Privacy Interfaces for Collaboration

JJ Cadiz & Anoop Gupta

September 14th, 2001

Technical Report
MSR-TR-2001-82

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Privacy Interfaces for Collaboration

JJ Cadiz & Anoop Gupta
Microsoft Research, Collaboration & Multimedia Group
One Microsoft Way, Redmond, WA 98052 USA
{jjcadiz; anoop}@microsoft.com

ABSTRACT
Privacy is a hot topic today in the worlds of technology, e-commerce, and public policy. However, the vast majority of the public debate about privacy has pitted the consumer against corporations and citizens against their government in terms of collecting and sharing data. While this debate is an important one, researchers should not ignore the technological advances that are providing people with the ability to share information with friends, family, and co-workers in valuable ways. Unfortunately, these technological advances have significant potential to infringe on people's privacy, and we don't yet know how to create interfaces that facilitate the use of these technologies while preserving people's privacy. This paper discusses the privacy problem and results from two lab studies that explored methods to overcome this problem.

Keywords
Privacy, information sharing, collaboration, privacy settings

1 INTRODUCTION

Sharing information can enable a number of simple and valuable collaboration scenarios. We share our telephone numbers in phone books so people can call us. We share our instant messaging online status so people can chat with us. We share our calendar so others can schedule meetings more easily with us.

Technological advances will likely enable even more of these types of scenarios. For example, many people currently have the luxury of sharing their business calendar with their co-workers, but sharing one's personal calendar with friends and family isn't nearly as easy. In addition, in the future we'll likely be able to share types of information that we could never share before. Soon cell phone companies will have the ability to determine the location of all their phones [18], and initiatives like Microsoft's .NET and Hailstorm frameworks seek to place a variety of information about people on-line [17].

However, the same technology that allows us to share information also poses a threat to our privacy. In the past few years, tremendous amounts of energy have been spent by both technologists and public policy makers on the topic of privacy. In September 2001, The New York Times ran a series of front page stories about privacy and reported that at least 50 privacy-related bills could be considered by the US Congress in the next year, and that 67% of Americans identify online privacy as a big concern (only 55% say the same about fighting crime) [16].

Unfortunately, the majority of the energy spent on the topic of privacy has been focused on the battle between consumers and businesses, or citizens and the government. While this debate is an important one, underlying it is an opportunity for many of today's interface researchers: while people are extremely concerned about sharing information with corporations or their government, people are often happy to share information with friends and family if it provides them with enough value. In the future, people will likely desire interfaces that allow them to easily share a wide variety of information with others.

However, creating these types of interfaces is a hard problem. Providing people with privacy means providing them with easy ways to control information about themselves, but the decisions people make about how, when, and with whom to share information are highly nuanced. As Ackerman [1] and Neumann [13] have noted, in the realm of privacy, there is a gap between what we need and what is possible with today's interfaces when it comes to privacy. Starting to address this gap is the goal of this paper.

In the next section we'll provide a review of the current literature on privacy. We present our approach to privacy interfaces in section 3, and in sections 4 and 5 we discuss two preliminary studies we performed. In section 6 we outline various future directions for addressing the privacy problem.

2 RELATED WORK

Perhaps the best place to start when discussing privacy is a definition. When people say they want privacy, what do they mean? Tavani [19] discusses several definitions of privacy, including, "being free from unwarranted intrusion", "being alone", and "being able to limit access to information about oneself." However, Tavani believes the best definition is one that recognizes that people like the ability to share information with others. Thus, Tavani writes, "privacy can best be understood as the condition of having control over information about oneself."

2.1 Privacy with Businesses and Governments

As noted in section 1, much of the research and debate about privacy has concentrated on situations where people haven't been given sufficient control over the information that businesses and governments possess. Specifically,

people are worried about large organizations collecting highly detailed information about them and using it in ways that were never approved.

Addressing these privacy issues is important, but it's not the focus of this paper. Instead of focusing on how people can control information possessed by businesses and governments, this paper's focus is on how we can give people more control over their information so they can share it in valuable ways with their friends, family, and co-workers. However, examining the literature with regard to businesses and governments is still valuable as it reveals principles and solutions that are applicable to our problem.

For example, Laudon [10] discusses the fundamental principles that US and European law are based on and how technology makes these principles obsolete. For example, many of the principles rest on the notion that people should know all the databases in existence that have information about them, but with the sheer number of databases today, it's impossible for any one person or organization to know about all such databases.

In addition to studying legal principles, researchers have examined user interface principles when creating privacy systems. Ackerman, Cranor, and Reagle [3] reported a survey of privacy preferences of 381 Internet users. Perhaps the most valuable finding of this survey was a confirmation of Westin's [20] finding that people tend to fall into one of three categories when it comes to privacy: *privacy fundamentalists* (people so concerned about privacy that they are usually unwilling to share information with web sites), *marginally concerned* (people usually willing to share information with web sites), and the *pragmatic majority* (people whose concerns fall in between the previous two groups). As noted in [3], an important implication of this finding is that the technique of hiding system complexity by using intelligent defaults won't work with such a diversity of privacy concerns across people.

Data from this survey supported what is perhaps the most recent innovation with regard to privacy and technology: the Platform for Privacy Preferences Project, known as P3P [14]. The goal of P3P is to give people an easy method of controlling what information they share with web sites. People enter their privacy preferences into web browsing software (or copy preferences from a trusted source), and when they visit web sites, the sites pass the browser their privacy policy in a machine-readable format. If a user's privacy preferences match the web site's policy, then there are no issues. If there isn't a match, then negotiation has to occur to determine if the user can browse the site in an acceptable way.

As [5] note, part of the difficulty of creating P3P was that they were creating a social protocol, not a technical protocol. With technical protocols, it's possible to create an incredibly complex system that allows any user to set any preference when it comes to sharing information, but with social protocols, systems must be far less complex to be tolerable for users browsing the web.

A related example of a simple interface for protecting privacy is the "privacy critic" developed by Ackerman and Cranor [2]. Privacy critics watch the web pages you're browsing and alert you if you visit a web site that is known to do things that you consider violations of privacy.

2.2 Privacy with Friends, Family, and Co-Workers

As noted above, while controlling information that businesses and governments have is an important problem, the focus of this paper is on how people can share information with friends, family, and co-workers to enable rich collaboration scenarios. Fortunately, there has been some research in this area. A good portion of this research has examined methods for people to share video with each other while respecting their privacy. For example, Hudson and Smith [8] proposed filtering video such that instead of showing people moving throughout a room, the video would only show shadow figures. Zhao and Stasko [21] examined similar techniques, along with Boyle et al [4] who performed a controlled study to determine how much a video should be altered to preserve both privacy and utility.

The idea of purposely obscuring perfectly good information to preserve privacy is an important principle. Although technologists (who often have the job of creating reliable, accurate systems) may find the concept counterintuitive, when it comes to systems that track sensitive information, adding noise and ambiguity can be helpful. For example, researchers who studied the use of instant messaging systems found that users liked that the systems were far from perfect when it came to reporting someone's status [12]. Specifically, if people received an instant message but didn't want to reply to it, they could ignore it, knowing that the other person couldn't be 100% sure that they were actually at their desk.

Greenberg and Fitchet [6] also examined the problem of sharing video streams, but they approached the problem differently. Instead of altering the video to preserve people's privacy, they created a video camera with a proximity sensor such that the quality of audio and video that was broadcast depended on how close the person was to the camera. This work is an excellent example of using social protocols to mediate the amount of information shared between people.

Video is an important case to examine because of its potential both for benefit and for violating people's privacy, but clearly there are other types of information that people want to share. Grudin's study of shared calendar systems [7] resulted in the classic finding that information sharing can break down when the people who must do the work to share the information aren't the same people who will benefit.

Other research on additional types of information to share includes Lau et al. [9], who discussed CollabClio, an

interface for people to share information about which web pages they had visited. CollabClio looked at mechanisms for people to describe which information they wanted to share with whom, as well as mechanisms for indicating to people that they were performing actions that could potentially be viewed by others.

3 OUR APPROACH

The literature reviewed above provides some excellent individual solutions for specific types of information. However, the types of information we'll be able to share in the future will likely increase to the point where creating individual solutions will no longer be feasible.

Ideally, we'd like an interface that would allow any person to share any type of information with any other person. Furthermore, this interface has to be easy enough such that it doesn't get in the way of normal social interactions. For example, today if I want to share my work contact information with a person, it's as easy as giving that person a business card.

Perhaps the most straightforward interface to imagine is a two-dimensional matrix where each type of information is a row and each person is a column, with each cell being a checkbox that people could use to enable or disable access for a person for a type of information. However, clearly such a matrix would be overwhelming for users, and creating intelligent defaults doesn't work well when the literature shows that people's privacy preferences don't fall cleanly into a single pattern. As Ackerman [1] notes, this is a "wicked" problem.

Unfortunately, we don't have all the answers to this problem, but we ran a set of exploratory studies to learn more about it. We report results from two of these studies in sections 4 and 5. In our first study, we wanted to examine how people made privacy decisions. We also wanted to study whether people would be more comfortable sharing information if 1) they could be notified when others accessed their information, and 2) if they could place time limits on when people would have access to their information.

We hypothesized that being notified when information was accessed would make people more comfortable with sharing some types of information. For instance, I normally do not want my co-workers to have access to my spouse's cell phone number, but in the event of an emergency, it would be ok if they looked it up. When they accessed the number, I would receive a notification and be able to evaluate whether it was appropriate. In the case that it wasn't I could talk to them.

We also hypothesized that providing time-bounded access would make people more comfortable with sharing information. For example, normally I would never want my tax consultant to have detailed information about whether I'm in my office or not. However, during tax season, often he has brief questions and needs to call me, so often he tries (and fails) to reach me by calling my office. During a few days out of the year, it would be valuable for both of us if I could share information with my tax consultant so he could tell when I could be reached via my office telephone.

Our first study examined whether adding these two features would make people more comfortable with sharing information. In our second study, we explored simple methods that people could use to specify which people in their lives could have access to their personal information.

4 FIRST STUDY

As noted in the previous section, our first study explored whether providing notification that information had been accessed would make people more comfortable with sharing information, and whether providing the ability to make information available only during certain times would make people more comfortable with sharing information.

4.1 Methodology

For this study, we recruited seven participants (4 women, 3 men). Participants were from a variety of backgrounds and had varying computer expertise. Participants received a software gratuity for their participation.

The main tool for the study was a spreadsheet. Along the top was a list of types of people (see Table 1). Participants were asked to fill in names (first names only) of these people, skipping any people that weren't relevant. Along the left column were a variety of types of information, ranging from basic contact information to low-level status information. The types of information and the groups they were presented in is shown in Table 2.

For all of the types of information (especially the medical information), we told participants that we did not want them to reveal the information, but that we wanted to know if they would share this information with the various people on the spreadsheet. Participants were also told that they could refuse to answer any question for any reason.

After participants filled in the first names of various people and reviewed the types of information, participants were then asked to imagine a secure web site that would allow them to share any type of information on the spreadsheet with any person. Participants were then asked to indicate how comfortable they would be sharing each type of information with each person by placing a number in each cell of the spreadsheet, using a 5-point scale where 1 = "Very uncomfortable sharing this information with this person" and 5 = "Very comfortable sharing this information with this person". As participants filled out the spreadsheet, they were asked to think aloud.

Table 1: Types of People Participants Considered

People you live with
Family members
Close friends
Other friends
Your manager
People who work with you
People you work closely with
Other people you work with
A person at your company whom you don't know
A stranger ("John/Jane Doe")

After participants filled in the spreadsheet for the first time, they were told to imagine that the system now had a feature where they could chose to be sent an e-mail when a person accessed their information (participants were told this e-mail would go to an inbox separate from their main inbox). Participants were then given the opportunity to change any of their comfort levels on the spreadsheet.

Participants were then told to imagine that the system had the ability to allow them to restrict access to information to certain times (for example, only sharing information about your location with your co-workers during work hours). Once again, people were given the opportunity to change any of their comfort level scores on the spreadsheet.

To make the task of filling in the spreadsheet easier, the types of information were placed in logical groups (as shown in Table 1) and, if appropriate, participants could mark a single cell to indicate that the score would be the same for an entire group of information (for example, all information pertaining to one's medical condition). In addition, the experimenter and the participant shared the spreadsheet via NetMeeting and participants had the option to either work on the spreadsheet themselves or have the experimenter fill in the numbers for them.

4.2 Results

Before presenting results from this study, it's important to note that participants were asked to make decisions in a lab setting about hypothetical situations, and thus data from this study should be reviewed with caution. At the same time, technology could make these hypothetical situations a reality in the future and people will be faced with similar decisions, thus these results are an interesting snapshot of how people might make these decisions.

First, for a variety of reasons, we choose not to report tables of numbers here showing, overall, how comfortable people were with sharing the various types of information with the various types of people. Combining all the scores from our participants about how they would share information with their friends, families, and co-workers seems dangerous given the diverse types of relationships people have with people whom they consider friends or co-workers, and given the highly nuanced decisions we'll discuss below. However, from a qualitative standpoint, we can comfortably say that overall, people were open to the idea of sharing information. No one was categorically opposed to the idea, except for the cases where people considering sharing information with people whom they didn't know ("unknown co-worker" and "John/Jane Doe").

Table 2: Types of Information Participants Considered. Information was presented in the groups listed below, and to simplify the study, people could chose to work with the groups instead of the individual types of information.

Home Contact Information	My home address My home phone number My personal e-mail address
Work Contact Information	My work address My work phone number My work e-mail address
Other Work-Related Information	My list of work contacts My work calendar (free/busy info) My work calendar (all info)
Misc Contact Info	My cell phone number My pager number
Medical Information	My doctor's name and phone number What chronic health conditions I have What medications I'm currently taking
Personal Leisure Information	My favorite hobbies My favorite music My favorite movies My favorite vacation spots
Status while at work	My current physical location Whether I'm on the phone right now Whether I'm in a meeting right now Whether I'm in my office Whether I'm typing at my keyboard The best way to contact me right now Whether I can be interrupted right now How many people are in my office Whether my phone is off the hook Whether there's activity at my computer
Status while not at work	My current physical location Whether I'm at home Whether I'm in my car Whether I'm on the phone Whether people are in my house Whether my car is moving Whether my phone is off the hook

Our main finding from this study was that the two features we studied—having the option of being notified when others accessed their information and being able to specify time limits on when information would be available—were both features that people said they wanted, but neither feature noticeably increased people's comfort with sharing information. In fact, when given the chance to make changes to the numbers on the spreadsheet, only one person made substantial changes, and three others made very minor modifications. When we asked people about these features, their general feeling was that the features would be nice, but they wouldn't cause them to share any information that they

wouldn't have already shared. People seemed to make general judgments about whether they trusted people and whether people needed to have the information, and thus the additional features of notification and time-bounded access had no effect on these judgments.

In addition, when listening to how people made judgments about whether they wanted to share information with a person, we noticed a pattern of four questions that people tended to ask themselves. These four questions were:

Does this person already have this information?

Does this person need to know this information?

Do I care if this person has this information?

Is this person trusthworthy?

However, the relative importance of these questions was not always consistent. For some people (perhaps people that would be classified as "marginally concerned" [3]), the most important question was whether they cared if anyone had the information. Thus, even if they couldn't think of a single reason that the person needed the information, they would share it anyway if they didn't care. For other people (who might have been classified as "privacy fundamentalists"), the most important factor in their decision was whether they could think of a credible reason that the person needed to know the information.

We noticed several other interesting nuances. For example, we thought that everyone in a household might be treated the same when it came to sharing information, but kids tended to be treated quite differently. Some parents seemed to make judgment calls based on whether they thought their kids were mature enough to have access to some of the information, while other parents knew their kids were mature enough but they still didn't want their kids to know certain things. In particular, one father didn't want his teenage daughter to be able to see his location because he didn't want her to be able to determine when he'd be coming home.

As another example, people tended to be open with sharing information with their friends, but one person shared relatively little information with one friend because the friend was a known gossiper. Overall, these observations support Ackerman's observation that decisions about privacy can be highly nuanced [1].

5 SECOND STUDY

The goal of the second study was to start exploring methods that would make it easier for people to share information with others. Specifically, we explored two possible concepts.

First, people are accustomed to sharing information using business cards. When we print business cards, we make decisions about what information to put on them. Furthermore, the act of giving someone a business card explicitly gives the person permission not just to have the information, but to use it as well. Thus, one goal of this study was to see if the business card concept could be used for a broader range of information and people.

Second, we hypothesized that people are able to determine that one person is like another person when it comes to privacy preferences. For example, if I make a set of judgments about which information a friend should have access to, I might be able to say that another friend should have the exact same access. We call this method the "person-prototype" method.

5.1 Methodology

As in the first study, we recruited seven participants. Participants came from a variety of backgrounds and had varying computer expertise.

The methodology for this study was similar to the methodology for the first study, with a few key differences. First, instead of indicating comfort level with sharing information using a 5-point scale, people were asked to make a binary decision abut whether they would or would not share each type of information with each person. Second, at the beginning of the study, people were asked to list the first names of a variety of friends, family, co-workers, and other people whom they knew. These names were then split into three different sets.

For the first set of people, participants were asked to indicate whether they would share each type of information with each person (much like the first task of the first study). Next, we explained to participants the idea of "business cards" and had them create a variety of cards that they could give to people on the spreadsheet. Participants then were given the second set of names and asked to give each person a card. As participants selected a card, the experimenter pasted the card's values into the spreadsheet, showing the participant which information would now be shared with that person. Participants were then asked how satisfied they were with the values (using a 5-point scale where 1 = "very unsatisfied" and 5 = "very satisfied").

For the final task, participants were shown the last set of names and asked to determine which information to share with each person by selecting a person from the first list. For example, if a person was deciding which information to share with her brother, she may choose her father as the person most like her brother. As in the card-based task, after participants selected a person, the experimenter copied and pasted the values to show the participant which information would be shared with the person, and participants were then asked to specify how satisfied they were with the values on a 5-point scale.

Note that the order for the second two tasks was reversed for half of the participants. Thus, half the participants did the card-based task first and the person-prototype task second, while the other half did these tasks in the opposite order.

5.2 Results

The main finding from this study is that people were quite able to set their information sharing preferences using both the card and the person-prototype technique. In addition, for both methods, satisfaction ratings were very high. When using cards, participants rated the method an average of 4.6 out of 5, and when using the person-prototype method, the average rating was 4.7. A Wilcoxan Signed Ranks test found that the difference between these two ratings was not significant ($z = -0.73$; $p = .47$).

When looking at the types and numbers of cards created, the data weren't surprising. People created an average of 3.7 cards (one person created just 2 cards, while one created 5), and the cards were labeled either using groups of the people whom the person associated with ("West Hill Community Council", "Spiritual People", "work"), or levels of relationships ("closer friends", "general friends", "relative strangers").

Of course, although the data from this study are positive, both of the techniques used require a startup process. No cards can be assigned and no people can be used as prototypes when the system is used for the very first time. However, once the initial configuration is completed, both of these methods are promising as techniques for providing people with easy methods of sharing information.

6 DISCUSSION

While the studies we discussed in the previous section contribute some interesting information to the privacy problem, it's clear that much ground remains to be covered. In this section, we'd like to outline several future design directions to consider.

6.1 Not All Information is the Same

As we've considered the variety of types of information that could be shared in the future, it's become clear that not all information has the same properties, and thus not all information can necessarily be treated the same way in interfaces. Following are two dimensions we've considered.

Fast vs. Slow Changing Information

One distinction is between fast and slow changing information. Once someone finds out your birthday, they'll have access to it forever regardless of whether you revoke access to it in some computer system. However, the same isn't true of your current online status: the second you revoke access to the information, a person can't be sure of its state. Of course, there are several types of information that fall in between these two examples, ranging from e-mail addresses to telephone numbers to (for some people) hair color.

The implication of this distinction is that interfaces may need to treat information differently based on whether access to the information can be revoked. Perhaps interfaces should be more careful when dealing with slow-changing information because of the difficulty in revoking access.

Descriptive vs. Communicative Information

Another distinction is between descriptive information and communication information. Descriptive information tells others about your personality and lifestyle. Examples include age, hobbies, job, and friends. In contrast, communication information is information that people can use to get in touch with you, but it doesn't necessarily tell others about the type of person you are. Examples include your phone number, e-mail address, and mailing address. Of course, information can rarely be purely communicative or descriptive, given that your e-mail address may indicate where you work, or your phone number may indicate the region where you live.

However, there's a type of information that clearly overlaps both descriptive and communicative information: status information. Examples include my on-line status, my location, and so on. This information is often used to contact people but is also inherently descriptive when considering the dimension of time. For instance, just knowing whether you're in your office right now is mostly communicative but not very descriptive. However, if I have historical data on when you're typically in your office, the data becomes much more descriptive and less communicative. The implication is that when dealing with status information, designers should not to create systems that store or aggregate status information over time as it can lead to information being used for reasons that it wasn't originally intended for (which is a violation of one of the basic principles of information collection [10]).

The other implication is that there's a different cost for losing control of different types of information. If someone gets access to communication information without my consent, the worse they can do is annoy me with phone calls, e-mails, instant messages, etc. If someone gets access to descriptive information without my consent, they'll simply know something about me that I didn't want them to know, which I may or may not find damaging.

6.2 Using Proxies

While it may be difficult to determine exactly who should have access to which information in all circumstances, it's often easier for us to identify one or more people who could make the judgment call if necessary. For example, managers may give their assistants access to a tremendous amount of information (entire calendars, contact lists, and e-mail contents) so that the assistants, not the managers, can have the responsibility of determining who should have access to which information. For example, if I need to find the vice president of my division and he's not in his office, I can ask his assistant and she can make the decision as to whether I should have access to that information. Similarly, parents and employees are often asked to provide an emergency contact that can be called in case of an emergency.

Proxies are sufficiently useful enough that entire businesses have been formed around the concept. MedicAlert (www.medicalert.org) is an example of a proxy service for medical information. Most of us wouldn't pay to share our intimate medical information with a large business, but people do exactly that with MedicAlert. For an annual fee, people disclose their medical records to MedicAlert and are given a bracelet to wear. The bracelet lists their medical conditions, an ID number, and a phone number to call for more information. If a person is ever found unconscious, paramedics can use the information on the bracelet to increase the chances of saving the person's life.

In addition, some software systems have started to take advantage of the proxy concept. The Wildfire system (www.wildfire.com) is a proxy service for incoming phone calls. In addition, some of the agent-like systems covered in section 2.1 could be considered software proxies that control access to information.

Of course, it would be wonderful if we were capable of creating software proxies that had the same good sense as a world-class administrative assistant, but this isn't a solution that we'll likely have soon. However, human proxies remain a possible component of privacy-preserving systems, in addition to systems that combine both human and software proxies.

6.3 Inferring Privacy Preferences from Social Actions

Perhaps the ideal system would be able to weave itself into our everyday lives and allow us to share information using known social actions. For example, as noted earlier, giving a person a business card is a known social action that gives someone access to your information.

Of course, inferring privacy preferences from social actions is difficult for at least three reasons. First, identifying which actions should lead to which privacy preferences is hard. Second, while there may be social actions for giving someone access to information, there aren't always actions for revoking access. For instance, if you give someone a business card with your cell phone number on it, but three months later you no longer want that person to have that information, how do you revoke access? Third, there aren't social actions for sharing all of the information we may be able to share in the future. For example, the only way I can share my current location with you is to call and let you know where I am. This is often hardly worth the effort unless there's an exceptional condition.

Thus, one potential solution to the privacy problem is designing new social actions for sharing information, or modifying current actions such that they carry the additional function of sharing information. For example, soon we may have "smart cards" that are like credit cards with the addition of a computer chip. These cards could be used to temporarily give access to information about myself to others.

For instance, if I was at my dentist and he wanted to schedule a follow-up appointment for me in two weeks, I could give the receptionist my card and he could slip it into a reader to get access to my calendar. Then, we could look at our calendars side-by-side and schedule the appointment on both our calendars. When finished, the receptionist would remove the card and give it back to me, which would revoke the receptionist's ability to see my calendar.

Meeting Requests with Location Awareness Permissions

Another example has to do with calendaring systems that allow people to send each other meeting requests (like Microsoft's Outlook). If I wanted to meet with John next Monday, I could send John a meeting request and he could accept, tentatively accept, or decline the meeting. If he accepted, the meeting would be entered on his calendar. John's decision is sent back to me via an e-mail message.

However, one problem I have when I meet with John is that it takes me about 20 minutes to drive to his office. When traffic is ok, there's no problem. However, when traffic is heavy, it can take me as long as 40 minutes to get to John's office. When I get stuck in traffic, I worry about making John wait and appearing rude.

However, if we had technology that provided awareness of people's whereabouts, this problem could be alleviated. My calendaring software could be configured such that attached to all my meeting requests is a token providing permission to the meeting attendees to see my physical location 15 minutes before and during the entire meeting time. If I have lunch with John every Tuesday from 12:30pm to 1:30pm, he would be able to see my physical location from 12:15pm to 1:30pm.

This functionality would help John because he could continue working until he saw that I was about to reach his building, and this functionality would help me because I wouldn't have to worry about making John needlessly wait for me if I was running late.

6.4 Groups of Similar Information and People

Another area of research that could be fruitful is determining groups of people and groups of information that users tend to treat similarly. For example, if I always share my home phone number with people who have my home address, systems could allow people to treat both types of information as one unit. The same concept could work for groups of people. While we started to use this concept in our studies by allowing people to work with the information in the groups shown in Table 2, our study did not systematically examine which types of information tend to be treated the same. The same is also true of exploring groups of people: the business card concept allowed people to share information with others using a concept related to groups, but determining which groups of people are treated the same was not an explicit goal of our study.

7 CONCLUDING REMARKS

Privacy is a problem that has no easy answer, but with technological advances on the horizon, it's also not a problem that's going away. Furthermore, even if researchers design several excellent ways for people to share information in one nation, it's probable that these methods will not work in other nations due to cultural differences [11]. Regardless, filling the gap between what society needs and what is technologically feasible in the domain of privacy is an important challenge for our community to undertake. It's our hope that the results and ideas presented in this paper will help others pursue this challenge.

ACKNOWLEDGMENTS

We thank Microsoft's usability coordination center for scheduling all the people who participated in our studies.

REFERENCES

1. Ackerman, M. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. To appear in *Human-Computer Interaction.* Preprint available at http://www.ics.uci.edu/~ackerman/
2. Ackerman, M., and Cranor, L. (1999). Privacy Critics: UI Components to Safeguard Users' Privacy. *Extended Abstracts from the ACM Conference on Human Factors in Computing Systems* (CHI 1999).
3. Ackerman, M., Cranor, L., and Reagle, J. (1999). Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences. *Proceedings of the ACM Conference on Electronic Commerce* (E-COMMERCE 1999).
4. Boyle, M., Edwards, C., and Greenberg, S. (2000). The Effects of Filtered Video on Awareness and Privacy. *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW 2000.
5. Cranor, L., and Reagle, J. (1997). Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project. *Proceedings of the Telecommunications Policy Research Conference*.
6. Greenberg, S., and Fitchet, C. (2001). Phidgets: Easy Development of Physical Interfaces through Physical Widgets. *Proceedings of the ACM Symposium on User Interface Software and Technology* (UIST 2001).
7. 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* (CSCW 1988).
8. Hudson, S., and Smith, I. (1996). Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW 1996).
9. Lau, T., Etzioni, O., and Weld, D. (1999). Privacy Interfaces for Information Management. *Communications of the ACM*, 42(10), October 1999.
10. Laudon, K. (1996). Markets and Privacy. *Communications of the ACM*, 39(9), September 1996.
11. Milberg, S., Burke, S., Smith, J., and Kallman, E. (1995). *Communications of the ACM*, 38(12), December 1995.
12. Nardi, B., Whittaker, S., and Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW 2000).
13. Neumann, P. (1994). Expectations of Security and Privacy. *Communications of the ACM*, 37(9), September 1994.
14. Reagle, J., and Cranor, L. (1999). The Platform for Privacy Preferences. *Communications of the ACM* 42(2), February 1999.
15. Rotenberg, M. (1992). Inside Risks: Protecting Privacy. Communications *of the ACM*, 35(4), April 1992.
16. Schwartz, J. (1999). "Giving the Web a Memory Cost Its Users Privacy", "As Big PC Brother Watches, Users Encounter Frustration", and "Government is Wary of Tackling Online Privacy." *The New York Times*, September 4th – 6th, 2001.
17. Sullivan, B. Microsoft to Charge for "Hailstorm." MSNBC, http://www.msnbc.com/news/546578.asp
18. Sullivan, B. They'll Know Where You Are: cell phone Tracking Technology Raises Privacy Issues. MSNBC, http://www.msnbc.com/news/536116.asp
19. Tavani, H. (1996). Computer Matching and Personal Privacy: Can They Be Compatible? *Proceedings of the Symposium on Computers and the Quality of Life* (CQL 1996).
20. Westin, A. (1991). Harris-Equifax Consumer Privacy Survey 1991. Atlanta, GA: Equifax Inc.
21. Zhao, Q., and Stasko, J. (1998). Evaluating Image Filtering Based Techniques in Media Space Applications. *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW 1998).