expressions. For some distributions over possible requests, the TIME mechanism is more expressive than the LOC mechanism, but for other distributions the opposite is true. In other words, neither the LOC mechanism nor the TIME mechanism is more expressive for all possible request distributions.

  Under this mechanism, Alice will need to create a rule sharing her location with her friends between 9am and 5pm, regardless of where she was.

- **Location & time-based (LOC/TIME).** The location and time-based mechanism combines the expressions of the LOC and TIME mechanisms. It allows users to express time intervals during which they would be comfortable sharing specific locations with each group. This is the most expressive mechanism we explore in this paper, however it is not fully expressive because it does not allow for different expressions based on the day of the week.

  Alice would be able to express her true privacy preferences under this mechanism.

4 Results and findings

Before we present our analysis comparing the efficiency of different privacy mechanisms, we will present some results that describe the data that we collected and some relevant survey findings.

4.1 Survey results

Our 30 subjects were all students at our university. The sample was composed of 74% males and 26% females, with an average age of about 21 years old. Undergraduates made up 44% and graduate students made up 56% of the sample.

In the pre-study survey, participants were asked to rate on a 7-point Likert scale how concerned they would be for their privacy when using a location-sharing service (ranging from not concerned to extremely concerned). We found that, in general, people were concerned about their privacy ($M = 4.66$,), but also felt that it would be useful for other people to find them ($M = 4.69$, $\sigma = 1.7$), based on a rating on a 7-point Likert scale ranging from not useful at all to extremely useful.

We also surveyed participants about how comfortable they would be if their close friends, immediate family, members of the university community, or strangers could view their locations at anytime, times they had specified, or at locations they had specified. Based on ratings on a 7-point Likert scale (ranging from "not comfortable at all" to "fully comfortable"), we found that, in general, participants were more comfortable with their close friends and family locating them than people within their university community or strangers. Within each group, we also found that respondents had equal levels of comfort within each relationship type when offered time-based restrictions or location-based restrictions (the differences were not statistically significant when measured in a paired t-test).

In general, subjects reported that location and time-based rules would increase their levels of comfort by a factor of about 1.25. For example, our users indicated that they would not be comfortable if strangers could check their locations at anytime (M = 1.93); but at times or locations they had specified, their comfort levels would slightly increase (M = 2.41). These results indicate that our participants feel that location or time based rules restricting access to their location information would reduce their privacy concerns.

After using the system, we asked a subset of our participants how bad they thought it would have been on a 7-point Likert scale from "not bad at all" to "very, very bad" if the system had shared their information at times when they did not want it to be shared. Our subjects reported significant levels of dis-utility at the prospect of their locations being inappropriately shared with the university community (M = 4.29), and strangers groups (M = 5.43). In contrast, our subjects reported relatively little dis-utility at the prospect of their locations being inappropriately withheld.

The corroborates our assumptions in the development of the utility function where the cost function is much larger in cases of accidental disclosure than to accidental withholding.

We also asked our subjects if they would have answered the questions differently if we had actually been sharing their locations on the web, and almost all of the subjects (93.1%) responded that they would not have answered differently.

4.1.1 Location Trails

On average, our subjects were accurately observed for just over 75% of the time during our experiment. The graph in Figure 3 shows that our observations were distributed relatively evenly throughout most of the day.

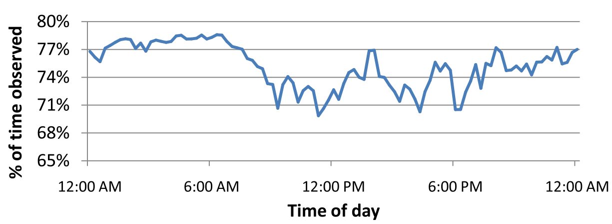

Figure 3: A graph showing the percentage of the time that we observed subjects on average during each 15 minute interval during a day.

We also found that most of our subjects visited 8 or fewer distinct locations throughout the week. A subject was considered to have visited a distinct location only if it was at least 200 meters from all other locations that the subject visited. Figure 4 shows the distribution over the number of distinct locations visited by our subjects.

We found that, on average, subjects spent significantly more time at one location than any others (most likely their homes). We also found that the time spent at a location appeared to

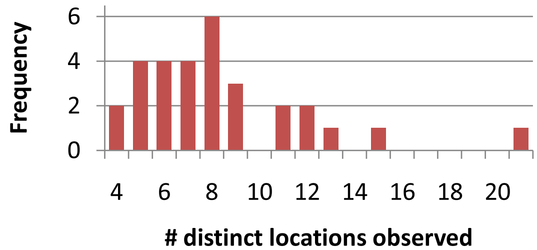

Figure 4: A histogram showing how many distinct locations subjects visited during our experiment (a location was considered distinct if it was at least 200 meters from all other locations the subject visited).

drop off exponentially for the second, third, fourth and fifth most visited locations (Figure 5). This result is similar to that of mobility patterns observed by Gonzalez *et al.* who found that human trajectories are very patterned with people visiting a small number of highly frequented places [12].

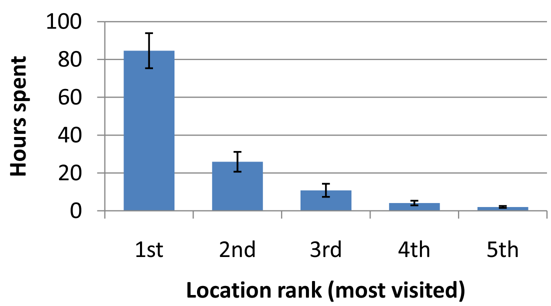

Figure 5: A plot showing the average amount of time that a subject spent at his or her five most visited locations.

Finally, we found that on average subjects would have been comfortable sharing their locations about 89% of the time with friends, 86% of the time with family, 46% of the time with other individuals at our university, and 26% of the time with the general population.

4.2 Rule usage

Based on the audits provided by our participants throughout the study, we are able to compute the number of rules they would have defined under each of the different privacy mechanisms we compared. For these calculations we assumed that the cost of mistakenly revealing one's location was five times greater than the reward for correctly revealing it (i.e., $c = 5$), and that users would then choose the rules that provided them with the optimal amount of utility.

Under the BL mechanism, each user would have only one rule composed of a list of people that can see them at all times. For the LOC mechanism, the number of rules necessary for each group is equal to the number of locations a user would have shared with that group (e.g., Alice shares her location with friends when she is at school or at home, so she has two rules for the friend's group). In the TIME mechanism the number of rules needed for each group is the number of contiguous blocks of time that a user would have opened for that group. To determine the number rules per group for the LOC/TIME mechanism, a similar evaluation is performed combining both time and location (e.g., Alice shares her location with relatives from 9am to 10am and 5pm to 7pm when she is at home, and from 11am-3pm when she is at school, so she has 3 rules).

We found that as the amount of expressiveness in a mechanism increased, so did the average number of rules that would have been used by our participants. Table 2 shows the average number of rules that would have been used by our participants in each mechanism for each group.

	Friends	Family	University community	Anyone	Total
Black list	N/A	N/A	N/A	N/A	1
Time	1.97	2.03	1.50	0.70	6.20
Location	6.90	6.23	3.30	1.37	17.80
Time/Location	7.97	7.97	5.23	2.73	23.90

Table 2: The average number of rules that would have been used by our participants for each mechanism and each group, assuming they acted optimally according to a utility function where the cost of mistakenly revealing their location was five times the reward for correctly revealing it (i.e., $c = 5$).

On average, we found that our users would have made a total of 6.20 rules in the TIME mechanism, 17.80 rules in the LOC mechanism, and 23.90 rules in the LOC/TIME mechanism. The number of rules for each group is statistically significant between groups (i.e., for friends, there are a greater number of location and time-based rules as compared to the other types of rules). Generally, people had the same number of rules for friends as they did for family (based on a series of t-tests, where p ¿ 0.2). For time-based rules, participants did NOT have a significantly different number of rules for people in the university community ($M = 1.5$) as compared to the number of rules for friends ($M = 1.97$), $t(29) = 1.46$, $p = 0.16$).

The number of rules that we generated for each user supports our hypothesis that people's privacy preferences are very nuanced.

4.3 The Efficacy of Expressiveness

We will now discuss our results regarding the complexity of our subjects' reported privacy preferences. In comparing the performance of different privacy mechanisms, we assume that each subject provided a ground truth privacy preferences when auditing his or her location information. We also assume that each subject is equally likely to use the mechanism, and that requests are equally likely to be made at all times.

We report the expected efficiency of each mechanism, assuming that subjects have policy-based utility functions (described in Section 2). The utility functions we study provide a reward of $r = 1$ unit per hour whenever a location is correctly shared (i.e., given to a group during a time that was marked as allowed). We assume that the subjects would receive 0 utility whenever their locations are blocked (i.e., $c' = 0$), rather than penalizing them for any missed opportunities. However, subjects pay a cost c whenever their locations are inappropriately shared (i.e., shared with a group during a time that was marked as not allowed). We report results with several different utility functions by varying the value of c.

For each utility function, we exhaustively search for the expression that a subject would have optimally specified.[8] Thus, the expected efficiency values that we report can be taken as upper bounds on the actual expected efficiency of these mechanisms, since subjects may not behave optimally in practice.

More expressive mechanisms have greater expected efficiency. The first set of results, presented in Figure 6, explores the performance of different mechanisms for each of the four different groups about which we asked our subjects. For this set of results, we fixed $c = 5$ as the cost associated with inappropriately revealing a subject's location (recall that this is 5 times the reward for correctly revealing a subject's location). Under our assumptions, these results confirm the hypothesis that subjects' privacy preferences are complex enough to warrant mechanisms with higher levels of expressiveness. For three of the four groups we asked about, each increase in expressiveness lead to significantly[9] higher expected efficiency.

For the friends, family, and university community groups the LOC/TIME mechanism has significantly higher expected efficiency than all of the other mechanisms. This confirms that location-based and time-based forms of expression are not redundant. Furthermore, in all of these cases, the LOC and TIME mechanisms both have significantly higher expected efficiency than the BL mechanism. For the anyone group, the only significant difference in expected efficiency is between the BL and LOC/TIME mechanisms. Interestingly, the LOC mechanism had significantly higher expected efficiency than the TIME mechanism for the colleague group (this is probably due to the fact that many of our subjects were comfortable sharing their locations with this group while they were on campus).

The results presented in Figure 6 clearly show that the most commonly used privacy mechanism for web-based location sharing services, the black list mechanism, is too simple to capture users' complex privacy preferences. By replacing this mechanism with a more expressive one, these services would be able to better capture the privacy preferences of their users.

Expressiveness is more important when information is more sensitive. Our second

[8]The exhaustive search for expressions decomposes in a straightforward way since each group, time, location and location/time pair can be considered independently. For example, a subject's utility for sharing a particular location does not depend on the other locations he or she has decided to share.

[9]We used a non-parametric bootstrap method to test for statistical significance between means with 95% confidence [38].

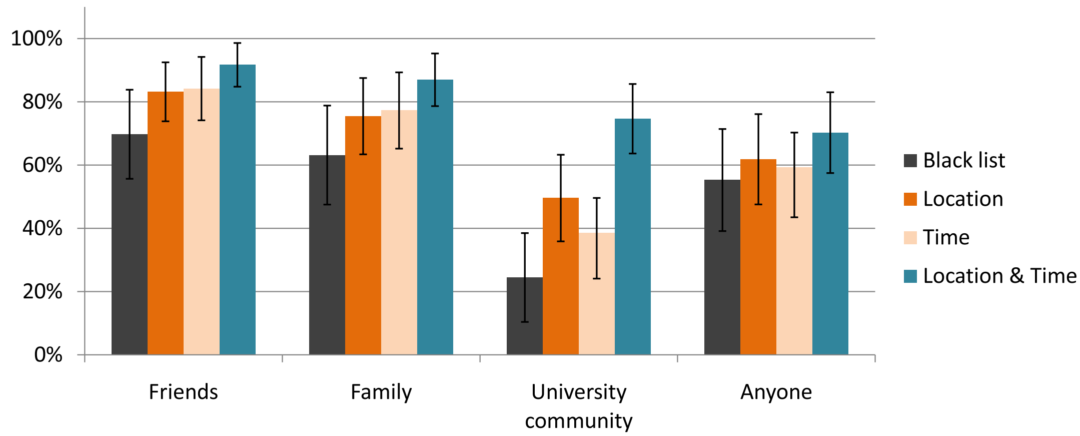

Figure 6: The percent of optimal expected efficiency (bars indicate 95% confidence intervals) achieved by the different mechanisms we tested broken down by group. These results assume that the cost for inappropriately revealing a location is $c = 5$, that the reward for appropriately revealing a location is $r = 1$, and that subjects would have made the best possible expression to each mechanism.

set of results explores the impact of varying the cost associated with inappropriately giving out a subject's location information. For this analysis we restrict our attention to the university community group, since preferences regarding this group were the most diverse. However, our findings with respect to this analysis were similar for all of the other groups.

Figure 7 shows that the efficiency of each mechanism drops as the cost of inappropriately revealing one's location increases. As this cost goes up subjects would be forced to make more restrictive expressions (e.g., by hiding more of their locations), and would receive lower utility from using the mechanism. However, as the mechanisms become more expressive their expected efficiency deteriorates far less rapidly. This is because more expressive mechanisms allow subjects to make more precise expressions. In the location and time-based mechanism, subjects would be able to avoid specific times or locations that are sensitive while still revealing substantial amounts of information when appropriate.

Discussion

Based on this research, we see that there is a need for greater levels of expressiveness in the design of privacy controls for location-sharing systems. While the efficacy of rules increases, so does the number of rules needed. We see that, in most cases, the efficacy of time or location-based rules is the same, but there is a much smaller burden on the the user to create a smaller number of time-based rules ($M = 6.20$) as compared to location-based rules ($M = 17.80$).

Also interesting to note is the similarity to how users feel about location and time-based rules and the similar efficacy of both. Our participants indicated that they were equally comfortable sharing their locations with groups of people with either location or time-based rules.

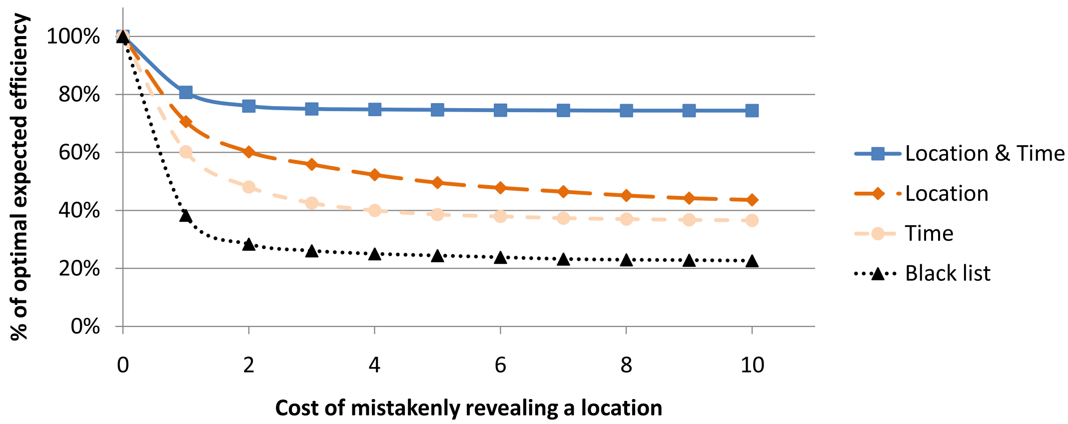

Figure 7: The percent of optimal expected efficiency achieved by the different mechanisms we tested for the "Colleagues" group. For these results we varied the cost associated with inappropriately sharing a location from $c = 0$ to $c = 10$. We assumed that the reward for appropriately revealing a location was fixed at 1, and that subjects would have made the best possible expression to each mechanism based on c.

5 Related work

Prior to our original work on expressiveness in mechanisms [5], there had been relatively little work on expressiveness specifically. We will discuss some related papers in the body of this paper. Here we will briefly summarize existing location-sharing services and other applications that have benefited from increased expressiveness.

5.1 Location-sharing services

Location-sharing services are very much in vogue, and are anticipated as part of the expected billions of dollars in marketing revenue from location-based services [13]. Despite the number of location-sharing applications that have been developed, none have yet to capture significant market share, and many people are still wary of sharing their locations online due to privacy concerns [2, 25].

Many research groups have have developed location-based services: PARC's Active Badges [37], ActiveCampus [3], MyCampus [29], Intel's PlaceLab [15], and MIT's iFind [19]. However their focus has been on increasing the accuracy of reported locations.

The commercial location-sharing services in existence currently include Loopt[10] and Google's Latitude[11], as well as several other, less successful offerings. Loopt allows users to create a person-specific whitelist to control who has access to the user's information. Similarly, Latitude also uses a whitelist, but allows location-sharers to set either exact location or city-level granularity for the locations that they share. Similarly, technology developers

[10]Loopt. http://loopt.com/

[11]Latitude. http://www.google.com/latitude

are allowing third parties to leverage location information to develop applications for their mobile phones; these platforms include the iPhone SDK,[12] and Google's Android SDK,[13]. Similar services are offered for online applications including Skyhook Wireless's web application Loki[14] (which facilitates Wi-Fi positioning) and Yahoo's FireEagle that facilitates privacy-enhanced location-sharing on a per-service level of control.[15]

To explore privacy concerns surrounding the sharing of location information, diary studies and laboratory experiments [3, 8, 28], small group testing [2, 20, 35], and interviews [14, 17, 21] have all been used extensively. This research emphasizes the concrete nature of the privacy concerns people have regarding their location information, finding that the context of a request is key to the willingness of someone to share their location [6, 7, 20, 23, 24, 35], and that having people create groups is a feasible method of access control [18, 28]. A field study of a location-sharing system found that having feedback, or being provided with information on who had viewed your location, had a significant impact in how comfortable people were with sharing their information with friends and strangers, and on reducing participant's levels of privacy concerns after using the location-sharing technology [36].

5.2 Applications of expressiveness

One of the first applications to benefit from expressiveness was strategic sourcing. Sandholm [31, 32] described how building more expressive mechanisms—that generalize both CAs and multi-attribute auctions—for supply chains has saved billions of dollars that would have been lost due to inefficiency. Success with expressive auctions in sourcing has also been reported by others [9, 16, 26].

Some work on expressiveness has begun to appear in the context of search keyword auctions (aka sponsored search). Benisch, Sadeh and Sandholm directly addressed the question of expressiveness in this domain [4]. They showed that adding slightly more expressiveness to traditional ad auction mechanisms, in the form of an extra bid for premium slots, leads to a significant efficiency improvement for some simulated advertiser preferences. Even-Dar, Kearns and Wortman examined an extension of sponsored search auctions, whereby bidders can purchase keywords associated with specific contexts [10]. Under certain probabilistic assumptions they are able to prove that the system becomes more efficient when this extra level of expressiveness is allowed. In a working paper, Milgrom explores the equilibria of sponsored search auctions with limited expressive power [27]. He finds that by *limiting* expressiveness the auction excludes some bad equilibria. This raises an important counterpoint to our work. In another recent paper on sponsored search auctions, Abrams et. al. studied the impact of inexpressive bids on efficiency [1]. They show that an inexpressive mechanism can have an efficient *full information* Nash equilibrium even when bidder valuations are complex.

[12] iPhone Dev Center. http://developer.apple.com/iphone/

[13] Android. http://code.google.com/android/

[14] Loki. http://loki.com/

[15] Fire Eagle http://fireeagle.yahoo.net/

Another application area that has received recent attention with regard to expressiveness is wireless spectrum trading. For example, Gandhi *et al.* [11] described a prototype wireless spectrum market mechanism. They stressed the importance of allowing spectrum bidders enough expressiveness to communicate their needs, and demonstrated—using synthetic demand distributions and various *ad hoc* bidder behavior models—that their mechanism has good efficiency properties.

6 Conclusions and future work

Over the past few years we have seen an explosion in the number and different types of websites that allow individuals to exchange personal information and content that they have created. These sites include online social networks, photo and video-sharing sites, and location-sharing services on the Internet. While there is clearly a demand for users to share this information with each other; recently, we have started to see a change in attitude, with users demanding greater control over the conditions under which their information is shared. In this paper, we looked in particular at the need for expressiveness in privacy policies allowing users to control the conditions under which they are willing to share their locations with others. Our results suggest that as web sites begin to expand their privacy controls, it is imperative that they include expressiveness that captures their user's true preferences.

While existing commercial applications in this space rely primarily on blacklists, results obtained based on a study involving 30 users carrying location-enabled cell phones for a week suggest that users of location sharing applications could benefit from richer privacy policies. The research reported in this paper combines an empirical work with the introduction of a new theoretical frameworks that enables one to quantify the potential increases in efficiency associated with different levels of expressiveness in policies. While the results reported in this article focus on privacy policies in the context of location sharing applications, our theoretical framework and methodology can easily be adapted to other types of security and privacy policies and offers a methodology for comparing the benefits associated with different levels of expressiveness in policies. Clearly, as policies become more expressive, users may have to spend more time specifying their preferences, if they are to take full advantage of the expressiveness of the policies they are given access to. Different interface technologies, such as expandable grids or user-controllable policy learning, as well as better use of user-centered design principles offer the prospect of potentially mitigating these tradeoffs.

One interesting finding from out work is that blacklist-based policies such as the ones used in most location sharing applications available today are very limited in their ability to capture peoples location sharing preferences. As users tend to err on the safe side as they specify their privacy policies (generally prefering not to reveal their location across a broader range of scenarios rather than risk sharing their location under a small set of situations that they may not be comfortable with), the effect of inexpressive location sharing policies generally gives rise to usage scenarios where location sharing remains very limited (e.g. see our own research as reported in [22]). More expressive policies offer users the ability to better qualify the conditions under which they share their location with others, thereby generally resulting in more sharing overall. This in turn means that inexpressive location sharing mechanisms are likely to result in less location sharing. This lack of expressiveness is in our view one reason why the many location sharing applications deployed so far continue to see fairly limited use: they just do not offer their users sufficient value.

Our empirical results confirmed that i) most subjects had relatively complex privacy preferences, and ii) that privacy mechanisms with higher levels of expressiveness are significantly more efficient when information is sufficiently sensitive. Thus, the fact that most location sharing services use simple black list mechanisms, which do not match the privacy preferences revealed in our study, may help explain the lack of broad adoption encountered by these applications so far.

The findings in this paper open several avenues for future work. We can explore additional dimensions of expressiveness, such as allowing expressions based on the day of the week, or the resolution at which the location information is provided (e.g., neighborhood, city, or state). Future work should also address the increase in user burden associated with increasing expressiveness. This increase in user burden could potentially lead to a discrepancy between a mechanism's optimal efficiency and the actual efficiency achieved by real users.

7 Acknowledgments

This work has been supported by NSF grants CNS-0627513 and IIS-0427858. Additional support has been provided by Nokia, France Telecom, Nortel, the CMU/Microsoft Center for Computational Thinking, ARO research grant DAAD19-02-1-0389 to Carnegie Mellon University's CyLab, and the CMU/Portugal Information and Communication Technologies Institute. The authors would also like to thank Lucian Cesca, Jialiu Lin, Tony Poor, Eran Toch, and Kami Vaniea for their assistance with our study.

References

[1] Zoe Abrams, Arpita Ghosh, and Erik Vee. Cost of conciseness in sponsored search auctions. In *Proceedings of Workshop on Internet Economics (WINE)*, 2007.

[2] L. Barkhuus, B. Brown, M. Bell, M. Hall, S. Sherwood, and M. Chalmers. From awareness to repartee: Sharing location within social groups. In *CHI '08*, pages 497–506, April 2008.

[3] L. Barkhuus and A. Dey. Location-based services for mobile telephony: a study of users' privacy concerns. In *INTERACT'03*, pages 702–712, 2003.

[4] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. The cost of inexpressiveness in advertisement auctions. In *Proceedings of ACM EC Workshop on Advertisement Auctions*, 2008.

[5] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. Theory of expressiveness in mechanisms. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 2008.

[6] B. Brown, A. Taylor, S. Izadi, A. Sellen, J. Kaye, and R. Eardley. Location family values: A field trial of the whereabouts clock. In *Ubiquitous Computing (Ubicomp '07)*, pages 354–371. Springer-Verlag, 2007.

[7] S. Consolovo, I. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location disclosure to social relations: Why, when, & what people want to share. In *CHI '05*, 2005.

[8] Jason Cornwell, Ian Fette, Gary Hsieh, Madhu Prabaker, Jinghai Rao, Karen Tang, Kami Vaniea, Lujo Bauer, Lorrie Cranor, Jason Hong, Bruce McLaren, Mike Reiter, and Norman Sadeh. User-controllable security and privacy for pervasive computing. In *Proceedings of the 8th IEEE Workshop on Mobile Computing Systems and Applications*, 2007.

[9] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[10] Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Sponsored search with contexts. In *Workshop on Internet Economics (WINE)*, 2007.

[11] Sorabh Gandhi, Chiranjeeb Buragohain, Lili Cao, Haitao Zheng, and Subhash Suri. A general framework for clearing auctions of wireless spectrum. In *IEEE DySPAN*, 2007.

[12] Marta C. Gonzalez, Cesar A. Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.

[13] Kelsey Group. BIA's The Kelsey Group Forecasts U.S. Mobile Local Search Advertising Revenues to reach $1.3B in 2013, February 24 2009.

[14] R. H. Harper. Why people do and don't wear active badges: A case study. In *In Proceedings of Computer Supported Cooperative Work (CSCW96)*, pages 297–318, 1996.

[15] Jeffrey Hightower, Anthony LaMarca, and Ian E. Smith. Practical lessons from place lab. *IEEE Pervasive Computing*, 5(3):32–39, 2006.

[16] Gail Hohner, John Rich, Ed Ng, Grant Reid, Andrew J. Davenport, Jayant R. Kalagnanam, Ho Soo Lee, and Chae An. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces*, 33(1):23–35, 2003.

[17] Jason I. Hong. *An Architecture for Privacy-Sensitive Ubiquitous Computing*. 2005.

[18] G. Hsieh, K. Tang, W. Low, and J. Hong. Field deployment of imbuddy : A study of privacy control and feedback mechanisms for contextual im. In *Ubiquitous Computing (Ubicomp '07)*, pages 91–108, 2007.

[19] S. Huang, F. Proulx, and C. Ratti. iFIND: a Peer-to-Peer application for real-time location monitoring on the MIT campus. In *CUPUM 07 - 10th International Conference on Computers in Urban Planning and Urban Management*, July 11-13 2007.

[20] G. Iachello, I. Smith, S. Consolovo, G. Abowd, J. Hughes, J. Howard, F. Potter, J. Scott, T. Sohn, J. Hightower, and A. LaMarca. Control, deception, and communication: Evaluating the deployment of a location-enhanced messaging service. In *UbiComp 2005*, pages 213 – 231. Springer-Verlag, 2005.

[21] E. Kaasinen. User needs for location-aware mobile services. In *Personal and Ubiquitous Computing 2003*, pages 70–79, 2003.

[22] Patrick Gage Kelley, Paul Hankes Drielsma, Norman Sadeh, and Lorrie Faith Cranor. User-controllable learning of security and privacy policies. In *AISec '08: Proceedings of the 1st ACM workshop on Workshop on AISec*, pages 11–18, New York, NY, USA, 2008. ACM.

[23] A. Khalil and K. Connelly. Context-aware telephony: Privacy preferences and sharing patterns. In *CSCW '06*, 2006.

[24] S. Lederer, J. Mankoff, and A. K. Dey. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *CHI '03*, number 724-725, 2003.

[25] C. McCarthy. The mobile social: Not ready for prime time?, February 13 2008. http://www.news.com/8301-13577_3-9870611-36.html.

[26] Theresa Metty, Rob Harlan, Quentin Samelson, Tom Moore, Thomas Morris, Ron Sorensen, Avner Schneur, Olga Raskina, Rina Schneur, Joshua Kanner, Kevin Potts, and Jeffrey Robbins. Reinventing the supplier negotiation process at Motorola. *Interfaces*, 35(1):7–23, 2005.

[27] Paul Milgrom. Simplified mechanisms with applications to sponsored search and package auctions. Working paper, 2007.

[28] S. Patil and J. Lai. Who gets to know what when: Configuring privacy permissions in an awareness application. In *CHI '05*, pages 101 – 110, 2005.

[29] Norman Sadeh, Fabien Gandon, and Oh Buyng Kwon. Ambient intelligence: The my-Campus experience. Technical Report CMU-ISRI-05-123, Carnegie Mellon University, July 2005.

[30] Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, and Jinghai Rao. Understanding and capturing people's privacy policies in a people finder application. *The Journal of Personal and Ubiquitous Computing*, 2008. Forthcoming.

[31] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted $35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.

[32] Tuomas Sandholm, David Levine, Michael Concordia, Paul Martyn, Rick Hughes, Jim Jacobs, and Dennis Begg. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces*, 36(1):55–68, 2006.

[33] Barry Schwartz. *The Paradox of Choice: Why More is Less*. Ecco/HarperCollins Publishers, 2004.

[34] Herbert A Simon. *Models of bounded rationality*, volume 2. MIT Press, 1982.

[35] I. Smith, S. Consolovo, A. LaMarca, J. Hightower, J. Scott, T. Sohn, J. Hughes, G. Iachello, and G. Abowd. Social disclosure of place: From location technology to communication practices. In *Pervasive '05*, pages 134 – 151. Springer-Verlag, 2005.

[36] Janice Tsai, Patrick Kelley, Paul Hankes Drielsma, Lorrie Faith Cranor, Jason Hong, and Norman Sadeh. Who's viewed you? the impact of feedback in a mobile-location system. In *Proceedings of Computer Human Interaction 2009 (under review)*, 2009.

[37] Roy Want, Veronica Falcão, and Jonathan Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10:91–102, 1992.

[38] A. M. Zoubir. Bootstrap: Theory and applications. In *Proceedings of the SPIE 1993 Conference on Advanced Signal Processing Algorithms, Architectures and Implementations*, pages 216–235, July 1993.

8 Appendix

In a privacy mechanism, an impact vector describes the impact of a particular expression by an agent under all possible requests that could be placed for the agent's information.

Definition 1 (impact vector)**.** *An* impact vector *is a function, $g : \vec{A} \to \{0, 1\}$. To represent the function as a vector of outcomes, we impose some strict order on the possible requests in $\vec{A}$, then g can be represented as $\{0, 1\}^{|\vec{A}|}$.*

We say that an agent can *express* an impact vector if there exists at least one expression that the agent can make in order to cause each of the outcomes in the impact vector to be chosen by the mechanism.

Definition 2 (express)**.** *An agent can* express *an impact vector, g, if* $\exists\theta, \quad \forall\vec{a}, \quad f(\theta,\vec{a}) = g(\vec{a})$.

We say that an agent can *distinguish* among a set of impact vectors if it can express each of them by changing its expression under the same collection of possible requests.

Definition 3 (distinguish)**.** *An agent can* distinguish *among a set of impact vectors, G, if* $\forall g \in G, \quad \exists\theta, \quad \forall\vec{a}, \quad f(\theta,\vec{a}) = g(\vec{a})$. *When this is the case we write $D(G) = \top$.*

The adaptation of the impact dimension measure for the privacy mechanism setting captures this intuition; it measures the number of different impact vectors that an agent can distinguish among.

Definition 4 (impact dimension)**.** *A privacy mechanism has* impact dimension *d if the largest set of impact vectors, G^*, that an agent can distinguish among has size d. Formally,*

$$d = \max_G \left\{ |G| \;\middle|\; D(G) = \top \right\}$$

Theorem 1. *For any utility function, distribution over agent types, and distribution over request attributes, the expected efficiency (given in equation 1) for the best privacy mechanism limiting an agent to impact dimension d increases* strictly *monotonically as d goes from* 1 *to d^*, where d^* is the minimum impact dimension needed to reach full efficiency.*

Proof. The set of mechanisms with impact dimension d is a super-set of the mechanisms with impact dimension $d' < d$. Thus the fact that the efficiency for the best mechanism increases weakly monotonically is trivially true. The challenge is proving the strictness of the monotonicity.

Consider increasing d from $d^{(1)} < d^*$ to $d^{(2)} > d^{(1)}$. Let $G^{(1)}$ be the best set of impact vectors that an agent could distinguish between when restricted to $d^{(1)}$ vectors (i.e., the set of impact vectors that would maximize the mechanism's expected efficiency). We know that there are at least $d^* - d^{(1)} \geq 1$ impact vectors needed to reach full efficiency that cannot be expressed, and thus at least that many impact vectors that are absent from $G^{(1)}$. When we increase our expressiveness limit from $d^{(1)}$ to $d^{(2)}$, we can add one of those missing vectors to $G^{(1)}$ to get $G^{(2)}$. Since $G^{(2)}$ allows an agent to distinguish among all the same vectors as $G^{(1)}$ *and* an additional vector which corresponds a more efficient set of outcomes, the new mechanism with impact dimension $d^{(2)}$ has a strictly higher expected efficiency. □

Theorem 2. *There exists a utility function, a distribution over types, and a distribution over request attributes such that the best privacy mechanism limited to impact dimension d is arbitrarily less efficient than that of the best privacy mechanism limited to impact dimension $d + 1 < d^*$, where d^* is the minimum impact dimension needed for full efficiency.*

Proof. Since an agent's utility function can depend arbitrarily on its type and the attributes of a request, we can construct a scenario in which the agent requires impact dimension at least $d + 1$ or it will experience an arbitrarily high cost. First we must ensure that the agent has at least $d+1$ types with non-zero probability. Next we choose a set of impact vectors, $G^{(1)}$, of size $d+1$. For each of the distinct impact vectors in $G^{(1)}$ we can ensure that it gives the agent arbitrarily more utility than all other impact vectors for at one of the agent's types. By the pigeon hole principle, the agent will be unable to express at least one of the impact vectors in $G^{(1)}$ in any mechanism with impact dimension d. Thus increasing a limit on impact dimension from d to $d+1$ will lead to an arbitrary increase in efficiency. $\square$