When Privacy and Utility are in Harmony: Towards Better Design of Presence Technologies

Jacob T. Biehl, Eleanor Rieffel
FX Palo Alto Laboratory, Inc.
Palo Alto, CA, USA
{biehl, rieffel}@fxpal.com

Adam J. Lee
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA, USA
adamlee@cs.pitt.edu

1 Abstract

Presence systems are valuable in supporting workplace communication and collaboration. These systems are effective, however, only if they are widely adopted and candidly used. User perceptions of the utility of the information being shared and their comfort in sharing such information strongly impact both adoption and use. This paper describes the results of a survey of user preferences regarding comfort with and utility of sharing presence data in the workplace; the effects of sampling frequency, fidelity, and aggregation; and design implications of these results. We present new results that extend some past findings and challenge others. We contribute new insights that inform the design of workplace presence technologies to increase both the utility and adoption of these systems.

2 Author Keywords

Presence systems, privacy, sharing, collaboration.

3 Introduction

Presence systems have proven to be valuable in supporting communication and collaboration in the workplace [4,5,24,28,32]. They function by fusing physical sensing capabilities with social and communication software to provide increased awareness of colleagues' physical presence, ongoing activities, and available communication channels. In order to prove useful and effective, these systems must have wide adoption and use within an organization to. Decisions regarding privacy and utility made during the design and deployment process can significantly impact their adoption and use. Widely publicized privacy issues with popular services such as Facebook and Google Buzz illustrate the complexity of this design space [e.g. 1,7]. As Patil et al. found, "user opposition due to privacy concerns can translate into minimal use or even the abandonment of the system" [19]. Furthermore, users who are uncomfortable with a system or its features may develop means of circumventing the system. Such actions waste users' time and make the system itself less useful, since its data become suspect. Patil et al. found that "users who were forced to use instant messaging … often resorted to circumvention tactics." Others "set their status to 'away' or 'busy' even when they were not", "or changed their preferences so as to appear online even when they were away from their desks." They point out that "such underuse or circumvention results in suboptimal use of awareness systems."

Understanding expectations and preferences with respect to the tension between privacy and information sharing has been explored previously [e.g. 2,10,20]. Furthermore, many studies of *location* sharing and privacy have been conducted [e.g. 3,6,9,11,12,17,30]. Nevertheless, these preferences and expectations remain poorly understood, especially in the workplace. Many of these studies were done on student populations, while others focused on social, rather than workplace, interactions. As such, much remains to be understood about presence and context sharing within the workplace. As existing work has shown, there is great potential for workplaces to benefit from presence technologies. Unfortunately, the limited understanding of the design space for these types of systems has resulted in few systems being used in practice, let alone commercially developed.

In this paper, we seek to better characterize this design space by surveying and analyzing user sentiments regarding presence data collection and sharing specifically *in the workplace*. We chose a survey-based approach for this investigation, as it allows us the greatest coverage of topics and diversity of participation. The survey covered a multitude of dimensions, including not only sentiments with respect to specific sensors and collection mechanisms, but also how frequency, fidelity, and grouping of data impact both comfort and utility. To our knowledge, this investigation is one of the first to examine all of these dimensions collectively while focusing specifically on workplace settings. Furthermore, it is one of the first studies to solicit user's thoughts not only on what types of information they would be willing to share, but also what information they would find useful for others to share in each context. This methodology resulted in many high-level insights, including:

- Users are concerned with the entire data pipeline. In particular, they are concerned about the types of data that are collected, not just which data are eventually shared.
- User preferences often follow bimodal distributions along a comfort/discomfort continuum. This observation suggests that finding a single, appropriate default data collection or sharing setting for all users is often not a good design.
- Users' comfort in sharing certain types of information often mirrors their perception of its utility. Thus, presence systems that can estimate use, and share information accordingly, will be preferred.
- Aggregation at the user or group-level can increase comfort while having little or no negative impact on utility.
- Users were substantially more uncomfortable having their presence data stored than they were with its initial collection and sharing. Preferences for data collection and dissemination do not always carry over to long-term storage.

These findings have led to the development of a new set of design implications. As we discuss below, many of the decisions currently being made in the design of workplace presence systems are too simplistic and should be reexamined, and in some cases, entirely rethought.

4 Related Work

Many studies explore privacy preferences and concerns. Iochello and Hong have an excellent survey of privacy in the context of human computer interaction [13]. This section discusses recent work most closely related to our concerns, work that focuses on privacy with respect to technologies that share presence or location information.

While we found a lot of prior work looking at what people are willing to share, we found little work examining what shared information people find useful. Davis and Gutwin [12], like us, asked participants what they would choose to gather about someone else as well as what they were willing to disclose about themselves. They considered seven roles (superior, peer, spouse, secretary, friend, subordinate, and acquaintance) and found that while willingness to share varied with role, the information their participants wished to gather did not vary with role. Our findings

differ from theirs in that we found differences in perceived usefulness of gathering information depending on role. We also distinguish close colleagues from distant colleagues, instead of grouping the two together under peer, and similarly we distinguish between direct manager and higher-level executives instead of grouping them together under superior. In both cases, we found significant differences between these groups that lead to important design recommendations for workplace awareness systems. Davis and Gutwin did find, as we did, that their participants' preferences for the fidelity of information depend on role for both sharing information about themselves and gathering information about others. While the preferred fidelity for gathering information was somewhat higher than for disclosing information, the two were close, and participants mentioned that they were unwilling to gather high-fidelity information in case where it seemed like an intrusion on their part or where they felt they did not want to learn things that were not public.

Consolvo *et al.*'s study [9] is similar to our work in context and setting: their sixteen participants were not students, and some of the survey questions were work oriented. The intent and design of their study were quite different from ours, however. In the core part of their study, the participants received 10 randomly timed, hypothetical requests each day for two weeks nominally from someone on a buddy list they had previously specified. The buddy lists consisted of up to 17 people that included at least two coworkers, at least one manager, and at least two family members. The participants then chose how to respond to each request. In our setting, we were interested in situations in which users were broadcasting their information through their participation in a presence system rather than responding to requests. Thus, in the setting considered in our survey, users would have less control over what information a given individual received at a given time, but more control over general settings including what data would be collected in the first place. Also, we investigated participants perceived usefulness in receiving information about others, as well as their comfort in sharing it. Furthermore, we asked participants about their comfort with presence data being stored, while Consolvo et al. did not. Consolvo *et al.* found that the level of detail participants disclosed depended on what level of detail the participants thought would be useful to the responder. Participants "chose to disclose less specific information because they thought something less specific would be most useful to the requester and not because they were uncomfortable giving the requester more detailed information." Their results lend support to our contention that one reason privacy and utility can be aligned is that less detailed information can be more useful than more detailed information.

Consolvo et al. found that the most important factors determining sharing preferences were "who was requesting, why the requester wanted the participant's location, and what level of detail would be most useful to the requester." Like us, they found a significant difference between sharing with coworkers and with managers, but the coworkers in question were on their buddy list of people with whom "they might want to exchange location information," so they did not examine the difference between coworkers participants work closely with and coworkers they do not. Similarly, they did not distinguish between direct managers and higher-level executives. Many of their participants rejected location requests if they felt that the requester did not have a good reason to need to know. They found that, when at work, participants disclosed location to coworkers 80% of the time, and to managers 69% of the time. The willingness to disclose dropped for other settings. Co-workers were willing to disclose location, for example, to coworkers 47% of the time they were at home, but only 24% of the time to a manager.

Like Consolvo, Lederer *et al.* [17] ask about hypothetical location requests, and found that sharing was based largely on who sent the request. They recruited from both UC Berkeley engineering students and "from websites across the US." They do not report what fraction of their participants were students. They only considered reporting from two types of events (working lunch and social evening) and four inquirers (spouse/significant other, employer, stranger, merchant). Our study looks at a wider variety of situations and provides more fine-grained information with respect to work relationships. They note that "when the inquirer is the subject's employer, situation becomes

a stronger determinant" than the requester, suggesting that finer grained inquiry is needed to understand privacy issues in a work setting. Indeed, our study shows that a more fine-grained investigation of work relationships was warranted.

Brush *et al.* [6] logged GPS data for 32 people in 12 households over a 2 month period who were paid for their data and also received software gifts. They then presented the participants with some of their data, and five different obfuscations of that data. For each obfuscation, people were asked whether they were willing to share it publicly by name, publicly anonymously, anonymously with Microsoft corporate partners, anonymously with researchers, or not at all. They also investigated how much participants valued their data, by asking if they would be willing to trade it for location-based services that use such data, and also how much they would want to be compensated in return for another months' collection of data. Their focus was on comparing obfuscation methods and determining the value placed on the data. Their study was not workplace oriented, unlike ours, and their investigation of sharing preferences considered only large, impersonal groups, with whom the participants had little relation.

Kaasinen [15] studied 55 people in Finland, a diverse group in terms of hometown, age, interests, and occupation. They focused on user perceptions of privacy issues with regard to location aware services such as locating nearby services or attractions. The found that while "people were worried about their privacy and the 'big brother' phenomenon when considering services enabling people to be located," they were not worried about privacy issues with location-aware services. "It did not occur to most of the interviewees that they could be located while using the service." Participants also expected that there would be regulations protecting them. We expect that because our study was carried out seven years later, included only participants with at least a minimum of experience with such technologies, and was focused on presence system intended to disclose location information, it would be difficult to generalize from their results to ours.

Many studies have been done entirely (e.g. [3,11,18,22]) or mostly (e.g.16,25,29,30]) on students. There are many reasons to believe that results of these studies do not carry over to more general populations, and to workplace settings in particular. For example, Danezis *et al.* [11] found that the median values for privacy concern were higher for the few students who traveled outside their college town once a week or more, and among the approximately one-third of the students who used their phones to talk with significant others. These trends suggest the results would likely have been quite different in non-student populations, such as business employees.

Khalil and Connelly [16] look at sharing and privacy issues related to the handling cell phone interruptions. They simulated cell phone interruptions over the course of a ten day study in which users were subject to a simulated interruption 13 times a day and were asked to fill out a questionnaire about what information they would be willing to share in response to that interruption. They were interested in the sharing of contextual information such as location, activity, who else was around, and who is speaking, and how what is shared is affected by the type of person. They did not ask the reciprocal question of what contextual information they would find useful receiving. Their participants were mostly students who had a part-time or full-time job so were also employees. The categories of people they considered were significant other, family friend, colleague, boss, and unknown. Again, they did not distinguish between distant and close colleagues, or immediate manager versus higher-level executives.

Barkhuus [3] found that who is asking the question is not so important, in contrast to the findings of Lederer *et al.*, and suggest that the reason may be that the study uses "a physically limited application and the potential enquirers very likely have some relation to the campus environment; this is likely the reason that some of the users do not worry about letting all possible users see their location." For this reason, the results of Barkhuus provide only limited insight into the settings of interest to us. Toch *et al.* [30] found that "users who are recorded at a large number of unique locations generally evolve more complex privacy preferences but also report finding location

sharing more useful." This finding supports our claim that privacy preferences may be quite different in the work setting than among students, and that awareness systems may also be more useful in a work setting than in a campus setting, depending on the work style and the granularity of the locations revealed. Lin *et al.* [18] studied granularity in location disclosure, and found that it varied with social relationship among other factors. They considered four relationship types (close friends, friends, acquaintances, strangers) and did not consider work relationships.

Ravichandran *et al.* [22] provided 30 undergraduate and graduate students with a location-tracking enabled mobile phone for a week, and presented the students with questions at the end of each day as to with whom they were willing to share specific pieces of data. Their main interest was in studying whether clustering techniques could be used to learn default policies. They comment that there is a marked decrease in error when going from one default policy to two default policies, supporting our point that a single default is not effective for many applications involving privacy. Sadeh *et al.* [25] showed that people have a hard time expressing effective privacy policies, and provided means (rule-based systems; interfaces; machine learning) to try to enable users to more easily specify and refine policies. Their work complements ours.

Tang *et al.* [29] looked at social-driven versus purpose-driven location sharing, and found significant differences in preferences. Our setting is purpose-driven, suggesting the need for plausible deniability, real-time feedback, and audit logs. While not our main focus, our results confirm the desire for such features among non-student, as well as student, participants.

5 Survey

Motivated by a desire to build a deeper understanding of users' sharing preferences within presence technologies, we designed and executed a user preference survey that examined these concerns across several dimensions. Our multi-faceted investigation sought to understand:

- concerns and preferences with respect to specific sensing technologies (e.g. GPS, cameras),
- differences in concern depending on how data are collected whether the data are stored,
- expectations of purpose for sharing presence information, and the effect of such expectations on sharing preferences,
- differences in concern when sharing with a variety of workplace and social groups, and
- concern and preference changes when presence information is aggregated into higher level presence states, shared with decreased fidelity, or aggregated across users/groups.

While a common and acceptable research practice, a survey-based approach has the obvious downside in that it gathers participants' *perceived* sentiments and reasoning, rather than actual use and behavior. The alternative is to build and deploy software where users can be observed and measured in their use of the tools. Because our focus was on building a broad understanding, across multiple dimensions of concern, comfort, and sentiments of use, a build-and-deploy methodology would be impractical as we would need to build, deploy, and study a variety of systems that each embody one particular instantiation of a function or policy that we sought to understand better.

Questions were a mix of multiple-choice and free-form. For instance, a multiple-choice question that appeared on our survey was:

> *Sensors and data sources can be sampled (checked for their current reading) at many different frequencies. For the following sensor and data sources, please indicate the maximum frequency that you would be comfortable with presence data from this source being sampled.*

Not at all Daily Hourly Every 15 minutes Once a minute Constant sampling

While a free-form question was:

> *Under what conditions would your feelings change about how frequent your presence data is accessed? For instance, would access by a particular individual or particular type of individual elevate concern? Please describe these situations – no need to provide specific names or events.*

Overall we asked 50 multi-part questions, which resulted in a total of 242 individual data points being collected *for each participant* excluding demographic data. See the online supplemental material for this article for the complete 50-question survey.

In choosing participants, our goal was to sample a broad spectrum of potential presence technology users. A total of 32 participants, 16 male, 16 female, were recruited through online and in-person solicitations. Unlike much past work that focused on undergraduate and graduate students, we sampled widely across age (mean=32.6, S.D.=7.7), profession, and even US geographic region (West, Midwest, Southeast, and East). It is important to note that participants were not all associated with any one company or organization (e.g., university). While smaller than some surveys (n=32), our population was well diversified and our analysis showed statistical significance within this population for the findings and underlying measures we report. A larger population would not have impacted the statistical significance of our results.

In addition to the survey checks mentioned above, we also placed conditions on participation and screened users so that we could ensure that each participant had a moderate understanding of Internet and location technologies. As a result, all of our participants were familiar with things like online banking, social networks, and various sensing technologies such as GPS. We took careful steps to design the survey such that specific questions sought not only to gain sentiment, but also to verify that the participants understood the questions they were being asked. Use of sensing technologies was not required, as we wanted a representative sample across use and non-use. Nevertheless, 93.8% (30/32) indicated having used online banking, 87.5% (28/32) used a location-based service at least monthly, 84.4% (27/32) used a social networking site at least monthly, and 28.1% (9/32) used a social networking site multiple times per day. At the end of our survey we asked whether the participants had been victims of identity theft (15.6%, 5/32), knew a victim (43.8%, 14/32), or had been stalked or watched (28.1%, 9/32).

The survey was administered online and required users to complete the survey in one session. All participants took between 45-60 minutes to complete the survey. Each participant received a $25 gift card upon completing they survey.

5.1 Results

We first discuss results concerning general properties of the preference distributions we obtained, noting that many of the distributions were bimodal (F1). We found alignment, as opposed to conflict, between comfort and utility for both frequency and fidelity, and for both traditional forms of sharing (calendars) and novel forms (presence states) (F2,F3,F5). We confirmed dependence of sharing preferences on location (F4). Role and relationship influence both comfort and utility (F6) and the fidelity of shared information and expectations of use (F7). Participants exhibited a high degree of concern about storage of raw sensor data (F8), and a lesser degree of concern about storage of higher-level presence states (F9). Participants recognized the value of stored data (F12), but were concerned about who owns the data, where it is stored, and who controls access to it (F10, F11). Aggregation of results over a group of users generally increased comfort (F13).

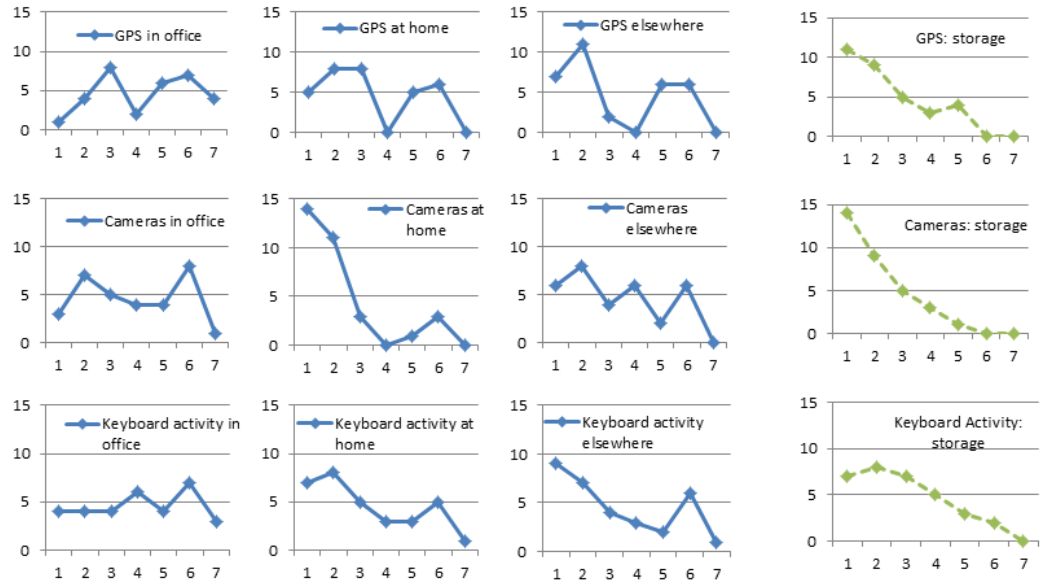

Fig. 1 The first three columns of this chart show user comfort with a specific technology (GPS, cameras, keyboard activity) *sensing* them in the specified location. The x-axis is the seven possible responses (very uncomfortable; uncomfortable; slightly uncomfortable; neutral, slightly comfortable, comfortable, very comfortable). The y-axis is the count for each response. The distribution is bimodal in almost all cases. The fourth (rightmost) column shows user responses to *storing* raw data from different the sensing technologies.

F1. Concern with sensing follows a bimodal distribution

We asked participants to rate their comfort level with three different types of sensing technologies (GPS/WiFi localization, cameras, and keyboard activity) in three different locations (in office, at home, elsewhere). For each sensor type and location pair, participants rated their level of comfort on a scale of 1-7, where 1 was defined as *very uncomfortable* and 7 *very comfortable*.

Concern with respect to the three sensing technologies follows a bi-modal distribution, though where the peaks are situated varied across sensor types and location (Figure 1). As expected, a large number of respondents indicated discomfort for each of the sensing technologies. A fair number of respondents, however, indicated comfort with those same technologies. For instance, for camera-based sensing 28.1% (9/32) of respondents indicated they were comfortable or very comfortable with the use of cameras in the office while the same number (9/32) indicated they were uncomfortable or very uncomfortable. Hence, over half of the responses (18/32) were at the extremes of the comfort scale.

There were significant differences in user comfort for each sensor type depending on location. We used a Friedman test to assess significance because the data are comprised of discrete values, are not normally distributed, and some cell counts are small. The differences were most pronounced for GPS (Friedman test: $\chi^2(2)=24.595$, $p<0.001$), with keyboard activity next (Friedman test: $\chi^2(2)=20.447$, $p<0.001$), and finally cameras (Friedman test: $\chi^2(2)=15.041$, $p<0.001$). To understand where the differences lay, we ran a post hoc Wilcoxon (See Table 1). To find significance at the 0.01 level, with Bonferroni correction, we look for p values of 0.003 or smaller. For GPS, there were significant differences between office and the other two conditions, but no statistical difference between home and elsewhere. For cameras and keyboard activity, there was statistical difference between home and the other two conditions, but not between office and elsewhere.

F2. Alignment between utility and comfort for frequency

		Office vs. Home	Office vs. Elsewhere	Home vs. Elsewhere
GPS	z	-3.452	-3.161	-0.48
	p	0.001	0.002	0.631
Camera	z	-4.046	-1.473	-2.973
	p	<0.001	0.141	0.003
Keyboard Activity	z	-3.57	-0.612	-2.708
	p	<0.001	0.541	0.007

Table 1: Differences in user comfort for each sensor type depending on location

We provided participants with brief scenarios illustrating how presence data could be used to improve communication and collaboration in the workplace. Participants were asked about scenarios in which sharing of *their own* data aided communication and collaboration, and scenarios in which the sharing was of *another person's* data. For instance, one scenario asked the participant to rate comfort in allowing others to access his presence data so that he (the participant) could be located for a face-to-face discussion. In a separate question we asked about the opposite direction: perceived usefulness of someone else's presence data in order to locate them for a discussion. For each scenario, each participant specified the frequency with which he or she was comfortable sharing the data, and later specified the frequency at which they would prefer information be sampled from someone else.

Responses to these questions showed surprising alignment between the frequency participants thought useful and their comfort sharing information at that rate (Figure 2). Participants overwhelmingly indicated that a sampling rate of every 15 minutes would be adequate, although there was some variability. The desired frequencies were strikingly similar for utility and comfort, with no statistical difference with direction of access as an independent factor for *any* of the scenarios. The Spearman correlation was $\rho = 0.692$, $\rho = 0.873$, $\rho = 0.793$, $\rho = 0.679$ for each of the four cases respectively, with significance at $p < 0.001$ in all four cases. This encouraging result provides an interesting case where the often-cited tension between privacy and utility appears to be small or non-existent.

F3. Alignment between comfort and utility for fidelity

We asked participants about sensing fidelity using a scenario-based approach similar to the one used for frequency. Participants indicated what granularity of location information they were comfortable sharing with respect to a variety of common situations that people encounter on a day-to-day basis in their office, on a corporate or university campus, at home, or out and about. Participants could choose one of four location granularities: no information, coarse (e.g., city name), medium (e.g., building name), and precise (e.g., the room). Participants also specified the fidelity they needed to find another person's presence data useful in the reciprocal set of scenarios.

For fidelity, responses for utility and comfort were also strongly aligned for most of the scenarios (Figure 3). Participants' desired level of granularity for information about others closely matched the granularity that they were comfortable providing themselves. For "in personal office," "in building," and "on campus," the Spearman correlation was $\rho = 0.784$, $\rho = 0.802$, $\rho = 0.655$ respectively, all significant ($p < 0.001$). The two non-work scenarios "at home not working" and "out and about" also had strong correlations of $\rho = 0.560$ and $\rho = 0.613$, also significant at the $p < 0.001$ level. For "working from home" and "working offsite" there was still significant correlation ($\rho = 0.435$, $\rho = 0.396$, both $p < 0.05$), but it was less strong than in the other cases. One factor may be that we as a culture are still working out what the expectations should be in these types of non-traditional work environment cases.

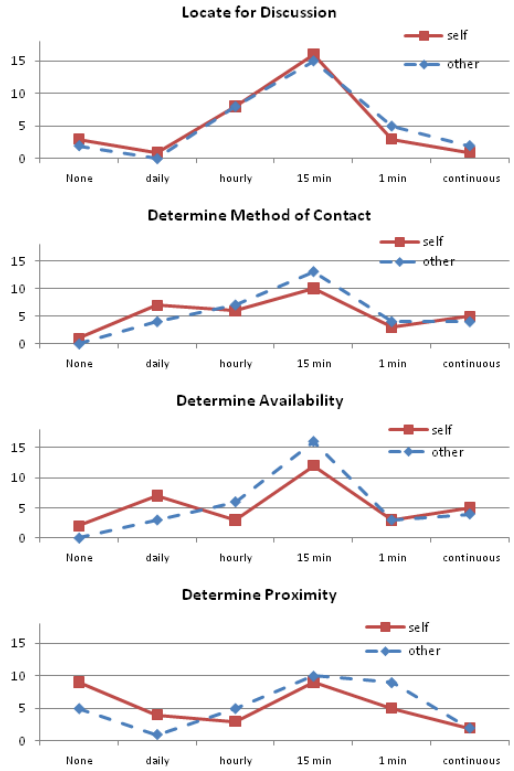

Fig. 2 The red curves show responses to the most comfortable frequency for information sharing for the given purpose. The blue curves show the desired frequency for receiving presence information for that purpose. The x-axis gives the six frequencies we asked about, and y-axis the number of responses.

F4. Fidelity preference depends on location

Participants were comfortable with high fidelity sensing in office-related locations (Figure 3). Many participants (78.1%, 25/32) chose medium-level fidelity or greater for all office-specific conditions (office at work, in building, and on campus). Fidelity ratings were lower for locations outside the office: only 40.6% (13/32) were comfortable with medium-level fidelity or higher when working from home, and 37.5% (12/32) when working off-site. For non-work-related locations, comfort was significantly lower with only 6.3% (2/32) willing to share medium-level or higher fidelity when not working, and 3.1% (1/32), when "out and about."

Fidelity preference is remarkably dynamic. A majority of participants, 53.1% (17/32), indicated comfort with precise, room-level positioning in their office at work, but only 40.6% (13/32) chose room-level fidelity for locations at work but outside his or her office. Comfort with precise positioning dropped even further for "on campus but not in one's normal building," with only 25% (8/32) participants comfortable with precise positioning in this setting. The differences in fidelity preferences across the seven locations are significant (Friedman test: $\chi^2(6) = 129.342$, $p<0.001$).
We found, however, that some fidelity preferences are subtle and may not be completely represented in these data. For example, our results suggest that fidelity preference depends on the meaning or context of the particular location, and that the boundaries do not necessarily align to physical boundaries and may change with context. As one example, room-level precision is preferred when someone is in her own office, but not preferred when she is in someone else's office.

F5. Preferences with respect to traditional presence sharing

We also asked participants about sharing preferences for more traditional presence sources, such as corporate and personal calendars. Participants were presented with generic, typical appointments, such as meeting with one's team, meeting with management, or on vacation. For

each type, participants specified how much information about that appointment should be shared: nothing; time occupied (e.g., busy); time and place of event; or time, place, and list of attendees.

Again there was alignment between the two directions of access (Figure 4). The Spearman correlation was $\rho = 0.847$, $\rho = 0.619$, $\rho = 0.623$, $\rho = 0.657$, $\rho = 0.649$ respectively, all significant ($p<0.001$). Sharing preferences varied greatly from situation to situation (Friedman test: $\chi^2(4)=81.635$, $p<0.001$). For example, many participants (41.9%, 13/31) indicated that listing attendees would be acceptable, even *preferred*, for team meetings, whereas no participant found it acceptable to provide attendee information for personal events.

F6. Role and relationship influence comfort and usefulness

A large number of questions in our survey focused on how role and relationship impact comfort and perception of the utility of presence technologies. Prior work established a link between role/relationship and sharing preference. We were interested in understanding how role and relationship impact not just what is shared, but how and by whom potential users think the information will be used, particularly to improve workplace communication and collaboration. We concentrated on roles and relationships commonly found in the workplace, such as close colleagues, direct managers, executives, distant colleagues, as well as friends, and members of the user's larger social network. Our role categories were more fine-grained than in previous work. We distinguished between close colleagues and distant colleagues, and between direct managers and higher-level executives, suspecting that we would find differences in sharing preferences to and from these groups.

We first asked users to rate their comfort in sharing a particular type of sensor data with a particular role/relationship. We then asked how comfortable each participant was with that data being shared with a particular role/relationship for a particular use, such as facilitating face-to-face interaction (Figure 5, first row). As with prior work, our results consistently showed that willingness to share presence information was closely related to the roles and relationships between pairs of users. Furthermore, our work shows a significant difference between roles that were lumped together in prior work, such as close colleagues and distant colleagues. We follow Wobbrock et al.'s approach [33] for performing a repeated measure analysis of non-parametric, discrete data. We first used their ARTool to apply an aligned rank transform to our data, after which we apply a standard ANOVA. We found that role had a significant effect on comfort ($F = 23.756$, $p<0.001$), as did sensor type ($F= 31.504$, $p<0.001$). Post hoc tests revealed that comfort sharing with close colleagues was significantly different from comfort sharing with each of the other roles ($p < 0.001$ in all cases). Similarly, comfort sharing with friends and family was significantly different from comfort sharing with each of the other roles ($p < 0.001$). No other role distinctions were significant. Note that comfort sharing with close colleagues ("A colleague working on a project with you") was significantly different from comfort sharing with distant colleagues ("A colleague that you do not work closely with"), two roles that have generally been lumped together in a general "colleague" category in previous work. Post hoc tests on sensor type showed significant differences in comfort between cameras and each of the other sensors ($p<0.001$), and also between calendar information and each of the other sensors ($p<0.005$), but no significant difference between GPS and accelerometer data.

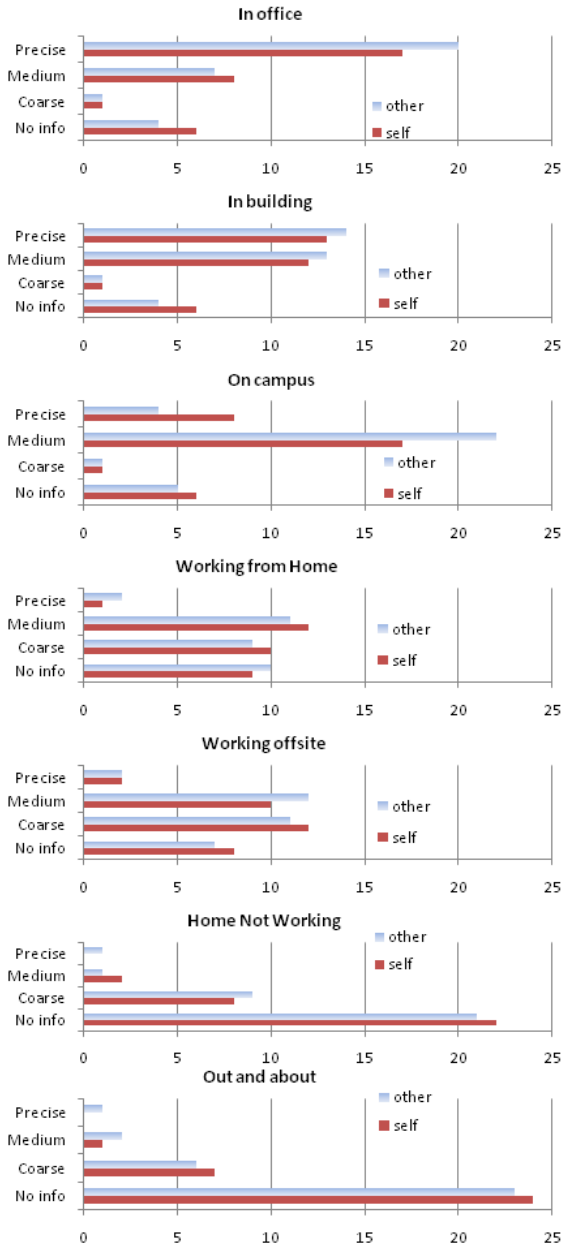

Fig. 3 The red bars show participants responses to the most comfortable level of coarseness for information the system collects in a given situation. The light blue bars show responses to the preferred coarseness for another person in the given situation. The x-axis is the number of responses.

The second row of Figure 5 summarizes responses to perceived usefulness of sharing information for a particular need or use. These results differ markedly from the first row, with most uses and role/relationship combinations having neutral to positive measures of usefulness. Within a particular use, there is variation across the roles and relationships. Perceived usefulness varies significantly with role ($F=13.694$, $p<0.001$). It also varies significantly with task ($F=4.302$, $p=0.007$). The roles broke into two sets, with no significant differences between distant colleague, executive, and broader social network, and no significant differences between close colleague, direct manager, subordinate, support staff, and friends and family. Differences between any two roles from the two different groups were significant at the $p<0.001$ level in all cases. The usefulness of sharing information with people in the second group of roles was recognized by most participants, but most participants did not perceive sharing with distant colleagues, executives, and broader social networks to be as useful. This finding conflicts with most existing presence systems, which are designed to share information broadly, with everyone at work, for example, or with a person's entire social network.

Fig. 4 The red bars show responses to what information a person is comfortable with the system collecting about a given type of event: none; time only; time and place; or time, place and attendees. The light blue bars show responses as to what information people would prefer to receive about another person. The x-axis is the number of responses.

F7. Role and relationship influence fidelity preferences and expectations of use

We also asked participants how they would adjust the fidelity of presence information for a given role/relationship. The responses were consistent with our other findings: friends and close colleagues were generally allowed higher fidelity information, while participants preferred lower fidelity for distant colleagues and members of their larger social networks.

Freeform comments from participants show they were more comfortable sharing with a role/relationship when they had a clear understanding of how such people would benefit from access to the data. One commented that if "*I'm working on a team with a project or paper due [soon] I would be willing to let my team members know my physical location even when I am not at work.*" Another states "*I would restrict sharing presence data in any instance where I felt that one or more parties were meddling or attempting to micro-manage me or my time*." One participant was very specific, "*I would [share] based on dependencies, derived from my view of another person or group's dependency on my involvement and the degree to which I judge that we depend on one another. I would probably develop a scale*."

F8. High concern with respect to storing raw sensor data

While participants were mixed in their comfort with various sensors *observing* them, they were more uniform in their concern about such data being *stored*. The bimodal distributions disappeared, with an overwhelming majority of participants indicating some level of discomfort

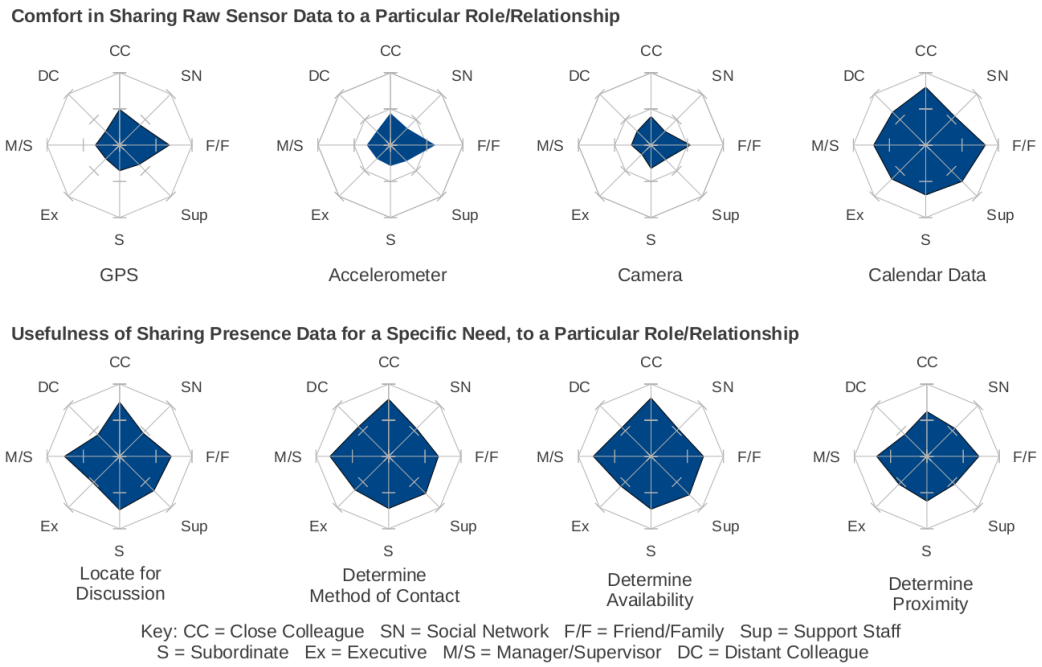

Fig. 5 These graphs shows the averages over the 7 possible responses, with tick marks are at 3.5 and 7. The first row shows the average response with respect to comfort sharing a type of raw sensor data with a particular type of person (1 – very uncomfortable, …, 7 – very comfortable). The second row shows average perceive usefulness levels (1 – clearly not useful, …, 7 – clearly useful) of shared presence data from a particular type of person for a particular purpose.

with data storage for each of the sensor types (Figure 6). With respect to storage of GPS data, 78.1% (25/32) indicated some level of discomfort (1-3, on the 7 point scale). Results were similar for cameras (87.5%, 28/32), accelerometers (75%, 24/32), and keyboard activity sensors (68.8%, 22/32). There was no significant difference between types. When asked about raw sensor data being stored without regard to a particular sensor (general concern), 81.3% (26/32) indicated discomfort.

The rightmost column of Figure 1 summarizes responses to how frequently raw sensor values should be sampled and stored. A majority of respondents were most comfortable with "no storage at all" (56.3%, 59.4%, and 59.4% for GPS, camera, and accelerometer-based sensors, respectively). Storage of keyboard activity sensor information was less of a concern (37.5%).

F9. Storage concerns lessened when storing higher level presence states rather than raw sensor data

To gain insight into participants' concerns regarding storage of computed presence states, rather than raw values, we asked users to rate their comfort in storing states such as "in office" or "at work." These states are calculated from raw sensor data but the raw sensor data are not stored. Participant concern was still high, but significantly less than for general raw data: 68.8% (22/32) expressed concern compared to the 81.3% reported above (Friedman test: $\chi^2(1)=5.762$, $p=0.016$). Far fewer participants (31.3%) indicated "no storage at all" for presence data than for raw data.

F10. Concern with location and ownership of stored data

Participants were concerned not only with what data are stored, but where these data are stored. When asked if they had general concerns about the possibility of their presence data being shared accidentally, 97% (31/32) indicated at least "slight concern," and 50% (16/32) responded "very concerned," the highest level on the survey scale.

Who stores the data had a significant impact on participants' comfort (Friedman test: $\chi^2(5)=105.689$, $p<0.001$). Free Internet service companies, both large and small, were least trusted

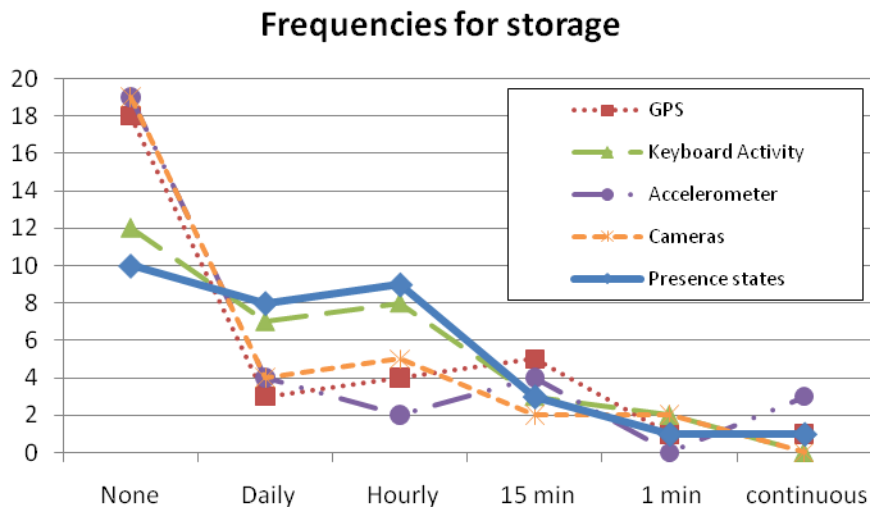

Fig. 6 Number of responses as to the frequency at which participants were comfortable having different types of data stored. Participants were significantly more willing to have the computed presence state shared than the raw sensor values.

by participants to store their data: 90.6% (29/32) indicated discomfort (collapsed ratings of "slightly uncomfortable", "uncomfortable" and "very uncomfortable"). This percentage reduced to 75% (24/32) for private, outsourced storage, 56.3% (18/32) for within-employer storage, and 53.1% (17/32) for workgroup-level storage. Participants overwhelming trusted themselves to store their own data, with only a single person (3.1%) indicating any discomfort.

F11. High concern with who controls access to stored data

Participants expressed concern with respect to which entities decide when and with whom stored presence information is shared. Nearly all respondents indicate discomfort with access decisions being made by free Internet companies (97%, 31/32) and private, outsourced companies (90.3%, 29/32). Decision-making by the employing organization was also of concern, but to fewer people: 81.3% (26/32) for within-organization and 78.1% (25/32) for workgroup-level. Participants again overwhelmingly trusted themselves to control access to their own data, with only one person indicating discomfort: 3.1% (1/32). These differences were significant (Friedman test: $\chi^2(5)=113.675$, $p<0.001$).

F12. Value perceived in storing presence data

In spite of the concerns about storing personal presence information, many participants indicated that stored data does have value. A majority (53%) of users thought access to their own presence data useful (28.1%, 9/32 maybe useful, 18.8%, 6/32 useful, and 6.3%, 2/32 clearly useful). On the other hand, 28.1% (9/32) thought it "clearly" or "likely not useful." Some of these participants contributed strongly worded freeform comments saying that they could see no positive uses for such information. Storing such data is a highly charged issue for many participants.

F13. Aggregation increases comfort

A final focus of our study was to understand how aggregating presence data across groups of users affects utility and comfort. Participants were asked about high-level presence state for a group of individuals, such as an overall presence state for their project team. We sought to understand when such aggregate presence states would be of value and what impact aggregation would have on concerns related to sharing, storing, and frequency of data collection. Survey participants were asked whether they felt having group presence data was generally useful and to indicate which specific groups' presence states they would find useful. They were also asked how their preferences with respect to data collection and storage would change if the data were used only to contribute to a group presence state rather than an individual one.

Overall, participants found aggregate presence information useful. We asked how useful aggregate summaries would be for each of the roles/relationships discussed above. Nearly all respondents (90.6%, 29/32) thought aggregate statistics useful for close colleagues (e.g. having a composited presence state for all colleagues in this category). Fewer participants found value for other roles: subordinates 56.3% (18/32), friends/family 46.9% (15/32), support staff 34.4% (11/32), and larger social group 21.9% (7/32). Only one participant (3.1%) thought aggregate information about distant colleagues useful. These results mirror those discussed above: aggregated data are perceived as useful for people in roles and relationship that participants view as having legitimate reasons for accessing the presence data.

Aggregation impacted participants' comfort with sampling frequency: 56.3% (18/32) were comfortable with more frequent sampling. A majority of participants (59.4%, 19/32) were also willing to have their data shared more frequently. Similarly, many participants (43.8%, 14/32) were comfortable having aggregated data stored for a longer period time.

6 Implications and Lessons

Our results provide a new corpus for reflecting on the design and evolution of presence technologies, and the extent to which they meet users' expectations of privacy and utility. This section discusses design lessons derived from these results. Our results indicate that adoption of these recommendations would lead to higher perceived value of presence technologies, thereby increasing user satisfaction and adoption. These recommendations are a key step toward bringing presence technologies into more widespread use.

6.1 A Sweet Spot in Design Where Privacy and Utility Align

Past work consistently refers to a tension between privacy and utility. Iachello and Hong [13] note that "a concept of tradeoff is implicit in most discussions about privacy." For instance, Chawla *et al.* [8] states that "there are two fundamentally conflicting requirements: privacy for the respondents and utility of the sanitized data." Unlike past work that concentrated on users' privacy concerns about data sharing, our investigation sought to build a deeper understanding through questions that linked comfort with the perceived utility, learning participants' preferences both as a data provider and a data consumer across a broad collection of usage scenarios. Responses to these more detailed questions show that the often-cited tension between privacy and utility is not always present.

In fact, most of our measures showed no significant differences between provider and consumer. For instance, our results suggest that constantly updating a persons' presence state not only decreases comfort, but also is not useful (*F2*). Similarly, sharing highly precise sensor readings is not desired by either data consumers or providers (*F3-5*). While our results hardly varied between data producer and consumer, there were differences based on context. For instance, a person's location greatly impacted the level of precision desired for reporting presence information (*F4*). These observations suggest a fundamental shift in design strategy for presence systems. Viewing privacy and utility as competing forces is too simplistic. A design process that seeks to understand where utility and privacy do (and do not) meet will enable designers to build better systems by leveraging a deeper understanding of how context determines sharing preferences. This section explores how incorporating context enables better design, and ultimately improves a system's utility and comfort.

6.1.1 Primacy of use in designing for presence sharing

Building from the argument above, our results motivate one significant finding for the design of workplace presence systems: comfort in sharing presence information depends on the existence of a plausible, responsible reason to share the information (*F7*). If the data provider sees no positive utility in sharing the data, the data provider exhibits an unwillingness to share. The data provider

need not fear a malicious use in order to be reluctant to share. Nor do data providers need guarantees that the shared data will be used in the way they imagine in order to be willing to share.

Freeform feedback underscores these points. One respondent expressed concern with any access of "*my presence data that cannot be explained by job-related or other explanations.*" Another worried about data "*at risk of being used for unforeseen purposes, by unauthorized people*." The choice of language (e.g. "*unforeseen*" rather than "*harmful*," and "*unauthorized*" as opposed to "*malicious*"), seen consistently in the comments, supports our claim that positive reasons for use are what make users comfortable sharing, rather than assurance against evil uses. This finding suggests that existing presence technologies fail to provide the sharing preferences that users most desire. Few systems provide users with means to control access based on how they expect others to use the data. Providing such controls is a complex design challenge that goes beyond access control and has been largely unaddressed.

6.1.2 Role and relationship must be grounded in models of use

Our results are in line with past work that suggests that role and relationship play a significant role in data sharing practices (*F6,7*). While past work [9,17] implies that role and relationship are fundamental to preference determination, our results suggest that the influence stems from their encapsulation of how presence data will be used (*F6,7*). For example, respondents were much more comfortable sharing awareness information with colleagues they worked closely with than to colleagues with whom they had no working relationship. Participants also perceive clear utility in sharing presence data with the first group, but not with the second. Previous studies had generally considered only a general "colleague" category. The significant difference that we found between close colleagues and more distant colleagues suggests that these two groups must be handled differently in workplace awareness systems. The changeable nature of these relationships implies that the systems must be able to flexibly adapt as co-workers move in and out of these roles. Participants also perceive utility in sharing awareness information with their direct managers, but not with higher-level executives, roles that were often grouped together in previous studies. Our data suggests these two groups must also be handled differently in workplace awareness systems. This finding presents another significant design challenge. Most existing presence systems are designed to share information broadly, with everyone at work or with all colleagues in a given location or department for example.

6.1.3 Support users in refining their understanding of use

Participants expressed concern at unexpected or unintended use of their data by others (*F7,8,10,11*). This concern was *not* focused on a specific misuse, but stemmed from general uneasiness with sharing presence information when a reason for use was not well understood. In freeform comments, users indicated that learning of unnecessary access to their data is one of only a handful of events that would trigger a change in privacy settings. (Others triggers were changes in organizational structure, job, or role.) These concerns suggest that to obtain sustained, effective use of presence systems, designers must enable users to build an accurate model of how their data are being used, and must provide means for users to change their sharing preferences to match their model of use. Violations of their model, or even uncertainty as to how the data are being used, negatively impacts use.

Participants expressed a strong desire for reciprocal awareness, to know when their presence information is accessed. Many participants described this reciprocity as a passive mechanism that would bring to light unexpected or anomalous access. This desire presents yet another design challenge: to design mechanisms that provide suppliers of presence information with appropriate awareness of how that information is being consumed. Prior work has studied this problem in the context of location sharing systems [26], providing means for users to gauge their *exposure*, but does not readily generalize to the richer data sharing models supported by most presence systems.

6.1.4 Good user-level design brings new system-level challenges

Even when there is no conflict between the desires of a provider and a consumer, there may be tension between what the system needs to collect in order to determine what to share and what users think is useful and comfortable to share (*F1-4,8-12*). For instance, to be able to accurately determine a users' state, samples from an in-office camera will likely need to be sampled more frequently than once every fifteen minutes (the most comfortable setting reported by our respondents). Thus, another design challenge is how to enable the system to collect accurate enough information to provide the desired utility without causing discomfort. A variety of emerging technologies provide novel, promising means to meet such challenges. For instance, [23,27] show how to perform privacy-preserving aggregation of time series data, such as presence data.

6.2 Simplified 'Middle Ground' or 'Average' Settings are Poor Designs for Presence Systems

To simplify systems and ease user burden, a common design strategy is to provide a single default setting that accommodates most users. Our results indicate that for many aspects of presence systems, a single setting is a poor design choice. As one example, for each type of sensor feed that we considered, our data exhibits a bimodal distribution, with some users comfortable with data collection by these sensor and others not (*F1*). In these cases, a "single setting" system results in either 1) a large number of uncomfortable users, with many opting out of the system altogether, turning it off frequently, or manipulating the sensor stream to protect their privacy at the expense of system accuracy, or 2) a limited system that would not have access to valuable data that many users are comfortable having collected.

Divorcing sensors from the information they provide improves the design of presence systems. Rather than designing the system to require specific information from a particular in-office sensor, the design should accommodate a wide variety of sensors that could provide some level of in-office information, with varying quality and accuracy. For instance, while camera-based sensing can detect occupancy events with high fidelity, simpler motion sensors can provide similar but less precise information. A presence system that *gracefully* degrades when receiving lower precision sensor feeds while still remaining robust could provide a better balance of privacy and utility. To help users make better sensing choices, they must be made aware of the quality/privacy tradeoffs inherent in their choice of sensor configuration. Users must also be made aware of how other people's sensing preferences impact shared presence states.

Our results exhibit the value of enabling presence systems to dynamically adjust what is being shared based on users' expectations of use (*F4-7,13*). Our results on sharing of calendar information illustrate this need well. Participants' sharing preferences depended strongly on the type of calendar event. This variation suggests value in, for example, presence systems that perform *automatic* detection and classification of face-to-face interactions and incorporate that information into decisions about what to share.

6.3 The Entire Information Pipeline Must Be Trusted

Our results yield several, smaller insights related to participants' concerns as to how presence data are collected, processed, and stored. Past work concentrated on enabling users to specify when and how information is shared *after* it is collected. Our results indicate that control at this level is important and necessary, but *not* sufficient for most users. Participant feedback strongly indicated a need to trust the entire data pipeline, starting with data collection. This section discusses a few special cases in more detail.

6.3.1 User-level control of information sources is desired

When a participant does not want a particular type of information shared (either directly or as part of an aggregate), the user wants to control whether such data are collected in the first place (*F1,8*).

For instance, given that many participants had strong concerns about tracking, each user should be able to prevent a wireless infrastructure from collecting data about his or her location. Similarly, she should be able to keep her office free of cameras.

6.3.2 Ownership and use guarantees are needed

Participants expressed strong concern about third-parties storing their data (*F10,11*). Most users (75.0%, 24/32) did not trust outside organizations with their data, and even more (90.6%, 29/32) did not trust large, global Internet companies such as Google or Yahoo. Central to the concern is fear of misuse and unauthorized sharing. As one user stated, "*My data … must never be made available to entities who might use it in ways that could put me at a disadvantage, such as insurance companies, credit card companies, health insurance companies, marketing companies, anyone who might sell the data.*" While data misuse and unintended sharing are not new concerns, no existing presence systems provide features to explicitly prevent these forms of use, or even alert users when such sharing happens. Formal use and retention policies will need to be designed and embedded into future presence tools, perhaps by leveraging prior results regarding the specification and enforcement of obligation and usage control policies (e.g., [14, 21]).

6.3.3 Aggregation can significantly relax concern

One of our most encouraging results is the impact that processing of data, specifically aggregation, has on measures of comfort and concern (*F9,13*). Nearly all participants indicated discomfort and concern with storing raw sensor data (*F8*). Participants were more positive about storing aggregated presence information derived from their raw sensor data (*F9*). Comfort ratings were even higher for information aggregated over a group of individuals (*F13*). With regard to utility, many respondents felt that access to aggregated data was sufficient, and that there was little or no need for more frequent or more granular data (*F2,3*). Most state-of-the-art presence systems do not provide means for aggregation, but many forms of aggregation can be achieved through enhancements to existing technologies. Care must be taken to ensure that users understand the levels of privacy *actually* afforded by these techniques, as not all aggregation techniques provide equal protection.

7 Conclusions

Presence systems have demonstrated utility for enhancing group awareness and communication in the workplace. Their ultimate success, however, depends on their widespread deployment and candid use. If users are made to feel as though these systems violate their expectations of privacy, this goal can be easily undermined. In this article, we reported on the results of an in-depth survey of user preferences regarding information collection, fusion, storage, and dissemination in presence systems. The implications of our findings are far-reaching, extending some prior results in the literature while challenging others. In particular, the findings of our survey help make the case that (i) privacy and utility are not necessarily situated along a single axis, and balancing these two goals is often possible in practice; (ii) user comfort with various sensing technologies and frequencies tends to follow bimodal distributions, indicating that simplified or "average case" sensing configurations are unsuitable, as they risk alienating multiple user groups; and (iii) the *entire* information pipeline – i.e., sensors used, frequency of sampling, types of fusion and aggregation supported, the means by which data is stored or not stored, and the mechanisms controlling the eventual access of sensed presence states – must be trusted and user-controllable to ensure broad acceptance of presence technologies. Existing presence systems fail to address many of these findings, and could be substantially enhanced by their inclusion.

One specific focus for our future work will be exploring the wide open design space for interfaces that can convey a user's level of *exposure* in an effort to address user concerns regarding acceptable use of their presence data. Preliminary results in this area have been successful for

conveying information about location exposure [26], but do not immediately generalize to the rich data model used by more advanced presence technologies.

The results of this study have already impacted the design of FXPAL's prototype presence system for workplace environments, myUnity [5,32], and will impact it further as more features are added and it is scaled up to support larger numbers of users. MyUnity has been in continuous use by more than 30 participants for over two years. It was designed to support collaboration by increasing workers' awareness of their colleagues' physical presence, activities, and preferred communication channels, and to build group awareness. Like many presence systems, myUnity's sensing and external data inputs are extensive and broad. For instance, the current version supports input from a variety of sensors that include in-office video cameras, motion detectors, wireless localization, and keyboard/mouse activity software monitors. Users run clients on their desktops or mobile devices to display a simple dashboard with color colored tiles that indicate the presence state of all users, and some aggregate presence states across groups of users. Users can click on a tile for more information about a user.

Figure 7: The client dashboard for FXPAL's myUnity workplace presence system, which uses colored tiles to display users' presence states.

One unique feature of myUnity's system design is that each end user can control which sensors collect information about them. Changes can be made at any time, allowing the end user to adjust the system as privacy desires and/or needs change. In this way, my Unity already provided good support for a variety of user preferences regarding presence data sources. MyUnity's presence engine is also designed to adapt as these information sources change, gracefully degrading the presence information to a correct but more general state when less information is known. We believe this feature had a significant positive impact on the system's early adoption and use. Future versions of the system will support a greater variety of sensors that will both support its functionality in a wider variety of environments and implement the lesson of divorcing sensors from the information they provide.

While myUnity initially did not store any data, extensions are slowly being added to myUnity that allow for presence information to be stored securely, and for statistics over this stored data to be computed. In particular, the results of this study with regard to the high level of concern about storing presence data, the lack of trust of third parties, and the benefits of aggregation, inspired new storage and sharing protocols, CollaPSE (Collaboration Presence Sharing Encryption) protocols [23], that allow an individual full access to her own data, enable third party processes to

compute on the data on behalf of the end user without learning anything about the data values, and provides mechanisms for sharing high-level, aggregate statistical information to other users. This feature, along with others, furthers myUnity's support of the critical design lessons related to long term storage and access of presence data.

Currently myUnity is deployed within a small research lab in which everyone knows each other and there is a high degree of trust between its members. Future versions will be deployed on a larger scale. The results of this survey are currently being used to refine and extend the features of the myUnity system, and underpin the design of the sharing framework for a larger deployment in more diverse work environments. A major limitation of this article is that our survey instrument captures users' *perceived* sentiments regarding information collection and dissemination, rather than their *actual* sentiments. Building out and deploying the necessary functionality is critical to understanding the true impact these design lessons have on the use and utility of workplace presence systems. Further, it provides a starting point for the development of future-generation presence technologies that consider user privacy and utility expectations from the outset.

8 References

1. Google Buzz settlement would give Google new privacy rules." The Los Angeles Times: 2 April 2011. http://lat.ms/exOfGn. Accessed 13 April 2011.

2. Acquisti, A., John, L., and Loewenstein, G. (2009) What is privacy worth? In Proceedings of WISE '09.

3. Barkhuus, L. (2004). Privacy in Location-Based Services, Concern vs. Coolness, in Workshop on Location System Privacy and Control, Mobile HCI '04.

4. Begole, J., Tang J.C., Smith R., Yankelovich N. (2002) Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups. CSCW pp.334-343.

5. Biehl, J., et al. (2010). "MyUnity: Building awareness and fostering community in the workplace," FXPAL-TR-09-21 and arXiv:1006.5024

6. Brush, A.J. B., Krumm, J., and Scott, J. (2010) Exploring end user preferences for location obfuscation, location-based services, and the value of location. Ubicomp '10 pp. 95-104.

7. Chansanchai, A. (2011) "Many sharing more on Facebook than they know." NBC Today: 11 April 2011. http://bit.ly/dPYACP. Accessed 13 April 2011.

8. Chawla, S., et al. (2005) Toward Privacy in Public Databases. Theory of Cryptography pp. 363-385.

9. Consolvo, S., et al. (2005), Location disclosure to social relations: why, when, & what people want to share. CHI '05 pp. 81-90.

10. Cvrcek, D., M. Kumpost, V. Matyas, and G. Danezis, (2006) A Study On The Value Of Location Privacy," WPES '06 pp. 109-118.

11. Danezis, G., Lewis, S., and Anderson, R. (2005) How Much is Location Privacy Worth? WEIS 2005,

12. Davis S., Gutwin C. (2005) Using Relationship to Control Disclosure in Awareness Servers. In Proceedings of Graphics Interface 2005 (GI '05), pp.145-152.

13. Iachello G., Hong J. (2007) End-user privacy in human computer interaction. Foundations and Trends in Human–Computer Interaction 1(1):1–137

14. Irwin K., Yu T., Winsborough W. H. (2006) On the Modeling and Analysis of Obligations. CCS'06.

15. Kaasinen, E. (2003) User needs for location-aware mobile services. Personal and Ubiquitous Computing 7: 1 :70-79.

16. Khalil A., Connelly K.(2006) Context-aware Telephony: Privacy Preferences and Sharing Patterns. CSCW'06, pp. 469-478.

17. Lederer, S., Mankoff, J., and Dey, A.K. (2003) Who wants to know what when? privacy preference determinants in ubiquitous computing. CHI '03 pp. 724-725.

18. Lin J., Xiang G., Hong J.I., Sadeh N. (2010) Modeling People's Place Naming Preferences in Location Sharing. Ubicomp'10, pp. 75-84.

19. Patil, S. and Kobsa, A. (2009) Privacy Considerations in Awareness Systems: Designing with Privacy in Mind. Awareness Systems, Human-Computer Interaction Series, Part 2 pp. 187-206,

20. Patil, S. and Kobsa, A. (2010) Enhancing privacy management support in instant messaging. Interacting with Computers 22(3):206-217.

21. Park J., Sandhu R.S. (2004) The UCONABC usage control model. ACM Trans. Inf. Syst. Secur. 7(1): 128-174.

22. Ravichandran, R., Benisch, M., Kelley, P.G., and Sadeh, N.M. (2009) Capturing Social Networking Privacy Preferences. PETS '09 pp. 1-18.

23. Rieffel, E.G., Biehl, J., van Melle, W., Lee, A.J., Secured Histories for Presence Systems. SECOTS 2011.

24. Romero, N., McEwan, G., Greenberg, S. (2007) A Field Study of Community Bar: (Mis)-matches Between Theory and Practice. GROUP'07, pp. 89-98.

25. Sadeh, N., et al. (2009) Understanding and capturing people's privacy policies in a mobile social networking application. Personal and Ubiquitous Computing 13(6):401-412.

26. Schlegel R., Kapadia A., Lee A.J. (2011) "Eyeing your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy," SOUPS'11.

27. Shi, E., et al. (2011) Privacy-Preserving Aggregation of Time-Series Data. NDSS 2011.

28. Szostek, A. M., Karapanos, E., Eggen, B., Holenderski, M. (2008) Understanding the Implications of Social Translucence for Systems Supporting Communication at Work. CSCW'08 pp. 649-658.

29. Tang, K.P., et al. (2010) Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing. Ubicomp '10 pp. 85-94.

30. Toch, E. et al. (2010) Empirical models of privacy in location sharing. Ubicomp '10 pp. 129-138.

31. Wagner, D., et al. (2010) Hide and seek: location sharing practices with social media. MobileHCI '10 pp. 55-58.

32. Wiese J., Biehl J., Turner T., van Melle W., Girgensohn A. (2011) Beyond 'yesterday's tomorrow': Towards the design of awareness technologies for the contemporary worker. MobileHCI'11.

33. Wobbrock J.O., Findlater L., Gergle D., Higgins J.J. (2011) The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only ANOVA Procedures. CHI 2011, pp. 143-146.