Privacy in Collaboration: Managing Impression

Sameer Patil, Alfred Kobsa

Department of Informatics, Donald Bren School of Information and Computer Sciences
University of California, Irvine
{patil, kobsa}@uci.edu

Abstract

The great promise of collaborative technologies that improve group awareness and communication is often overshadowed by accompanying privacy concerns. In such systems, the privacy concerns relate to the individuals one interacts with – colleagues, superiors, subordinates, friends and family. Using Instant Messaging as an example, we illustrate that addressing privacy issues in such systems can be quite challenging. Based on in-depth interviews with experienced IM users, we propose that the primary concern regarding privacy in collaborative awareness systems is an individual's desire to control how one appears to others. We advocate that designers empower users to appropriately monitor and manage the impression they project towards others by providing modifiable policies and settings, with suitable defaults and seamless interaction. We also provide a few design suggestions to seed further exploration.

1 Introduction

Timely information about current activities, work progress and availability of one's team members is very valuable for fostering informal communication within a work team, and for low-level synchronization of work activities (Dourish & Bellotti, 1992; Herbsleb, Mockus, Finholt, & Grinter, 2000). This is especially critical when team members are geographically distributed. Yet, people usually loath disclosure and surveillance of their whereabouts and activities (Want, Hopper, Falc_o & Gibbons, 1992). This results in tension between the need for awareness and the desire for privacy.

In systems devised for communication and collaboration, the privacy concerns in question are primarily with respect to other individuals one interacts with – colleagues, superiors, subordinates, friends and family – as opposed to big, nameless entities such as corporations and the government.

We start by describing what awareness and privacy mean in the context of collaborative work. We use Instant Messaging (IM) systems as an example to illustrate how even a conceptually simple collaborative system can be quite complex when analyzed from the point of view of privacy. Based on interviews with users of IM systems, we propose that users' primary concerns regarding privacy in collaborative work settings arise from their desire to manage their impression. Finally, we provide some suggestions for privacy-sensitive system design.

2 Awareness

Awareness of the activities of collaborators helps individuals plan, orient and coordinate their own work to fit the activities of their team, department or organization. This results in increased efficiency and effectiveness of individual work as well as the work that is carried out collaboratively (Dourish & Bellotti, 1992). It is no surprise then that more tightly-coupled the collaborative activity, higher the amount of effort and time individuals spend in seeking information about the availability and activities of others and in providing information to others of their own availability and activities (Herbsleb, Mockus, Finholt & Grinter, 2001).

Awareness information is multi-faceted. It includes information about people's presence, activities (past, present, and future), schedules, routines, deadlines, and availability. Moreover, such information may be provided and received through a variety of channels – from physical to social to electronic. For instance, by peeking through a partially open office door one may find out whether a colleague is busy. One may also use one's knowledge of a colleague's routine to infer his or her availability, or one can consult his or her online calendar.

Over the years a variety of software systems and hardware devices have been built with the explicit goal of supporting the collection and dissemination of awareness information. Examples of such systems include Shared Media Spaces (RAVE (Bellotti & Dourish, 1997), Portholes (Dourish & Bly, 1992), Thunderwire (Hindus, Ackerman, Mainwaring & Starr, 1996)), Shared Calendars, Mailing lists, Shared Workspaces (Polyteam (Mark, Fuchs, & Sohlenkamp, 1997), BSCW - http://bscw.gmd.de/), Shared repositories (Docushare - http://docushare.xerox.com, CVS - http://www.cvshome.org), Newsgroups, Instant Messaging (e.g. ICQ - http://www.icq.com), Sensors (e.g. Active Badges (Want et. al., 1992)), and Shared Displays (e.g. Notification Collage (Greenberg & Rounding, 2001)). Even systems generally regarded to be single-user, such as email and telephone may be employed for awareness purposes. For example, caller ID may be used to screen calls; automatic email replies may be used to indicate extended unavailability etc.

We find people typically using a combination of diverse systems and mechanisms in their efforts to generate, disseminate and receive awareness information. The manner in which various mechanisms are combined and used depends on the people and the task(s) involved, the granularity of awareness information, the frequency of changes in awareness information, resources, cultural norms, context and so on.

Awareness information assumes a more important role in the context of close collaboration, and even more so if the collaborators are geographically distant (Herbsleb et. al., 2000). Since our primary interest is supporting collaborative work of globally distributed teams, we focus here on awareness systems and on mechanisms encountered in this scenario.

3 Privacy

Privacy is currently a highly publicized and hotly debated topic. Yet, due to the complexities involved, there exists no commonly agreed-upon definition of 'privacy'. One source of difficulty in defining privacy is its highly situated, context-dependent nature (Palen & Dourish, 2003). Even in the same situation, different individuals involved may have different opinions and expectations of what privacy means. This fuzziness, context dependency and individual variability make dealing with privacy a rather difficult task.

Bellotti (Bellotti, 1996) has pointed out that two types of privacy definitions are common. She refers to them as normative and operational. Normatively, Warren and Brandeis (Warren & Brandeis, 1890) define privacy as "freedom to be left alone". Stone et. al. (Stone, Gardner, Gueutal & McClure, 1983) offer an operational definition of privacy as "ability of the individual to personally control information about oneself" whereas Samarajiva (Samarajiva, 1997) extends the definition to "the control of outflow of information that may be of strategic or aesthetic value to the person and control of inflow of information including initiation of contact".

In the physical domain, a variety of mechanisms and artifacts seem to have evolved over time to make privacy management easier. These embody certain social protocols based on shared assumptions, such as locking the door to prevent access to others, or knocking before entering even when the door is partially open. However, when the shared assumptions behind the embodied social protocols are no longer applicable, for whatever reason – individual, cultural, contextual, task-specific – privacy management once again becomes problematic and privacy violations may occur.

It is commonly observed that the consequences and risks involved determine the amount of (explicit) effort and time devoted to managing privacy. When the consequences are potentially severe, people devote considerable attention to preserving privacy. If, despite their efforts, a violation of privacy does occur, individuals typically negotiate until a commonly agreed upon state of privacy is reached for everyone involved.

4 Relationship between Awareness and Privacy

The above discussion regarding awareness and privacy makes the inherent interrelation between the two apparent. The general perception is that there is an inverse relationship between privacy and awareness: more awareness leads to less privacy and vice versa. Even though this may typically be the case, the reverse may also be true, i.e. providing more awareness provides more privacy. For example, maintaining a personal web page with electronic copies of their papers allows faculty members to limit the intrusion caused by requests for copies of their publications (Palen & Dourish, 2003).

Given the highly situated and context dependent nature of both awareness and privacy, it should be no surprise that the exact manner in which awareness and privacy depend on each other is also context dependent. However, regardless of the exact relationship between the two at any given time, it is certainly true that they influence each other greatly. The question, then, is how people manage the relationship between awareness and privacy – both in the physical and the electronic world.

Effectively dealing with privacy management in electronic domain is rather challenging. Part of the reason is that privacy runs into the social-technical gap referred to by Ackerman (Ackerman, 2000). Electronic systems frequently embody, or try to mimic, artifacts and concepts from the physical and social domains. However, the underlying assumptions of privacy may be partially or totally lost in the transformation from the physical or social world, to the digital. A break in expectations means either too much or too little privacy, compared to what is desired or expected.

The other part has to do with specific affordances of the domain itself. In the electronic domain it is much easier to mine data, and to combine distinct pieces of information in a way that the compiled information is of greater value than the sum of its parts. Additionally, digital information may be easily archived extending its temporal dimension indefinitely. Finally, digital information can be easily transmitted across distance, making its reach global. This ease of information mining, archiving, and sharing has far-reaching consequences for privacy management. As a result, it has heavily contributed to the widespread concerns regarding privacy in the electronic domain.

We illustrate our points in the next two sections by discussing how current awareness systems deal with privacy, and by presenting a comparison of the privacy management aspects of various IM systems.

5 Privacy Mechanisms in Current Collaborative Awareness Systems

As a start, it is instructive to study privacy mechanisms in current collaborative awareness systems. Designers and builders of such systems frequently tend to treat privacy either as a secondary consideration or as an issue for future exploration. This may be due to the underlying assumption that individuals who collaborate with each other have less stringent privacy expectations. As we will discuss below, the result is often systems with privacy mechanisms that do not match users' needs and expectations.

Current awareness systems provide for privacy management through a combination of a large number of mechanisms. The essence of these mechanisms seems to revolve around controlling access (to oneself and one's artifacts) through proper authorization. Different mechanisms differ in terms of who has control, who is authorized, and how the process of authorization works. Some examples of privacy mechanisms include access control (e.g., password-protected login), assignment of groups and roles, and summary and distortion (e.g. abstracting a document, blurring of a video stream (Boyle, Edwards & Greenberg, 2000)). These mechanisms may be enacted and enforced in a variety of ways including provision of defaults, generation of feedback, enforcing reciprocity, policies and procedures, as well as support for social consensus.

In reality, control and authorization decisions are hard to make upfront since they change dynamically with context. Incorporating this context dependence into the capabilities provided by present systems is difficult, to say the least. The adequacy of these mechanisms for privacy management and the manners in which they are utilized in current awareness systems need to be studied in detail. If we know what mechanisms people use in an effective manner, we can look into the reasons, and then use the findings to inform the design of awareness systems with improved privacy management solutions.

6 Privacy and Awareness in IM Systems

One of the most popular and widespread contemporary awareness mechanisms is Instant Messaging (IM). Although IM emerged in a non-work context, it is increasingly being used and studied in the context of supporting collaborative work (Herbsleb, Atkins, Boyer, Handel & Finholt, 2002). IM allows people to indicate their presence to others who are on their list of contacts. At the same time, it allows checking for presence of contacts. It is possible to provide finer-grained information than merely online/offline, by indicating one's current status through various predefined "status messages" (e.g. "busy", "on the phone", "out for lunch", "away from the desk").

Even a relatively simple system such as IM poses a multitude of concerns regarding privacy. The importance of managing this influence on privacy, even in situations presumed to involve familiar contacts, is evident from the fact that all popular IM systems include a "Privacy" option panels with settings and options to allow individuals to manage their privacy. To obtain a better understanding of how people handle awareness and privacy in IM systems, we first compared available mechanisms for adjusting awareness dissemination and privacy protection that are provided by the four most popular IM systems – ICQ (http://www.icq.com), Yahoo! Messenger (http://messenger.yahoo.com), MSN Messenger (http://messenger.msn.com), and AOL Instant Messenger (http://www.aim.com). The following features were considered in our review:

Sound notification: This refers to the capability of associating sound alerts to various events, such as incoming messages or someone logging in.

Grouping: Grouping functionality allows various contacts to be organized into different groups, such as "Family", "Friends", "Coworkers", "Project X Members" and so on.

Privacy option panels: Privacy option panels allow modification and customization of various settings in order to manage privacy. For instance, a user may specify whether he or she wishes to reveal status information on the web, or whether his or her phone number should be available to people on the contact list.
Blocking: Blocking a contact allows a user to prevent his or her awareness information from being provided to the blocked contact. The blocked contact will always see the user as being "offline".
Custom status: The ability to set custom status, such as "Working on Documentation for Project X", improves the limited flexibility of the pre-set status messages provided by the system.
Auto reply: Auto reply functionality allows a user to reply to an incoming message with an automatic reply message when he or she is away. The message may be chosen from ones provided by the system or may be custom defined by the individual. This functionality is analogous to a non-personalized or personalized answering machine greeting.
Offline messaging: Offline messaging refers to the ability to receive messages from contacts even while being "offline". Offline messaging allows messages to be stored on the server (akin to email) and delivered at next login.
Popup notification: This refers to the capability of receiving small, ephemeral popups in a corner of the screen to serve as notification of events such as a contact signing in or a new message session being started.
Individual settings: This refers to the ability to specify various settings on a per-individual basis. For instance, a user may wish to be always "Available" to a certain contact, regardless of his or her explicitly set status. Similarly, he or she may not want a particular contact to have access to his cell phone number.
Group settings: Group settings refer to the ability to specify settings on a per-group basis. Changing a setting for a group affects the individual settings for all contacts in the group. Thus, if a person chooses to always appear "Away" to the "Friends" group, all contacts that are grouped under "Friends" will see her as being always "Away".
Video connection: Video connection allows one to broadcast live video images of oneself with a computer-attached camera. The video may either be a continuous stream or a series of snapshots taken at regular intervals.
Reciprocity: Reciprocity refers to whether or not the system enforces policies in a reciprocal manner. For instance, if person A blocks person B, a reciprocal policy will require that person A is also automatically blocked by person B.
Web status integration: Integrating status information with the web allows publishing of status information to a web site that allows others to view it without having to log into the IM system.
Permission to add: This refers to whether or not a user needs explicit permission from others to add them to his or her contact list and vice versa. A user who wishes to avoid multiple individual requests for permission may choose to set a global option to allow anyone to add him or her without explicit permission.

Table 1: Comparison of IM Systems[1]

Feature	AIM	MSN	Yahoo	ICQ
Sound notification	Yes	Yes	Yes	Yes
Grouping	Yes	Yes	Yes	Yes
Privacy menu	Yes	Yes	Yes	Yes
Blocking	Yes	Yes	Yes	Yes
Custom status	Yes	No	Yes	No
Auto reply	Yes	No	No	Yes
Offline messaging	Yes	No	Yes	Yes
Popup notification	Yes	Yes	Yes	Yes
Individual settings	No	No	Partial	Yes
Group settings	No	No	Partial	No
Video connection	No	Yes	Yes	No
Reciprocity	Yes	Partial	No	No
Web status integration	No	No	Yes	Yes
Permission to add	No	Yes	Yes	Yes

Table 1 provides a comparison of the four popular IM applications in terms of these features. While several features are common to all IM systems studied, some of them are implemented in different ways. There also exist subtle differences than those included in the table. For instance, MSN Messenger automatically turns off sound and pop-up notifications when in "busy" status (with no option to modify this behavior); Yahoo! Messenger and ICQ allow one to log in directly in "invisible" mode, while others do not, and so on. Moreover, it proves difficult and

[1] The table is based upon analysis of IM versions available at the time of writing. Since IM systems evolve continuously, a few of the cells may have since changed.

cumbersome to incorporate even the most rudimentary aspects of context into the system. This is evident from the sparse support for group-level, and individual-level settings.

7 Interviews with IM Users

We interviewed seven frequent users of IM to understand their privacy attitudes, expectations and practices when using IM, and to see whether these differed with the location and purpose of IM usage (Patil & Kobsa, 2004). We solicited participants via a posting to a mailing list, as well as via word of mouth. Potential participants were sent a short (5 questions) multiple-choice screening questionnaire in order to get a sense of the IM system(s) they used, the usage frequency, and the number of contacts in their list(s).

7.1 Subjects

Seven subjects participated in the study:

- a software developer in a large corporation
- a graphic designer in the technical staff of a university
- a software engineer in a small Indian consulting firm (with offices and clients in India and the U.S.)
- a doctoral student whose native language is Spanish
- a technical support person in a large corporation
- an engineer at a large corporation that handles sensitive defense contracts
- a second-year undergraduate student in Social Science

We deliberately chose individuals with diverse backgrounds who were involved in different types of undertakings in different types of environments, in order to compare and contrast use of IM in a broad variety of situations. Two of the subjects (graphic designer and technical support person) were female. Subjects were between mid-20 to early 30, except for the undergraduate (20) and the engineer (above 50). All were experienced users of IM and had been using it for at least a year. All used IM from multiple locations (e.g. work, home, school), and had more than 20 people in their contact lists. The frequency of IM use varied from a few hours per month to more than 8 hours/day. Subjects participated in the study on a voluntary basis, and no compensation was provided.

7.2 Methodology

A semi-structured interview of about 1 to 1.5 hours was conducted with each subject. For the graphic designer, we did a second follow-up interview of about half an hour to probe more into some of the information provided in the first conversation. In order to get a sense for the physical environment in which the subjects use IM, we tried to conduct interviews at the place where the person used IM the most (However, three took place at different locations.) All interviews were conducted face-to-face, except for the subject from India who was interviewed by phone. The interviews were digitally tape recorded and then transcribed for analysis.

We used about 20 rather broad questions as a guideline for the semi-structured, conversational interviews. The questions were meant to gather information about people's tasks and routines, the manner in which they use IM in their daily lives, and their expectations and behavior regarding privacy – both in general and specific to IM. Questions were tailored to each subject based on their answers to the 5-question screening questionnaire. Additional questions were asked during the interview, as deemed necessary to gather relevant information.

7.3 Results

Despite the diversity of chosen subjects, their expectations and practices regarding privacy were strikingly similar. (While there also were quite a few differences, we will mostly focus on the similarities here.) In general, subjects had trouble articulating what "privacy" meant to them. They found it much easier to discuss privacy in terms of concrete situations and examples. This is to be expected, given the highly "situated" nature of privacy.

All subjects claimed not being overly concerned with privacy when using IM. Most operated with the general assumption that they do not have much privacy when working online. Yet, as will become clear in the following discussion, quite a few of their practices suggest a definite desire and concern for privacy, despite claims to the contrary.

Overall, we found that privacy concerns of subjects fell along three main dimensions: who, when & where, and what. These are described below:

7.3.1 Who (known vs. unknown)

Subjects reported the desire to have a very high degree of privacy from people not on the contact list. Non-contacts were often treated as strangers with unknown intentions. Subjects took pains to make sure that anyone not on the contact lists could not see any information about them. For instance, only one subject (undergraduate student) maintained a public profile. He also indicated that many of his friends had profiles as well. We believe that this difference is most likely due to the fact that undergraduates are at an age and stage in life where they actively engage in socializing and want to "advertise" themselves.

People on the contact list, on the other hand, were treated as trusted acquaintances. Given the greatly lowered privacy barrier for contacts as opposed non-contacts, it was hardly surprising that all subjects were quite careful about whom to add to their list. The graphic designer relied upon standardized screen-name conventions followed at her workplace, the software developer used the corporate directory which was integrated with the IM client at his organization, while the doctoral student and the software engineer reported adding only those people with whom they had had extensive face-to-face relationships for some period of time.

This careful screening of contacts at the outset also resulted in relatively few contacts being blocked or deleted later. Blocking occurred either when someone was added in error, or upon some significant external change(s). For example, the software developer mentioned blocking his ex-girlfriend after they broke up. Similarly, the doctoral student mentioned deleting contacts from his old job after he quit that job.

For known contacts, subject practices pointed to a desire for different levels of availability for different groups of people – such as co-workers, family, friends. For instance, some of our subjects had reservations about having their superiors on their contact list. The doctoral student collaborated with his supervisors only via email, as he did not want to always be accessible to them via IM. This is further corroborated when subjects mentioned using IM grouping mechanism to selectively monitor their contact list.

Graphic designer: "*The IWTT members are right here. It's the first thing that I see, and I can tell my team members are on.*"

Doctoral student: "*My wife logs in and only looks at the group of family members. If no one in that group is logged in she will disconnect. That's the only group of people she cares about at that time.*"

7.3.2 When & where (availability)

This dimension can also be looked at as the desire to manage availability to avoid interruption or distraction from current task. Expectations and practices regarding availability heavily depended on location, time and (work) context. Thus, subjects had different desires regarding their availability while working (from any location including their homes), as opposed to not working. ("Work" here is used in a general sense and includes schoolwork.) While working, subjects wanted to be as available as possible to co-workers for collaboration. They also paid more attention to the availability of the co-worker contacts on their list. Subjects tried to keep contacts informed of their availability via status indicators. The graphic designer left descriptive status messages even if she was away from her desk for only 5 minutes. The software developer turned off the "auto-idle" feature, because often he was around yet not using the computer, incorrectly creating the impression that he was away from his desk. The graphic designer also mentioned that she often guessed the location of her contacts based on changes in the picture or icon that they chose to associate with their name.

Graphic designer: "*Sometimes somebody will work from home in the morning and then come in the afternoon. But the only thing that distinguishes between locations is the different icons that people might have. They might have an icon when they're at work and an icon on their home computer. And when they log in you can tell just based on an icon.*"

Subjects frequently employed "plausible deniability" (Nardi, Whittaker & Bradner, 2000) as an indication of (un)availability. They chose not to respond immediately to incoming messages if they were otherwise occupied. Similarly, a non-response to a message they initiated was taken to mean that the contact was busy and will reply at a later, more convenient time. The software engineer, however, said that he tried to send a quick "busy right now" message whenever possible.

For subjects with working lives, IM allowed a limited extension of "home" into "work". Subjects reported having personal, non-work contacts in their list at work. However, while at work, IM conversations with friends, family, and significant others were reported to be few and far between, with primary attention being devoted to work-related matters. The occasional personal conversation seemed to serve the purpose of maintaining social bonds, and catching a moment of relief from stress of work.

Interestingly, the reverse was not typically true; "work" rarely extended into "home", unless specifically working remotely from home. Subjects made sure that work did not invade personal life. The software engineer almost never used IM from home as he wanted to "stay away from the computer". The graphic designer and the software developer had separate personal IM accounts, which they used from home, while the doctoral student piggybacked on his wife's account at home. Subjects did not have any work-related contacts in these accounts.

7.3.3 What (content)

Subjects were greatly aware of the sensitivity of the contents of their IM conversations. For the most part, IM was treated similar to email or written communication. Subjects were aware of, and had accepted, that IM may be monitored by system administrators, or be sniffed off the network. Yet, just as with email, subjects had a reasonable expectation that their conversations will only be read by the intended recipient(s). The undergraduate student believed that the chances of anyone grabbing his conversations were so miniscule that he was not concerned. Moreover, they expected the recipient(s) to follow the same common etiquette as for email if sharing conversations with a third party. In fact, the graphic designer's workplace had come to an unwritten consensus about the policies to be followed for sharing saved conversations with others not part of the original conversation.

Graphic designer: "*We created rules within our group. I work with 5 people. And the rule is, anything that is said in AIM or in email, if you want to forward it on to a third party you have to check with the person first, tell them exactly what you would be clipping and pasting and sending. If they okay it, fine. But you cannot do that under any circumstances, no matter how benign the conversation seemed. You can't do that unless you've asked first. And so we stick to that rule and have not had any problems.*"

Most subjects expressed unease at the prospect of their IM conversations being saved by their contacts. However, they had resigned themselves to the fact that this was something that they could neither know about nor control. At the same time, they all cited instances in which a previously saved conversation either by them or by a contact had been useful at some later point. All seemed to employ the strategy of consciously trying to avoid saying anything over IM that might be potentially harmful for them in the future.

All subjects reported switching to a different medium of communication for conversations that they deemed too sensitive for IM. Subjects resorted to the telephone or a face-to-face conversation in such cases – either because they did not want a written record of the conversation, or because they felt that IM was too impersonal a medium, or because they felt that written communication was not the best choice for the situation, or some combination of the above reasons.

Finally, all subjects reported being aware to some extent that others who walk up to their desks were able to glance at the contents of their screens. The software developer and the software engineer said that they minimized their windows whenever someone approached their desk.

Software developer: "*I'd rather have it minimized and blinking than there for everyone to see what I'm talking about.*"

The undergraduate student also minimized windows, but only when he was conversing about the person approaching him. The doctoral student mentioned that his conversations are in Spanish, providing him with an added layer of privacy in an English-speaking country. The graphic designer as well as the doctoral student initially denied being too concerned about others watching the screen. However, further probing revealed that the doctoral student often turned off the monitor if engaged in a conversation with someone physically at his desk, while the graphic designer mentioned occasionally using the "Show Desktop" button to minimize all windows. She also recalled an instance in which she felt quite awkward when her mother was watching the screen over her shoulder.

8 Privacy-Sensitive Design of Collaborative Systems

Given that many awareness and communication facets of IM overlap with capabilities present in other collaborative awareness systems, we believe that our findings regarding privacy in IM could be generalized beyond IM to any typical awareness system. We now discuss implications of these findings for designing collaborative awareness systems in a privacy-sensitive manner.

8.1 Managing Impression

Based on the above discussion, and drawing on the work of Goffman (Goffman, 1959), we suggest that privacy expectations and behaviors in collaborative awareness systems are primarily shaped by the desire to control how one appears to others, i.e. the wish to project an appropriate "impression" of self through the system to the various parties involved. As Palen and Dourish (Palen & Dourish, 2003) point out, "We seek to maintain not just a personal life, but also a public face. Managing privacy means paying attention to both of these desires". This may be seen in the subjects' practices of presenting themselves differently by being "available" to different extent to different groups of people. Subjects' wish to control sharing of their one-on-one conversations (with any party not part of the original conversation) also points to their desire for being in command of the impression they project about themselves to the third party in question. The impression that users want to present to someone seems dependent on the type of relationship with the person. The impression one would want to project to one's superior is quite different from that one would want to project to peers. Providing information to trusted colleagues raises fewer privacy concerns than to unknown third parties. This is highlighted in the interviews by the subjects' very strong desire for privacy from people not on their contact lists.

The particular practices that people employ to manage their impression on others seem to be influenced by a variety of factors such as system defaults, personal preferences, prior knowledge and experiences, group norms, organizational policies, and cultural expectations.

The desire to manage one's impression is likely to strongly influence the point of balance between the need for privacy and the consent to disclose awareness information. One is likely to demand more privacy in matters that could potentially reflect poorly upon oneself. On the other hand, one may tolerate, or even demand, less privacy when the situation creates a favorable impression of oneself from the point of view of others. For example, due to a general fear of monitoring, employees may be reluctant to distribute records of the exact time at which they arrive at work every day. However, an employee who consistently comes in early may in fact wish to have this fact known widely as a testimony of greater commitment to work.

8.2 Implications for Design

We advocate that designers view privacy considerations in collaborative awareness systems from the perspective of "impression management". This perspective will be helpful in empowering users to appropriately monitor and manage how they appear to others via the system. It should be noted, however, that we are not urging for a mere implementation of Goffman's words. Such a direct operationalization has not been successful (Lederer, Beckmann, Dey & Mankoff, 2004) because it forces users to make implicit practices explicit. A more sophisticated, nuanced and integrated approach is needed.

In order to manage the impression projected via the system, users ought to be given the opportunity to inspect various pieces of information about themselves that can be viewed by others. Some of the subjects we interviewed reported instances in which they were unsure what exactly others could see regarding their status. Moreover, users should be able to obtain summaries and statistics about information regarding themselves that is visible to others, and a comparison of this data with averages and variances of the group(s) to which they belong. However, we recommend against providing such summaries and statistics of awareness information about individuals other than the user himself or herself, in order not to facilitate surveillance. Users could also be provided with the capability to receive timely alerts and notifications regarding changes to various factors and parameters that are of particular importance to them.

In designing collaborative systems with impression management features and mechanisms, we need to pay particular attention to three major factors:

Defaults: Given the complex and context-dependent nature of privacy, the number of options and settings to be managed is quite large. As a result, defaults that are widely applicable across persons and situations must be provided. Alternatively, typical profiles – such as "manager", "student", and "home user" – with different defaults could be utilized. In IM, we found that default settings are seldom looked at, let alone modified. The only time subjects remembered changing something is either when they first installed the software, or when something concerned them enough to take action. For example, the software developer turned off pop-up notifications and sounds as these were causing him to be distracted. And, after getting a few messages from random strangers, the graphic designer figured out how to set restrictions to allow only her contacts to view her information. This underscores the importance of setting appropriate defaults at the time of installation. The values of defaults ought to be informed by detailed studies of users, their tasks and the setting(s) in which the system operates.

Modifiable policies: Since the notion of privacy is highly nuanced, it is impossible to devise universally applicable policies. For example, a system may have the policy of not revealing one's home phone number to anyone except one's family and personal friends. However, in case of an emergency, one is unlikely to expect a rigid enforcement of such a policy. The inability to turn notifications on while in "busy" mode is an example of a policy that is too rigid. Designers should allow for user modification of default policies to provide users the control and flexibility to adjust the system as pertinent to their situation.

Interface and interaction: A great deal of attention needs to be paid to the user interface and user interaction. Feedback should be provided in a context-sensitive, non-intrusive and seamless manner. Often, subjects were unsure about their current settings for various parameters, primarily due to very little feedback and due to lack of visibility regarding current settings. This makes it very difficult to gauge the impression one projects. A system ought to give users the opportunity to inspect the various pieces of information about themselves that can be viewed by others, and also to obtain summaries and statistics about them. At the same time, subjects frequently mentioned "distractions" caused by too much or inappropriately timed feedback. Interaction should be designed in such a way that specifying and modifying one's status, settings and policies requires little or no time and effort. Automatically setting the status message to "Away" based on keyboard inactivity is an example.

8.3 Specific Design suggestions

Finally, we provide a few specific initial design explorations for future collaborative systems. Our aim is to merely stimulate exploration by providing a few seed ideas. Concrete implementations of these or other design suggestions can only develop from prototyping and iterative improvement, via user testing and feedback.

Time limits: Work and social practices may sometimes require contacts with whom one interacts only for a limited period of time. Collaborative systems that support time-limits for contacts would allow users to create short-term collaborations with a pre-specified (but later modifiable) expiration date.

Group-level settings: Given that, at any given time, people's preferences for availability to different groups of people are vastly different, a collaborative system could be designed to provide the ability to specify various settings differently for each group, in addition to a single set of global options. For example, having group-level settings in an IM system can allow one to appear busy to friends while at work, and at the same time remaining available for interruptions from colleagues at work.

Context-based settings templates: Group-based settings could be further enhanced by providing the ability to define templates for settings. Such templates could be utilized to manage availability and interruptions in different contexts or locations. For instance, a template for "work" could be defined to allow only urgent messages from family, and to grant full access to co-workers. On the other hand, when at "home", a mere change of template can make one unavailable for work-related matters, and fully available to friends and family.

Restricted archiving abilities: Some of the privacy concerns regarding content could be addressed by restricting the (digital) archiving abilities of users. These restrictions could be based on permissions or time-limits or both. For example, saving an IM conversation could be signaled to all parties involved, or may be allowed only upon explicit permission of all parties. The saved conversation could additionally be flagged by an expiration date beyond which it would not be accessible.

9 Conclusion

Awareness of the presence and activities of colleagues is valuable for effective collaborative work, all the more so when team members are geographically distributed. Systems that seek to capture, maintain, and provide pertinent information in order to raise awareness of collaborators' activities need to deal with the thorny issue of privacy. Due to its complex nature, privacy remains difficult to define. Yet, dealing with privacy concerns in the electronic domain is essential for the design of collaborative systems. Even a relatively simple system such as IM presents a multitude of privacy considerations. We found that privacy concerns of IM users fell along three main dimensions: who, when & where, and what. Users reported dealing with these issues by employing some combination of self-governance, carefully evolved practices (self, group, or social), mandated or mutually agreed-upon policies and guidelines, and software settings. We have proposed that an important driving force behind the need for privacy in collaborative awareness systems is the desire to control how one appears to others. System designers should empower users to appropriately manage the impression they project towards others via the system. This could be achieved by interface features that allow users to monitor and control how they appear to others, and by providing easily modifiable privacy policies along with suitable default settings. Time limits for contacts, group-level settings for privacy preferences (and appropriate templates for these), and restrictions on message archiving abilities of users could be explored as possibilities for improving control over impression management.

10 Acknowledgements

We wish to thank Heather Pulliam, Bonnie Nardi, Beki Grinter, Gloria Mark, and all of our interviewees.

11 References

Ackerman, M. S. (2000). The Intellectual Challenge of CSCW: The Gap between Social Requirements and Technical Feasibility. *Human-Computer Interaction*, 15, 179-203.

Bellotti, V. (1996). What You Don't Know Can Hurt You: Privacy in Collaborative Computing. *Proceedings of Human Computer Interaction Conference on People and Computers XI*, 241-261.

Bellotti, V., & Dourish, P. (1997). Rant and RAVE: Experimental and Experiential Accounts of a Media Space. In Finn, Sellen and Wilbur (eds.), *Video-Mediated Communication* (pp. 245-272). New Jersey: LEA

Boyle, M. C., Edwards, C., & Greenberg, S. (2000). The Effects of Filtered Video on Awareness and Privacy. *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 1-10.

Dourish, P., & Bellotti, V. (1992). Awareness and Coordination in Shared Workspaces. *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 107-114.

Dourish, P., & Bly, S. (1992). Portholes: Supporting Awareness in a Distributed Work Group. *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 541-547.

Goffman, E. (1959). The Presentation of Self in Everyday Life. New York: Doubleday.

Greenberg, S., & Rounding, M. (2001). The Notification Collage: Posting Information to Public and Personal Displays. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 514-521.

Grudin, J. (1988). Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, 85-93.

Herbsleb, J. D., Atkins, D. L., Boyer, D. G., Handel, M., & Finholt, T. A. (2002). Introducing Instant Messaging and Chat in the Workplace. *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, 171-178.

Herbsleb, J. D., Mockus, A., Finholt, T. A., & Grinter, R. E. (2000). Distance, Dependencies, and Delay in a Global Collaboration. *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 319-328.

Herbsleb, J. D., Mockus, A., Finholt, T. A., & Grinter, R. E. (2001). An Empirical Study of Global Software Development: Distance and Speed. *Proceedings of the 23rd International Conference on Software Engineering*, 81-90.

Hindus, D., Ackerman, M. S., Mainwaring, S., & Starr, B. (1996). Thunderwire: A Field Study of an Audio-only Media Space. *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 238-247.

Lederer, S., Beckmann, C., Dey, A. K., & Mankoff, J. (2003). Managing Personal Information Disclosure in Ubiquitous Computing Environments. *University of California, Berkeley, Computer Science Division, Technical Report UCB-CSD-03-1257*.

Mark, G., Fuchs, L., & Sohlenkamp, M. (1997). Supporting Groupware Conventions through Contextual Awareness. *Proceedings of the 5th European Conference on Computer Supported Cooperative Work*, 253-268.

Nardi, B. A., Whittaker, S., & Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 79-88.

Palen, L., & Dourish, P. (2003). Unpacking "Privacy" for a Networked World. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 129-136.

Patil, S. & Kobsa, A. (2004). Instant Messaging and Privacy. *Proceedings of Human Computer Interaction Conference*, 85-88.

Samarajiva, R. (1997). Interactivity as Though Privacy Matters. P. Agre & M. Rotenberg (eds.), *Technology and Privacy: The New Landscape* (pp. 277-310). Cambridge, MA: MIT Press.

Stone, E., Gardner, D., Gueutal, H. G., & McClure, S. (1983). A Field Experiment Comparing Information-Privacy Value, Beliefs and Attitudes Across Several Types of Organizations. *Journal of Applied Psychology*, 68(3), 459-468.

Want, R., Hopper, A., Falc_o, V., & Gibbons, J. (1992). The Active Badge Location System. *ACM Transactions on Information Systems (TOIS)*, 10(1), 91-102.

Warren, S., & Brandeis, L. (1890). The Right to Privacy. *Harvard Law Review* 4(5), 193-220.