

Using Relationship to Control Disclosure in Awareness Servers

Scott Davis and Carl Gutwin

Department of Computer Science, University of Saskatchewan
110 Science Place, Saskatoon, SK, S7N 5C9, Canada
scott.davis, carl.gutwin @usask.ca

Abstract

Awareness servers provide information about a person to help observers determine whether they are available for contact. A tradeoff exists in these systems: more sources of information, and higher fidelity in those sources, can improve people’s decisions, but each increase in information reduces privacy. In this paper, we look at whether the type of relationship between the observer and the person being observed can be used to manage this tradeoff. We conducted a survey that asked people what amount of information from different sources that they would disclose to seven different relationship types. We found that in more than half of the cases, people would give different amounts of information to different relationships. We also found that the only relationship to consistently receive less information was the acquaintance – essentially the person without a strong relationship at all. Our results suggest that awareness servers can be improved by allowing finer-grained control than what is currently available.

Key words: Availability, privacy, awareness servers, availability awareness, relationship-based control .

1 Introduction

Awareness servers are distributed groupware systems that provide information about others – information like presence, location, and activity [1,6,19]. Awareness servers are now commonplace, because they are integrated with instant messengers (IM). Information from these systems is often used to determine another person’s availability for contact: for example, when a person is online, they can be sent a text message.

However, most current awareness servers provide very little information about people. IM systems usually show only a user’s login status and idle time. Even research systems usually use only a single source of information such as video snapshots [4], audio [8], location [18], or calendar status [17]. The limited range of information means that people have more difficulty making judgments about availability, and as a result, may interrupt a person at inappropriate times, or may miss opportunities to interact.

In co-located situations, people use a wide range of information sources to maintain awareness of another person and to determine their availability. For example, people may make decisions based on a combination of

a person’s current activity, location, behaviour, and the state of their local environment (the state of the office door, lights, coat, or desk). Therefore, it seems likely that by increasing the range and detail of what is shown in an awareness server, people will be able to better understand when someone is available, and recent awareness systems are starting to provide this [1,5]. Unfortunately, a tradeoff exists between awareness and privacy: more sources of information and more detail also mean that privacy is reduced, and few people are willing to let detailed information about them be sent out as a broadcast through an availability server.

In this paper, we explore the idea of using relationship to provide finer-grained control over the disclosure of information, thus helping to manage the privacy tradeoff. By giving different amounts of information to different types of people, an awareness server could increase disclosure without compromising privacy. Relationship is already used for this purpose in the real world – it is clear that we disclose different amounts of personal information to different types of people, and that our willingness to let others gather information about us is strongly related to who they are and what their relationship is to us.

To determine whether relationship can be a useful design concept for awareness servers, we conducted a survey that asked people how they would disclose and gather awareness information. We asked people what level of information from six sources they would be willing to disclose, and also what level they would like to see about others, for people in seven different canonical workplace relationships.

The study showed that people do differentiate by relationship for more than half of the information sources. However, people do not differentiate the same way: the only consistent result was that acquaintances received significantly less information than other relationships. The study also showed that people are interested in seeing more information about others than they are willing to disclose, and that there are significant differences between sources: people are less willing to share information about computer screens and telephone calls, and more willing to disclose information about calendar and location. Our findings suggest that future awareness servers should provide people with fine-grained control over disclosure, and that relationship could be used as a way to provide this control.

2 Background

2.1 Informal Awareness and Awareness Servers

Informal awareness of a work community is the general sense of who's around, what they are up to, and whether they are available—the kinds of things that people know when they work together in the same building [4]. Informal awareness helps people recognize and capitalize on opportunities for collaboration.

Traditionally, physical proximity has been necessary for the maintenance of informal awareness [12]. For distributed groups, CSCW researchers have attempted to restore informal awareness with awareness servers that gather, distribute, and display information about the members of a group. The first types of informal awareness systems were media spaces, systems that provide long-term audio and video links between work or social sites. Media spaces, however, are oriented towards places rather than people; other awareness systems have focused on information about individuals.

Several research systems have been built that provide a range of different information sources at a variety of different levels of detail. For example:

- Portholes [4] provides periodic video snapshots of individual offices, but without audio;
- Thunderwire [8] provides a high-quality always-on audio link, but without any visual information;
- Active Badges [18] shows people's locations in a building at the room level;
- Augur [17] allows people to see others' online calendar entries;
- Peepholes [6] (and most current IM systems) shows who is logged in to their computer, and whether they are active or not.

The information provided by even minimalist systems helps people track who is around and who is available. However, all of the examples above use only a single information source, despite the fact that people use a combination of information in real-world situations. More recent systems do use multiple information sources: for example, Awarenex [16] shows login status, general location, and online calendar; Lilsys [1] indicates motion, sound level, telephone use, door status, login status, and current calendar entry; and MyVines [5] uses audio, location, computer activity, and calendar information.

One main distinction in these awareness systems is the difference between simply presenting the state information and allowing the observer to interpret it [4,6,8,16], and inferring certain things automatically from the state data [5,17]. In particular, these systems attempt to infer whether or not a person is available for contact.

2.2 Availability and Interruption

People use awareness information for a variety of reasons, but one in particular is to help them determine when a person can be contacted. Since people may be busy, the decision of whether to interrupt someone can be a difficult one. There are many social factors governing people's willingness to interrupt others, but it is clear that we are sensitive to a person's level of engagement [7], and that it is rude to interrupt at a time when the recipient feels that they are not available. Despite this social pressure to avoid unwanted interruption, the interrupter usually gains more benefit than does the person being interrupted; this cost imbalance suggests that interruptions will occur despite social pressure to avoid them [3].

The amount of information that an observer has about a person appears to be a factor in interruption decisions. Begole and colleagues [1] observe that most current awareness servers show a person's presence, and even though presence is not the same as availability, the absence of additional information leads most people to equate the two. Dabbish and Kraut [3] showed that when observers had to interrupt a person playing an interactive game, people who had detailed information about the state of the game were more likely to wait before interrupting, and interrupted less frequently than those who had no information. People's increased ability to assess availability also had negative consequences, however, in that observers spent more time waiting for an appropriate moment to interrupt, and as a result got less work done themselves.

The type of information is also important to judgments of a person's availability for contact. In Dabbish and Kraut's study, information about activity on the computer screen was needed to determine availability. In other situations, social information is more important: for example, Hudson and colleagues [10] showed in a study of a professional assistant that audio was an important information source. If the assistant was talking with someone else, either in person or on the phone, then they were unavailable.

The difficulty of determining exactly what information will be needed to make a good decision about interrupting shows that including several information sources can be useful. However, the problem with providing more information is that it reduces privacy.

2.3 Privacy and awareness

Privacy has been investigated in several contexts in CSCW, and there are several different definitions for the concept. Despite this variation, two main aspects of privacy are commonly identified that are relevant to this work: confidentiality and solitude [2,15]. These are

‘control processes’ in which a person manages the flow of information towards and away from the self [2].

Confidentiality is control over information moving outward from the self, and affects other people’s access to personal information. One aspect of confidentiality that is particularly relevant to our work is control over information fidelity – the accuracy and specificity of outgoing information.

Solitude is control over information moving towards the self (including interruptions), and determines how much of one’s attention is taken by that information.

There are tradeoffs between awareness and both confidentiality and distraction: as the observer gets more information, the source has less privacy; and, as more information arrives about other people, there is a greater potential for distraction from other tasks [11]. Some systems attempt to resolve the first tradeoff by modifying the awareness information to provide only some of the original information – for example, by blurring a video image to the point where presence and identity can be detected, but not the specifics of activity [11]. This can work in some cases, but not in all, particularly when in a sensitive environment [14].

Boyle and Greenberg [2] suggest that because awareness and privacy cannot both be satisfied by these simple manipulations, better control over information is required. However, they also recognize that user effort is a crucial factor in the success of the control process: “There are few fine-grained yet lightweight strategies for controlling video media spaces... Heavyweight and coarse-grained privacy control interfaces prompt an ‘all or nothing’ trade-off [and] as a result, people often do not make changes when appropriate, and often end up configuring the system to grant all others either full access at any time, or no access whatsoever” ([2], p. 6).

Our goal is to see whether relationship can be used as a fine-grained, lightweight control – a way to increase the specificity of privacy control without the effort of specifying a preference for each individual. There is some evidence to suggest that relationship could be useful in this regard; for example, identity has been shown to be a more stable indicator of privacy preferences than situation [13].

3 Study Methods

The study was a survey of people’s preferences in disclosing information to, and gathering information from, different types of canonical relationships. Sixteen participants (six women and ten men) were recruited from a variety of work and organizational settings (software development, engineering, management, sales, teaching, and research). Participants’ ages ranged from 25 to 42 years, all participants were regular users of computer systems for their work, and all were commonly in situa-

tions where they had to determine the availability of someone in their organization.

A questionnaire (Figure 1) was given to each participant, asking them about what information they were willing to disclose, and what information they would choose to use, in a number of different situations. The questionnaire was organized by four factors:

- Information source: video, audio, telephone, location, calendar, and computer use;
- Information fidelity: four different fidelity levels, from a binary indication to the richest available for that information source;
- Relationship to observer: peer, supervisor, subordinate, secretary, friend, spouse, or acquaintance (i.e., any employee not in the other six groups);
- Direction: whether the awareness information is being disclosed (outgoing) or gathered (incoming).

	Friend	Supervisor	Peer	Subordinate	Spouse	Secretary	Acquaintance
Type 1. Video							
Live video (20 frames/sec.)							
Snapshot image (2/min.)							
Fuzzy snapshot (1/min.)							
Motion detector							
Type 2. Audio							
High quality sound link							
Muffled sound link							
Sound level							
Sound detector							
Type 3. Location							
Moving dot on a map							
Room and location							
Type of room							
In/out of building							
Type 4. Telephone							
Sound link to phone call							
Call in progress & names							
Call in progress & type of call							
Call in progress only							
Type 5. Calendar							
Full calendar for today							
Booked/free & current entry							
Current calendar entry							
Currently booked or free							
Type 6. Computer activity							
Screen snapshot							
Application list & idle time							
Idle time only							
Active/inactive only							

Figure 1. Questionnaire questions (text labels have been shortened to reduce space).

Participants were first asked to think of a particular person in each relationship category. They were then asked to state, for each cell in the table, whether they would disclose the information to that person, and whether they would use the information when trying to determine that person’s availability. People were asked to imagine themselves in two scenarios, one for considering disclosure and one for considering what they would choose to gather. These scenarios were chosen to simulate an average, everyday setting, rather than extreme cases (either of being extraordinarily busy, or of being desperate to contact someone).

From the questionnaires, we extracted the maximum fidelity of information that would be given to each relationship type. This data was used in the analysis below.

4 Results

Results and analysis from the questionnaire are organized by the following research questions:

- Do people differentiate the information they are willing to disclose (or gather) by relationship?
- How much do people differentiate?
- Are particular relationships consistently given more or less information?
- Do people differentiate by information source?
- Are there differences between fidelity levels?

4.1 Do people differentiate by relationship?

For each person, we calculated the maximum fidelity they would give to (or use from) each relationship, for each information source. The number of differences within each information source is a simple indication of differentiation, and indicates whether a person would use different settings for different relationships in an awareness server.

In addition, we analysed the data using two separate groupings. First, we tested all relationships (including the ‘acquaintance’ category), which corresponds to an availability server where everyone can see each other (e.g., an internal company directory). Second, we also considered the data without the ‘acquaintance’ category, which corresponds more closely to current IM systems, where people specifically choose which people to include in their system.

Disclosure. In our data, people differentiated their disclosure more than half the time (Figure 2). When considering all relationships (i.e., including acquaintances), people would disclose different amounts to different relationships for 61% of the information sources. When considering only personal and professional relationships (i.e., friend, spouse, supervisor, peer, subordinate, secretary), people differentiated with 54% of information sources.

Gathering. People differentiated much less in their preferences for gathering information. As shown in Figure 2, people only indicated a relationship-based difference in about 25% of cases. Since there was so little differentiation, our analyses in the next section focus on disclosure only.

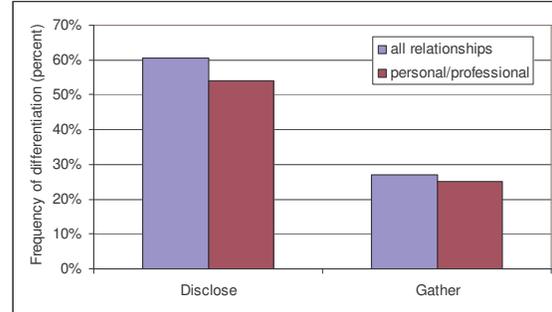


Figure 2. Histogram showing frequency of differentiation, for both disclosing and gathering.

4.2 How much do people differentiate?

We considered both the amount and the magnitude of differences. To measure the amount of differentiation, we counted the number of differences from the majority answer for each information level. For example, if a person would disclose level one video information (and only level one) to everyone, there is no differentiation; if they would give a higher level to their spouse only, then there is a difference of one from the majority.

Figure 3 shows amount of differentiation for disclosing information, and again considers two groupings of the data (all relationships, and professional/personal relationships only). Although the largest category has no differentiation, a large proportion of cases (more than 40% when considering all relationships) show more than two differences per information source.

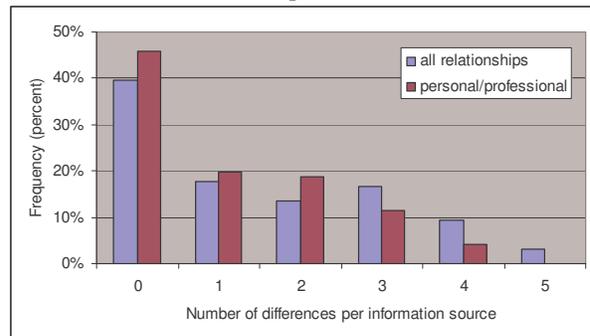


Figure 3. Histogram showing amount of differentiation per information source, in terms of the number of differences from the majority.

Finally, to measure the magnitude of the differences, we recorded the difference between the highest and lowest fidelity levels for each information source. Fig-

ure 4 shows that in several cases, there were fairly large differences: for example, in 30% of cases (considering all relationships) there was a difference of three or four between the highest and lowest fidelity levels. In these cases, people are giving some relationships the highest level of fidelity, and others either the lowest level or none at all.

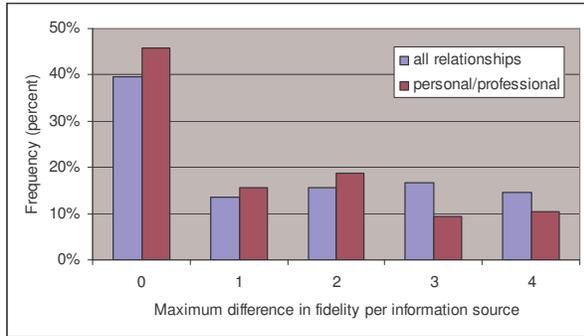


Figure 4. Histogram showing magnitude of difference between highest and lowest fidelity levels.

4.3 Are there differences between relationships?

Using the measure of maximum fidelity, we looked at whether people treated certain relationships differently.

Disclosure. For disclosing information, there was a significant effect of relationship on maximum fidelity ($F_{6,10}=5.26, p<0.001$). A post-hoc Tukey test was carried out to look for differences between relationships. There was only one significant difference found: participants were willing to disclose significantly less fidelity to the ‘acquaintance’ relationship than to any the others (all $p<0.01$). No other differences were found between pairs of relationships (see Figure 5).

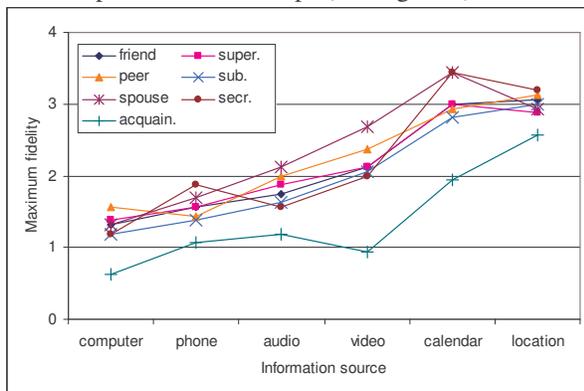


Figure 5. Mean maximum fidelity participants would disclose for each relationship type, by source.

Gathering. A similar analysis of variance was carried out for data from the gathering scenario. As can be seen from Figure 6, the relationships were treated similarly, and no significant effect of relationship was found ($F_{6,10}=0.48, p<0.82$).

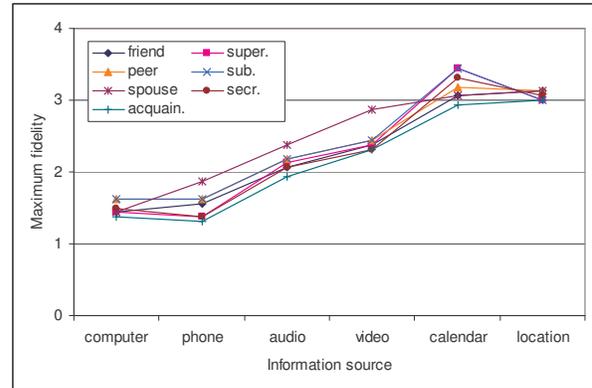


Figure 6. Mean maximum fidelity that participants would gather for each relationship type, by source.

4.4 Are there differences between sources?

We also used maximum fidelity level to look at whether people had different preferences depending upon the source. It should be noted, however, that fidelity is an inexact measure when comparing sources, since scales are relative to the source. For example, level two fidelity for video is not the same amount of information as level two for a calendar. Nevertheless, the scales all range from a binary representation to the maximum amount of information possible for that source, and therefore, we believe that the measure can be still used to indicate broad differences.

Disclosure. Figure 7 shows mean maximum fidelity for each information source. Analysis of variance showed that there was a main effect of information source ($F_{5,11}=15.42, p<0.001$).

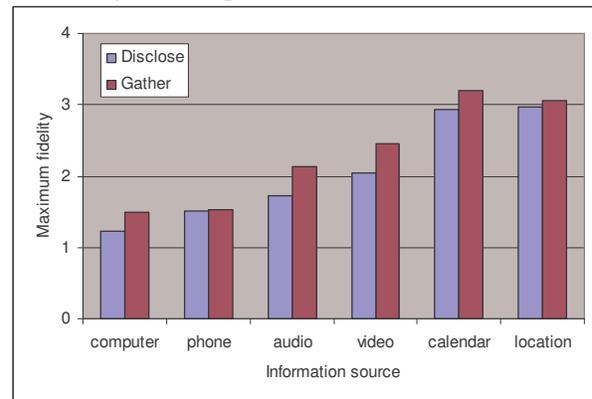


Figure 7. mean max fidelity for disclosure and gathering, by information source.

A post-hoc Tukey test showed several significant differences between individual sources (see Table 1). This test suggested that the sources can be divided into three main groups: people are willing to disclose the least fidelity for computer and telephone, significantly

more for audio and video, and significantly more again for calendar and location.

Gathering. There were also clear differences between sources in the data from the gathering-information scenario (see Figure 7). There was again a main effect of source on maximum fidelity level ($F_{5,11}=17.23$, $p<0.001$). Follow-up tests divided the sources into similar groups as those discussed above.

	computer	phone	audio	video	calendar	location
computer		-2.52	-4.43	-7.12	-14.86	-15.12
phone	n.s.		-1.91	-4.60	-12.34	-12.60
audio	$p<0.05$	n.s.		-2.69	-10.43	-10.69
video	$p<0.01$	$p<0.05$	n.s.		-7.73	-7.99
calendar	$p<0.01$	$p<0.01$	$p<0.01$	$p<0.01$		-0.26
location	$p<0.01$	$p<0.01$	$p<0.01$	$p<0.01$	n.s.	

Table 1. Tukey pairwise comparisons (disclosure data). Statistic value above the diagonal, significance below.

4.5 Differentiation by fidelity

As expected, people were less willing to disclose higher-fidelity information (see Figures 8 and 9). However, people were also unwilling to gather high-fidelity information; people commented that they would see this as an intrusion on their part, or that they did not want to see or hear things that were not public. Significant main effects of fidelity level were found for both disclosure ($F_{3,13}=40.03$, $p<0.001$) and gathering ($F_{3,13}=17.78$, $p<0.001$). However, in both cases, level one and two information was treated almost equally.

5 Discussion

The study provides four main results:

- people differentiate information disclosure by relationship, for more than half of the information sources;
- the only difference within relationships is that people would reveal less to acquaintances than to personal and professional relationships;
- people were willing to reveal higher fidelity information about their locations and calendars than they were about their computer use, telephone conversations, or office audio;
- most people were willing to disclose lower levels of fidelity, and few were willing to disclose the highest level (see Figures 8 and 9).

In the following paragraphs, we consider explanations for these findings, discuss a particular problem that is raised by relationship-based control, and make recommendations for the design of awareness servers.

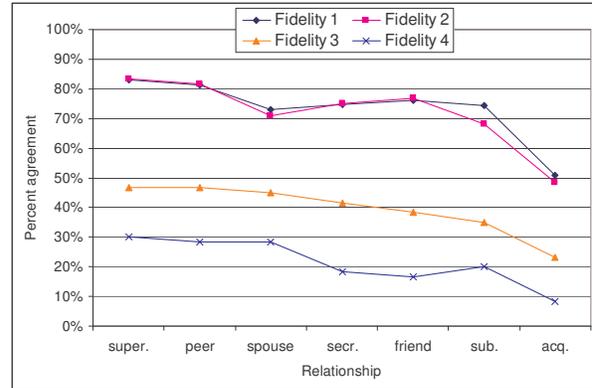


Figure 8. Disclosure: percent of participants agreeing to each fidelity level for disclosing information.

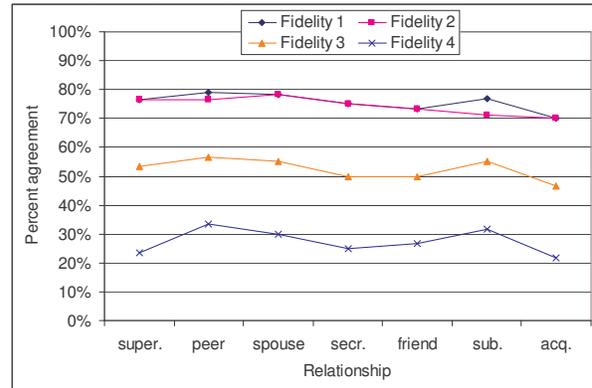


Figure 9. Gathering: percent of participants agreeing to each fidelity level for gathering information.

5.1 Explanations of results

We expected people to differentiate by relationship, but it was at first surprising that only the ‘acquaintance’ relationship was consistently different. However, the similarity in personal and professional relationships can be explained by simple variation within the participants: each of them has had different experiences with relationship types, and so different people would be likely to treat them differently.

We also expected differences between fidelity levels, but were interested to see the similarity between levels one and two. Comments from participants suggested that level one information would not be very useful, although there are certainly cases where binary values would help people determine availability.

Finally, the differences in disclosure between sources appear to reflect the degree to which the information is already public. People were willing to disclose less information about computer and telephone use, activities which are usually private (it is seen as rude to look at a person’s computer or listen in on a phone call without invitation). People were willing to give out more information about location and calendar:

location is already public information (at least for people who are nearby); calendars are less so, but electronic versions are often made public within the organization for purposes such as scheduling meetings. Several participants mentioned that they had no problem disclosing information about their calendar because it was already available in their organization.

5.2 A dilemma in relationship-based control

Our study suggests that there is a correspondence between the amount of information given to particular relationships, and a person's availability to that relationship. For example, some participants gave higher-fidelity information to a spouse or a secretary, and when questioned, stated that these were also the people who had the most rights to contact them. Similarly, mere acquaintances were given less information, and these were also the people from whom interruptions would be the least welcome.

However, tying fidelity to availability raises a problem, because *less* information can sometimes lead to *more* opportunity to interrupt. The reason is that when people have more information, they are better able to determine availability, and so are less likely to interrupt inappropriately; with less information, there is little basis for deciding that a person is too busy to contact, and so interruptions are more often timed badly.

This is illustrated by Dabbish and Kraut [3], who gave observers different levels of information and studied their interruption patterns. People who saw more information interrupted less often, because they were better able to determine that the person was busy. The problem is that less information gives the observer the 'benefit of ignorance:' since they do not have enough information to determine availability, they cannot be blamed for interrupting at a bad time. This benefit can be seen in a story related by one participant: "I saw that he was with someone, so I went back to my office and phoned him, because then it wouldn't seem rude." With more information, people have both the ability to make good decisions, and the social pressure to do so.

The dilemma shows that sometimes privacy can be *increased* by disclosing more information [2]: that is, disclosing more information reduces confidentiality, but increases solitude by enabling people to make better decisions about interruption.

However, in the case of an awareness server, the problem cannot be solved by giving more information to everyone, since people are unwilling to do this. We believe the only solution is to also allow fine-grained control over solitude. This implies that in addition to giving different amounts of information to different relationships, different groups would also have differential access – that is, more or less ability to interrupt.

Differential access already occurs when a person has a secretary (and only certain types of people are allowed in), but is not as common in other communication means. The telephone, for example, provides equal access regardless of relationship; similarly, anyone on an IM contact list can pop up a window on their target's computer screen. Relationship-based control over access could mean, for example, that only one's spouse could make the telephone ring, or that mere acquaintances would not be able to open chat windows on the screen. Interface design for differential access is still in the early stages [9], but we note that it is already happening in mobile phones, where custom ring tones can be attached to particular incoming telephone numbers. Along with the ability to organize contacts into groups, this allows a person to give different groups different rings that are more or less demanding of their attention.

5.3 Lessons for designers of availability servers

Several lessons can be taken from this research. First, the study suggests that awareness servers should provide a means for differentiating information disclosure. Relationship appears to be an attribute that people use in this process, although further study is needed to determine whether other attributes (or customizable groups) could be more appropriate in some situations. Some IM systems already allow limited differentiation (e.g., by allowing a person to 'appear invisible' to certain contacts) but we would argue that a richer range of differentiation is warranted.

Second, awareness systems should make it possible to disclose at least the amount of information represented by level two of our four fidelity levels. This level was acceptable to a large majority of our participants, and the additional information could help people to make better decisions about availability. Since level two and level one were treated so similarly, it is possible that the lowest level of fidelity may not even be needed in awareness servers; however, since the lowest level lends itself well to on/off type displays that fit into a small space, it may be necessary to retain it regardless of whether anyone will actually choose it.

Third, interfaces for controlling disclosure should be different from those for controlling information-gathering. The ability to differentiate information quickly and easily should be a fundamental part of the disclosure interface; but since differentiation was much reduced for gathering, a single setting could suffice in this control. However, the information-gathering interface will have to help people limit the potential for distraction. One way this could be achieved, for example, is to show only the lower levels of fidelity in the interface until the user indicates that they want more detail.

Fourth, awareness servers should allow differentiation by information source. As a default, these systems should provide less information about computer and telephone, and more about location and calendar. Again, however, the ability to change these defaults should be provided.

6 Conclusion

Awareness servers often provide too little information for people to accurately determine availability. More sources of information, and higher fidelity in those sources, can help observers to make better decisions about when to interrupt. However, more information means a reduction in privacy, and for computer-mediated awareness systems to be effective, better control over confidentiality is needed. We carried out a survey to investigate relationship as a mechanism for finer-grained control over the disclosure of awareness information. We were interested primarily in whether people would differentiate the information they disclosed or used by relationship type. We found that they do so more than half the time when disclosing information, but only in about one-quarter of cases when gathering information about others. We also found that people treat different sources differently, and that only acquaintances are consistently given less information.

In future, there are several further investigations that could be carried out in this area. Our first goal is to increase the size and scope of the survey, to see whether our initial findings will hold in a larger group, and whether other grouping attributes, the ability to create custom groups, or simple individual control will be preferred to relationship-based groups. In addition, we plan to investigate interactions between relationship, task situation, and urgency, as has been suggested in prior work [13]. Second, we plan to build a working prototype that puts our findings into practice, allowing people to differentiate disclosure by relationship. We are particularly interested in designs that will allow people to see different levels of information, but without causing undue distraction [11]. We are also interested to see whether people's preferences change as they gain more experience with the system. We also plan to add differential access control to the system, to test the idea that control over solitude can help to solve the dilemma of confidentiality.

References

- [1] Begole, J., Matsakis, N., and Tang, J., Lilsys: Inferring unavailability using sensors, *Proc. ACM CSCW 2004*, 511 – 514.
- [2] Boyle, M. and Greenberg, S., The Language of Privacy: Learning from video media space analysis and design. *ACM ToCHI*, in press.
- [3] Dabbish, L., and Kraut, R., Controlling Interruptions: Awareness displays and social motivation for coordination, *Proc. ACM CSCW 2004*, 182 – 191.
- [4] Dourish, P. and Bly, S., Portholes: Supporting awareness in a distributed work group, *Proc. ACM CHI 1992*, 541-547.
- [5] Fogarty, J., Lai, J., and Christensen, J., Presence versus Availability: The design and evaluation of a context-aware communication client. *IJHCS*, 61(3), 2004, 299-317.
- [6] Greenberg, S., Peepholes: Low cost awareness of one's community, *Proc. ACM CHI 1996*, 206-207.
- [7] Heath, C., Jirotko, M., Luff, P., and Hindmarsh, J., Unpacking Collaboration: The interactional organization of trading in a city dealing room. *JCSCW*, 3(2), 1995, 147-165.
- [8] Hindus, D., Ackerman, M., Mainwaring, S., Starr, B., Thunderwire: A field study of an audio-only media space, *Proc. ACM CSCW 1996*, 238-247.
- [9] Horvitz, E., Jacobs, A., and Hovel, D., Attention-Sensitive Alerting, *Proc. Conf. on Uncertainty and Artificial Intelligence (UAI) 1999*, 305-313.
- [10] Hudson, S., Fogarty, J., Atkeson, C., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J., and Yang, J., Predicting Human Interruptability with Sensors: A Wizard of Oz feasibility study, *Proc. ACM CHI 2003*, 257-264.
- [11] Hudson, S. and Smith, I., Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support, *Proc ACM CHI 1996*, 248-257.
- [12] Kraut, R., Egido, C., Galegher, J., Patterns of Contact and Communication in Scientific Research Collaboration, *Proc. ACM CSCW 1988*, 1-12.
- [13] Lederer, S., Mankoff, J., and Dey, A., Who Wants to Know What When? Privacy preference determinants in ubiquitous computing, *Proc. ACM CHI 2003*, 724-725.
- [14] Neustaedter, C., Greenberg, S., and Boyle, M., Blur Filtration Fails to Preserve Privacy for Home-Based Video Conferencing, *ACM ToCHI*, in press.
- [15] Palen, L., and Dourish, P., Unpacking Privacy for a Networked World. *Proc. ACM CHI 2003*, 29–137.
- [16] Tang, J., Yankelovich, N., Begole, J., Van Kleek, M., Li, F., and Bhalodia, J. ConNexus to Awareness: Extending awareness to mobile users, *Proc. ACM CHI 2001*, 221-228.
- [17] Tullio, J., Goecks, J., Mynatt, E., Nguyen, D., Augmenting Shared Personal Calendars, *Proc. ACM UIST 2002*, 11-20.
- [18] Want, R., Hopper, A., Falcao, V., and Gibbons, J., The Active Badge Location System, *ACM ToCHI*, v.10 no.1, 1992, 91-102.