ICST Transactions
on Security and Safety

Research Article ICST.ORG

A framework for usable end-user privacy control in social software systems

Maryam Najafian Razavi*, Denis Gillet

Ecole Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland

Abstract

Recent studies have shown that many users struggle to properly manage selective sharing of the diverse information artefacts they deposit in social software tools. Most tools define privacy based on the 'network of friends' model, in which all 'friends' are created equal and all relationships are reciprocal. This model fails to support the privacy expectations that non-technical users bring from their real-life experiences, such as enabling different degrees of intimacy within one's network and providing flexible, natural means of managing the volatile social relationships that social software systems confront. Furthermore, the model suffers from lack of empirical grounding and systematic evaluation. This paper presents a framework for building privacy management mechanisms for social software systems that is intuitive and easy to use for the average, non-technical user population of these systems. The framework is based on a grounded theory study of users' information sharing behaviour in a social software tool. Results inform the design of OpnTag, a social software prototype that facilitates personal and social information management and sharing. Preliminary empirical data suggest that our proposed privacy framework is flexible enough to meet users' varying information sharing needs in different contexts while maintaining adequate support for usability.

Keywords: grounded theory, people tagging, privacy, social software, usability, Web 2.0

Received on 15 February 2011

doi: 10.4108/trans.sesa.2011.e4

1. Introduction

Social software systems are a family of Web 2.0 applications, characterised by their primarily user-driven content and the ability to mediate personal and social information across collectivities such as teams, communities, and organisations. The advent of the 'Social Web' has made users producers as well as consumers of information, resulting in publishing and distributing huge amounts of user-created data. Examples of social software systems include social authoring tools (e.g. wikis), social bookmarking tools (e.g. del.icio.us), and social networking tools (e.g. LinkedIn and Facebook). Users are widely adopting these tools for personal and social information management because they provide significant enhancements in utility and cost over similar desktop tools, in the sense that not only they allow their users to create personal information spaces that are easily accessible from anywhere on the Web, but also give them the tools to share their various information artefacts with others and take advantage of others' shared artefacts. These two advantages are in many cases so strong that users are either explicitly willing to give up control of that information or do so without any real awareness of the degree to which they are doing so. Recent research in the area of knowledge management (KM), however, has recognised the need to improve people's ability to control who sees what from the information they deposit in their online personal spaces (Erickson, 2006). Nevertheless, the topic of personal privacy—how people manage privacy of their own information with respect to other individuals (as opposed to organisational privacy (Iachello and Hong, 2007)—remains largely unexplored in the research literature.

Current state of the art with privacy management in social software is that most tools either define access control as a private/public dichotomy, not accounting for the various other shades of privacy in between (e.g. del.icio.us), or they need users to self-administer fine-grained privacy control on their data through a privacy setting

*Corresponding author. Email: maryam.najafian-razavi@epfl.ch

page (e.g. Facebook), which is complicated and cumbersome for the average, non-technical users. These two approaches pretty much present the two ends of the spectrum for privacy management in social software systems, with most tools falling somewhere in between. While mechanisms at one end of the spectrum do not provide sufficient control over privacy of one's information, mechanisms at the other end are not usable enough for the average, non-technical user population of social software systems.

The purpose of this work is to present a usable framework for providing end-users with better control over selective sharing of information they deposit in a social software system. The objectives of the work fall into three general areas:

(i) Developing a better understanding of users' perspectives on personal privacy in social software domain, in terms of the extent of the problem, specific privacy needs and concerns, and strategies that users employ in order to achieve their desired levels of personal privacy.

(ii) Identifying factors that impact users' information sharing behaviour in this domain, in order to build a conceptual model of personal privacy that matches users' mental models.

(iii) Devising guidelines for building privacy management mechanisms in this domain that satisfy users' varying privacy needs, and yet, are usable for the average, non-technical user population of social software systems.

In order to meet these objectives, we first employed a grounded theory study to develop an understanding of the information sharing process in the context of social software systems. Based on the results of the study, we propose a set of design heuristics for privacy management in social software, and consolidate those heuristics into a framework for privacy in this domain. Our proposed framework supports per-artefact privacy management as opposed to per-category privacy management supported by most other tools, and enables definition of social contacts of non-equal weights through creation of egocentric groups.

In order to create a test bed where the suitability of the proposed principles can be tested, we next introduce OpnTag, a social tool we developed whose privacy management mechanism instantiates the proposed framework, and present the results of an empirical evaluation that provide initial validation that the framework is flexible enough to meet users' varying privacy needs and indicate areas for future enhancements.

2. Related work

In recent years, use of social software has moved from niche phenomenon to mass adoption (Gross *et al.*, 2005; Millen *et al.*, 2006). This increase in use has been accompanied by diversity of purposes and access patterns. As a result, researchers have studied several issues that pertain to these tools, including people's attitudes towards disclosing personal data. Gross *et al.* (2005) report on a study of patterns of information revelation in online social networks and their privacy implications. Their results are based on actual field data from more than 4000 users of Facebook. They report that patterns of information revelation depend on a number of factors, including pretence of identifiability, type of information revealed or elicited, and the degree of information visibility. Along the same line, Darrah *et al.* (2001) observes that people tend to devise strategies to restrict their own accessibility to others while simultaneously seeking to maximise their ability to reach people, and Westin (2003) argues that privacy management is the continuous act of balancing the desire for privacy with the desire for communication and disclosure.

Researchers have also studied users' attitude towards revealing information in several other contexts, including workplace and location-aware mobile services. Olson *et al.* (2005) took a quantitative approach in conducting an in-depth survey of people's willingness to share a range of everyday information (such as web sites they visit or their health status) with various others, including family members or co-workers. They pointed out that whether data are anonymised or can be tied directly to people play a major role in people's willingness to disclose. Other relevant factors reported include general attitude towards privacy based on Westin privacy indexes (privacy unconcerned, pragmatist, or fundamentalist) (Westin, 1991), and personal judgement regarding 'appropriateness' (i.e. relevance) of sharing certain information with certain groups.

In another work, Patil and Lai (2005) conducted a study on the privacy vs. awareness trade-off to identify the kinds of information that users of an awareness application are willing to share with various others (team-mates, family, friends, managers, etc.) for various purposes in the context of the workplace. They identified which clusters of awareness information are more likely to be shared with whom and in what context (i.e. 'team members' received comparable levels of awareness sharing with 'family' during work hours).

Whalen and Gates (2005) reported on a small-scale study of the type of personal information that users would be willing to disclose in open online environments, primarily focussing on uncontrolled spaces such as search engines. Their results, although limited in scope, point to the existence of consistencies in the way people treat certain classes of information, which suggests it might be possible to group related information into clusters that are treated similarly.

Recent work in KM has also recognised the need to improve people's ability to control who sees what in their

personal information. Erickson (2006) explored the concept of personal information management in group context, by arguing that when personal information is to be shared with a group, the way it is used, and managed changes. In that article, he defined Group Information Management and identified many research questions that need to be explored, including how personal information is shared within a networked group, the norms of personal information sharing within groups, and the way those norms are negotiated in the group.

Palen and Dourish (2003) clarified the difference between the problem of personal information privacy and that of access control, by arguing that privacy is a continuous process of negotiating boundaries of disclosure, identity, and time, rather than a definitive entitlement. They observed that people in social software systems might act simultaneously in different spaces: as individuals, as members of a family, members of some occupational group, etc. In each of these affiliations, they may choose to disclose different information to different audiences. Palen and Dourish then exposed the unsuitability of existing access control models for privacy management since the conventional separation of one's network into 'roles' (as done by existing access control models) fails to capture the fluid nature of these various genres of disclosure in which one acts.

We therefore follow Palen and Dourish's lead and adopt the term 'privacy management' to mean the user-centred expression of personal (and organisational) constraints on information sharing. This is distinguished from 'access control', which is the means by which systems enable and enforce these choices. The main problem of privacy management in social software systems then is how to reconcile the co-presence of various groups that one identifies with, by providing users with flexible, non-overwhelming means to control what to share with whom when, which is the focus of this work.

3. A study of information sharing behaviour in social software

As a first step in this research, we performed a grounded theory study to understand end-users' information sharing behavior in social software systems and to identify specific privacy needs and concerns in this domain. This section describes the study and the theory that was derived from it.

3.1. Methodology

The research method adopted in the study was grounded theory (Glaser, 1978, 1998; Glaser and Strauss, 1967). It is a primarily inductive investigation process in which the researcher aims to formulate a small-scale, focussed theory that is derived from the continuous interplay between data analysis and data collection. Rather than starting with a preconceived theory that needs to be proven, the researcher begins with a general area of study and allows the theory to emerge from the data. Such theory has been claimed to have a better chance at resembling reality, compared to theory that is derived by putting together a series of concepts from solely speculation on how one 'thinks' things should work (Glaser, 1992).

3.2. Locating the study

Since a grounded theory method looks for emergence of theory from the data, grounded theory researchers are advised to choose samples in a way that maximises access to the phenomenon under study by selecting selecting most evident cases (Glaser, 1992, 1998). Informants chosen for study must be expert participants with rich and extensive prior experience with the phenomenon in order to be able to provide the researcher with a valid account of their experiences. For these reasons, we needed to adhere to three criteria in locating our study:

Finding the right tool. Firstly, we needed to choose a social software tool that provides some form of privacy management, preferably at an advanced level. After an extensive review of the existing tools, we chose Elgg (Tosh and Werdmuller, 2004), an open source social software system with integrated blog, wiki, social bookmarking, and social networking functionality. The two key features that motivated our choice of environment were its support for the creation of *ad hoc* groups and communities where privacy issues potentially arise, and its strong emphasis on its permission architecture, which had resulted in reasonable support for privacy control at a fairly granular level that other tools simply did not have.

Finding the right users. Secondly, we needed to find a situation where the tool was used extensively, preferably over a long period of time, so that users were properly familiar with it and were not novice users. While exploring various options to identify such user community among the general user population of Elgg, we came across a community of high school students enrolled in a special program for gifted kids called Trans, who were using Elgg for over a year as a requirement for their curriculum. This provided us with a user community for our selected tool who were using it on an ongoing basis for a reasonably long period of time.

Finding the right context of use. Finally, we needed to locate a context of use where both concepts of information sharing with various groups and privacy were paramount. Students in the Trans community were required to fill in their personal profile, write reflections in their weblogs on the topics covered in the classroom on a daily basis, and join and participate in a special community created for their group. For each of these artefacts (weblog posts, profile items, and personal reflections posted to the community blog), they had the option of regulating

access (i.e. make it visible to only oneself, the instructor, a specific community, or everyone). Since active use of the environment was part of the their curriculum, these students had in fact a rich experience in using various features of the tools, which was an essential requirement for the emergence of the issue of privacy preferences and selective disclosure of information.

Confirmation of the suitability of the context of use was the final step in the process of locating this study. It must be noted, though, that even though the study is situated in the context of Elgg, constant effort has been made not to limit the discussions to the specifics of the application. Instead, we treated Elgg just as a focal point to ensure that the subjects had the experience with a system that allowed them to manage their privacy directly.

3.3. Data collection

Our initial set of participants included nine students from the Trans community. The participants' ages ranged from 15 to 17, and the gender balance was rather evenly split (five females and four males). All nine participants were quite confident with the tool and with the Web in general. We selected semi-structured, in-depth interviews as our data gathering strategy, suggested as one of the best fits with the grounded theory methodology (Glaser, 1998; Glaser and Strauss, 1967). Unlike structured interviews, semi-structured interviews have a flexible and dynamic style of questioning directed towards understanding the significance of experiences from the informants' perspectives. Our interview strategy involved asking open-ended questions to allow informants to discuss what would be important from their perspective. We then used both planned and unplanned probing to uncover details and specific descriptions of the informants' experiences. The interviews were structured around a list of topics based on the research questions that needed to be answered, including questions about sharing preferences with regard to the type of information, the person or group with whom the information was shared, and the purpose and incentive behind sharing or holding information. Although we started with the same set of questions with each informant, because of the open nature of semi-structured interviews, each interview took a different turn based on the specific ideas and experiences of the participant in question. When that happened, we followed the participant's lead, allowing for new and important issues to uncover. Each interview was between 30 and 40 min in length. All interviews were tape-recorded with the informants' permissions and later transcribed to provide accurate records for analysis. Standard procedures were followed to maintain the confidentiality of the interview data and the anonymity of the informants.

The analysis of the data gathered from our initial interviews with the nine participants resulted in identifying the basic social processes (BSPs), which are the core concepts around which the grounded theory is built. After identifying the BSPs, we used a procedure called theoretical sampling (Glaser, 1978) to develop new insights and refine the insights we had already gained. For this stage, we consciously selected three more participants from the same group (one female and two males) who had extensive experience with other social software applications in addition to Elgg (i.e. other ePortfolios, forums, and weblogs). We also redirected the interview questions in a way to reflect our new goal of verifying the emerging theoretical themes and their relationships. The experiences of these three participants particularly helped in identifying places where the current privacy mechanism was considered insufficient and users felt the need to switch to other platforms in order to achieve their goal.

After analysing the data from all 12 interviews, we realised we could identify interchangeable examples showing the same phenomenon in different instances, and there were not any new concepts and/or relationships being developed. This was an indicator of theoretical saturation (Glaser and Strauss, 1967), the point at which we ceased data collection.

3.4. Data analysis

Our grounded theory was formulated from data using a constant comparative method of analysis with three stages: the first stage of analysis, called open coding, involved breaking the interview transcripts down into discrete incidents (i.e. ideas, events, and actions) which were then closely examined and compared for significant concepts. These concepts were abstractions in the sense that they represented an aggregated account of many participants' story. We used the qualitative analysis software NVivo at this stage to label incidents in the data with code words and to write theoretical notes that captured momentary thoughts.

The second stage of analysis, called theoretical coding, involved taking the concepts that emerged during open coding and reassembling them with propositions about the relationships between those concepts. The relationships, like the concepts, emerged from the data through a process of constant comparison. Neither the concepts nor the relationships were preconceived or forced upon the data.

The third stage of analysis, called selective coding, involved delimiting coding to only those concepts and relationships that related to the core explanatory concept reflecting the main theme of the study. At the end of this stage, we were able to produce a more focussed theory with a smaller set of high-level concepts.

4. The grounded theory

The concept map in Figure 1 represents the theory that emerged from the data in this study. The concepts in

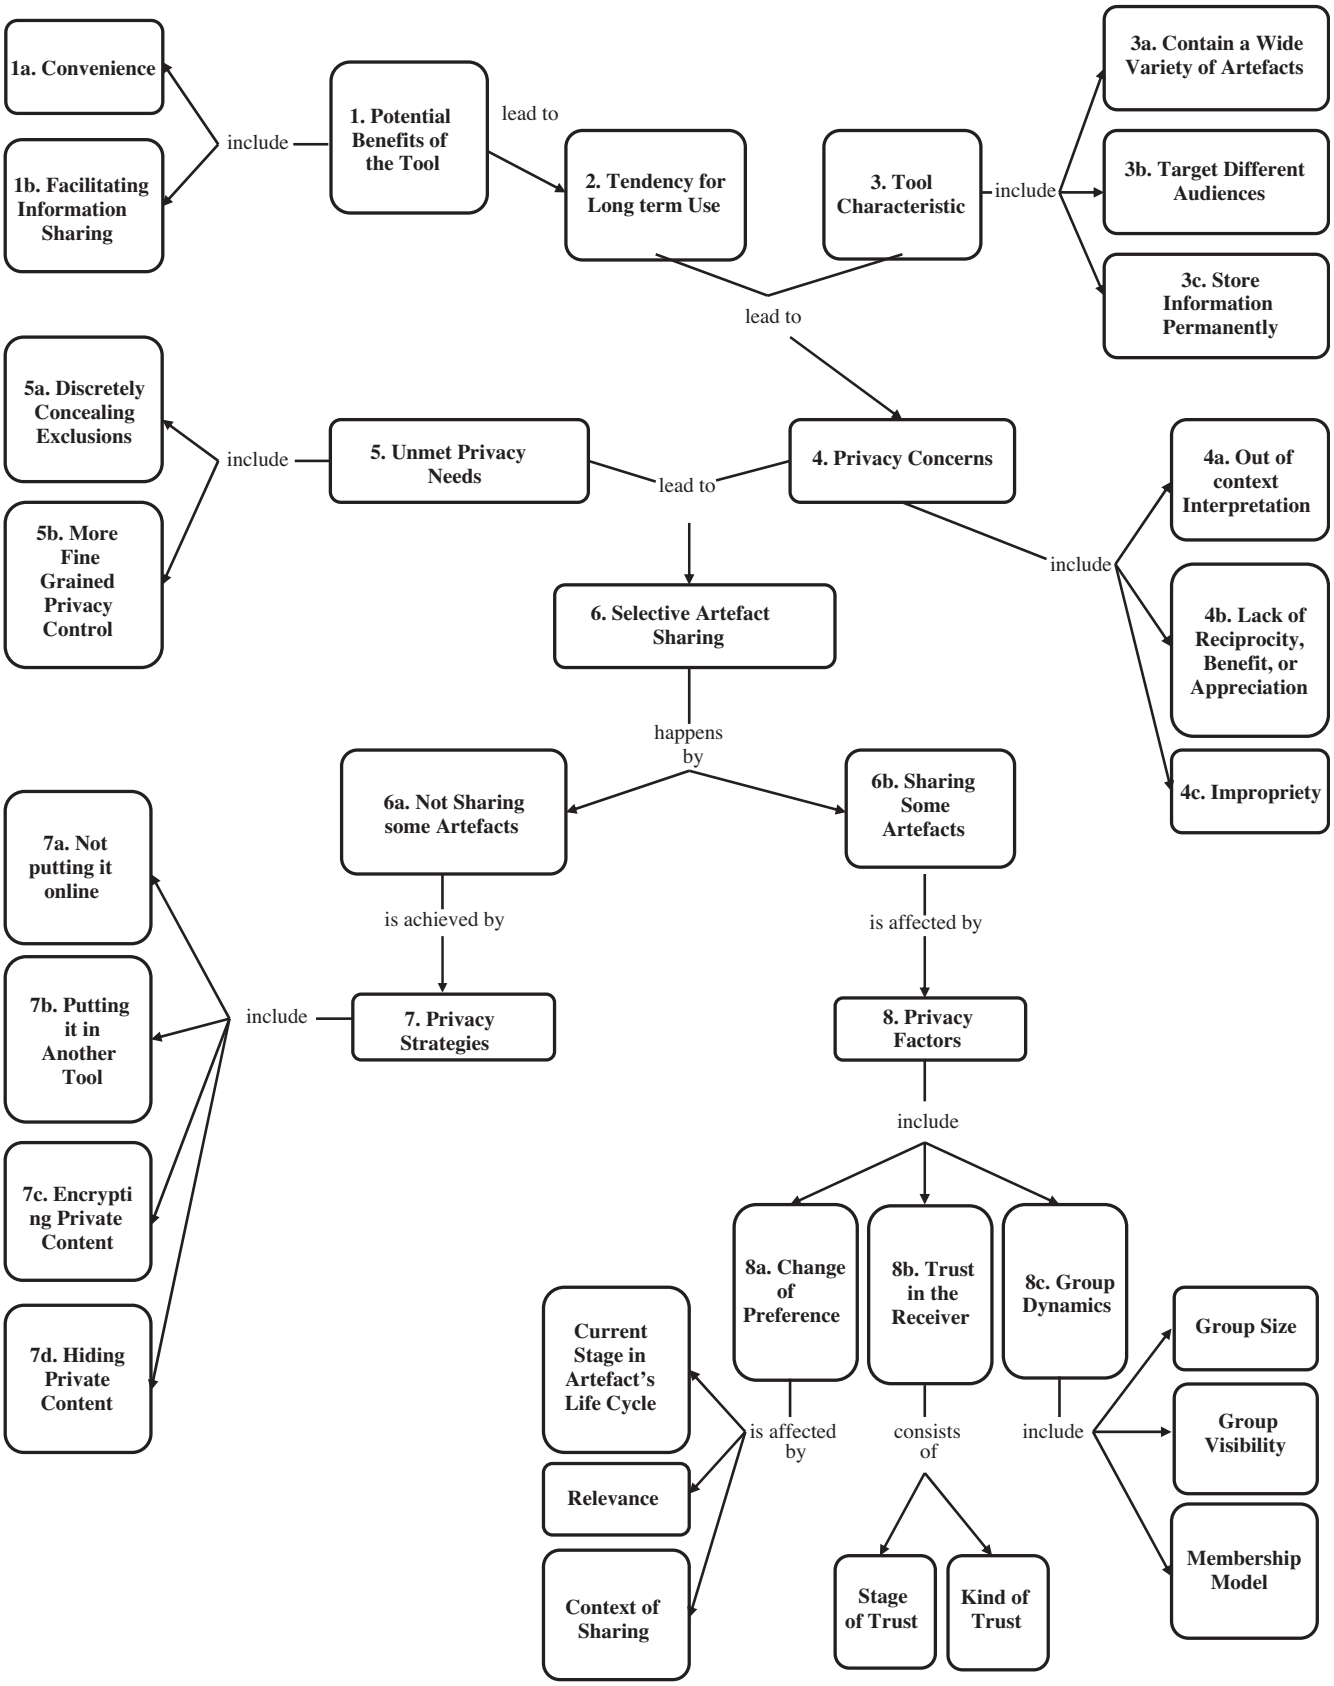

Figure 1. A concept map of information sharing behaviour in a social software system.

the map have been marked by numbers so that they can be referred to in the following description by their corresponding numbers (used in brackets), thus making navigation of the map easier. In summary, the theory suggested that because of the many potential benefits of the tool [1], such as convenience [1a] and ease of information sharing [1b], participants were willing to use it over a long period of time. However, there were certain inherent characteristics of the tool [3] that in combination with users' willingness for long-term use [2] gave rise to privacy concerns [4]. First, the artefacts contained in social tools cover a wide spectrum, ranging from personal to professional to social information [3a]. Second, various information artefacts may be targeted to different groups of audiences that are not necessarily static [3b], and finally, information disposed in social tools is persistent and permanently searchable [3c]. Some of the privacy concerns mentioned by our participants included out-of-context interpretation [4a]; lack of reciprocity, benefit, or appreciation [4b]; and impropriety (considering certain content inappropriate for some audiences) [4c]. We also observed that there were certain privacy needs of users that the tool failed to support [5], including the need to control privacy at a more fine-grained level [5b] and the need to discretely conceal exclusions [5a].

The combination of privacy needs and concerns gave rise to the need for selective information sharing [6], which happens by sharing certain artefacts with selected lead to individuals [6b] while holding back form sharing others [6a]. To withhold from sharing, users were employing certain strategies [7], including using other platforms with better privacy management mechanisms for their more private content [7b]; refraining from deploying certain content in an online environment because of lack of acceptable privacy levels [7a]; putting the more sensitive content somewhere (i.e. a web page or weblog), but not providing a link to it from places where their real identity was known [7d]; and finally, writing their more private content in some sort of a 'code language' so that it is meaningless to anyone other than the user him/herself [7c].

Deciding on sharing an artefact on the other hand was affected by a set of privacy factors [8]. One such factor was change of preference [8a]. The study showed that rather than a binary scale of public vs. private, users' judgements of privacy of resources often reflected a transition from private, to semi-private/restricted share, to public, depending on the state of the artefact, the relationship between the owner and the receiver of artefact, and the context of sharing. We also found out that users' assessment of the persons or groups who will be the receivers of information played a strong role in making decisions about information sharing [8b]: users tended to share less with people/groups with whom they were in the initial stages of trust, and as their trust moved towards a more mature level over time, they began to feel more comfortable and share more. Furthermore, their decision was affected by both the kind of trust they had in the receiver (e.g. cognitive or emotional) and the stage of their trust with the receiver (e.g. initial, intermediate, or mature).

And finally, our results indicated that users' willingness to share something they have vested interest in also depends on their perception of how it will be used, with the dynamics of the groups or communities where the information is going to be shared being the most influential factor in deciding about information sharing [8c]. Our study revealed that users often hold back from sharing information in anticipation of lack of reciprocity, benefit, or appreciation, and loss of credit for their work. The theory suggested that when group/community dynamics are clear enough to convey to the users how their information will be used within the group, users are better equipped to make informed decisions regarding how much they want to share within the group. Moreover, this predictability may be critical to making the decision to share information in the given context at all. Tables 1–8 provide examples of users' comments that led to various

Table 1. Sample user quotes on the benefits of Elgg and tendency for long-term use.

P#	Example comment
4	It is very useful to have everything in my Elgg, because then other people can see what work samples I have. Like say, if I am applying for scholarships, then I want them to have samples of my work without me having to send them things, so they can access the site whenever they want to. It's like I have an online resume.
1	There are two benefits to using Elgg; first, it has good integration potential, meaning, I can access my course materials all in one house. I can even present from where my data is. Second, it makes it easy to connect with classmates because they are all on Elgg, too. So, if I want to, say, organise an event or something, I know they will all see it.
3	For me, the main use of Elgg would be after Trans: it is my ePortfolio; it will be part of the application package I send out to schools. That's the reason I am using it in the first place. Besides, all my Trans friends are using Elgg now; so we can stay in touch if we continue using it after Trans.

Table 2. Sample user quotes on privacy concerns.

P#	Example comment
1	. . .I don't always share information on scholarships; because they are highly competitive, by sharing I would just put myself in a less advantaged position. I know other people don't share such stuff, either; so. . .
12	I usually prefer people not to know that I am coming to this program because that sort of affects the way that people think about me. By keeping my educational and social information from certain people who really don't know a lot about me, I am treated more like an equal.
5	. . .even though that is an important part of my identity [referring to a certain interest], I just decided to take that off my ePortfolio, because although I don't mind my fellow Trans students know that, I don't want to find people who don't know me think I am weird. . .

Table 3. Sample user quotes on privacy needs.

P#	Concept	Example comment
8	Life cycle	Sometimes, I would put it [referring to samples of creative writing] on private because it has too much information about me that I don't want sharing over the Internet, or sometimes it has more private things 'to me' to go public. But then sometimes they become 'outdated' or I need to put them up as samples for assignments, or examples for a question.
3	Life cycle	My reflections are usually private, but for example, for [a particular course], we need to write down our reflections so that [the instructor] could see what we took out of the sessions. That's when I need to move something from private to public: it's because I need [the instructor]'s comments on it.
2	Relevance	My critical reflections are public for now; but I will change them to private when I want to provide my Elgg for scholarships. They don't need to see all my critical reflections. Not that they are self put-down or anything; in fact, criticism is the way we make progress, right? it will just be irrelevant for the purpose.
12	context	We have created a group for our [a course] group project in the past. There was this [...] assignment that we had and everyone needed to contribute by writing in the journal. So we uploaded the file into Elgg file repository and initially, gave access to it to only the group. Then when it was done, we also let [the instructor] see it, like we added her to the friends in the group. She was quite happy with the work, so she suggested we make it public so that others can see it, too.
1	Context	I move things between private and public in my ePortfolio, which is mostly schoolwork. For example, say we have a lab assignment due on Thursday; I would post it up for me to look at in the private one, just to check that everything is completed before I submit. Only after the due date I post it in the public one, because of copying.

Table 4. Sample user quotes on privacy factors—change of preference.

P#	Example comment
9	The problem I have with that [current privacy mechanism in Elgg] is that when I let some people see something, other people can see that there is something, but they don't have access to that. So they are like: oh, can I look at it? and then sometimes, you just don't know whether you want to share with them or not, and it's kind of weird to say no right away. So, then sometimes, I just rather keep it all private or all public so not to have to make that decision.
8	What would have been nice to have, is for people who don't have access to it to see a blank page instead of a message like, sorry, you don't have access to this.
11	I would rather keep my reflections private, but for example, for [a particular course], we need to write down our reflections so that [the instructor] could see what we took out of the sessions. Then I need to move my reflections from private to public.

theory constructs. A more detailed description of the grounded theory process can be found in Razavi (2009).

5. Study limitations

The unified demographics of our particular sample might raise a question as to whether the results were affected by the specific characteristics of this group. Particularly, an argument can be made as to whether the fact that this group of participants was under age 18 had any effect on the results. While there were certain categories of artefacts that our participants would not share because of their age (e.g. none of them were allowed to post a real photograph of themselves on their profile), we did not

Table 5. Sample user quotes on privacy factors—trust.

P#	Concept	Example comment
11	Kind of trust	Right now I am on a forum and I remember in the beginning I was really careful about exposing personal information, such as where I go to school or posting a picture. I would just ignore and leave myself out of it. After a while, you sort of trust them a bit more. I haven't been as far as putting a picture on, but I would say oh, I would get my license in a couple of years or something like that. But I won't make a reference to the fact that I am not old enough—I would just say I will get it in a couple of years. So, I am still pretty cautious about it; because after all, my trust just comes from interacting with these people over time. I mean, I just 'feel' more comfortable after being in the group for a while.
2	Stage of trust	After interacting in a group for a while I would feel more comfortable sharing with the group but its not always very comfortable, just more comfortable than before; Like, from 'not very comfortable' to 'sort of comfortable'. I am not the kind of person who gets too comfortable over the net.
6	Stage of trust	If you participate in an online community and you talk to people and they begin to give their opinions about something, you feel you begin to know who that person is by what they say are their ideas and what they like, and you develop a sense of knowing who they are, and they are no longer unknown; because we fear what we don't know and so if we get to know what that person stands for, maybe we can trust them some more.
10	Kind of trust	I am not the kind of person who makes friends over the Internet easily and I don't really connect with forums well; but once that happened, though, I actually had my friend who had visited the forum for a long time. So, it was easier to connect because I had a really strong connection there.

Table 6. Sample user quotes on privacy factors—group dynamics.

P#	Concept	Example comment
11	Size	[What I share in a community] also depends on the size of the community. Because some communities are really popular; there are lots of people; so you can't really get to know everyone. I am usually more comfortable when it is small, like say ten people. That's a bit more personal, and I get better credit for my contributions.
4	Size	I once created a community for [...], which was a closed community. My experience with that community was actually very positive: everyone would contribute actively and give others feedback on their work. But then, we all sort of knew each other, so it was more like chatting with friends... It was a small community, though.
2	Membership model	The problem with anonymous communities [where providing real information is not a requirement for membership] is that you have no way of knowing who the comment is coming from... you can't trust them with their judgment: it could be a grade one kid or it could be a Ph.D. so it's not worth anything.
3	Membership model	[What I share in a particular community] would really depend on who else is in there. In Trans [the particular community they have for all Transition students] I know the students [who the community consists of], so I would share my opinion on certain things that I wouldn't mind sharing with them in person; but for some stuff, I would definitely not share.
5	Visibility	To me there is a strong distinction between private and public groups. Private groups are invitation only, so I would appear with my real name and share practically everything. The public ones are open to everyone though; so I usually use a pseudo name and I am cautious not to reveal any personal information.
12	Visibility	[When sharing stuff in a community] I'd like to know what they are doing with it, but they don't have to tell me. I mean I am offering it, so they can use it if they want to. If they want to tell me what they are doing with it, I would like to know that, too. I don't mind as long as they give me credit for it.

focus on this group of artefacts. Rather, we tried to stay away from these obvious cases, and focus on the more general area of sharing information they had vested interest in. Also, although this study is situated in the context of an educational environment, our participants used it for managing and sharing many more varieties of personal and social information in different contexts. As such, their experiences reflected diverse information sharing habits in various contexts, as evidenced by the fact that none of the privacy factors that emerged in the study (the changing nature of users' privacy preferences, the effect of trust on information sharing decision, and sharing differently in group of different dynamics) were specific to the educational context. As such, we believe that even though our study has well-defined boundaries in terms of the user population, types of information artefacts, the intended

Table 7. Sample user quotes on privacy strategies.

P#	Example comment
11	[What I share in a community] also depends on the size of the community. Because some communities are really popular; there are lots of people; so you can't really get to know everyone. I am usually more comfortable when it is small, like say ten people. That's a bit more personal, and I get better credit for my contributions.
11	Besides Elgg, I have two other ePortfolios, and a couple of weblogs. One is private and one is public. On the private ePortfolio, I have things that are actually more private, like it has information about me, that sort of stuff. The purpose of that is that I just want to write some stuff down, so that it is sort of 'said' somewhere. Sometimes I don't want to keep stuff in my mind, like for example, a journal or something, I would put it on the private one.
10	I use [another platform] for more private stuff because there are settings for public or friends-only or you get to choose who gets to see it. If it is something you want the teacher to see but not anyone else, you can just set it that way.
9	[my private blog] is open, but it's sort of hidden, it's not obvious how to find the page. I have not provided a link to it from anywhere. So, it's open, but it's sort of hard to find.
12	I use LiveJournal for stuff that I want to share only with my closest friends; for things that I consider really private, however, I wouldn't write them down anywhere online; because the easiest secret to keep is the one that is never told.

audiences, and the context of use, it does provide meaningful insights into users' privacy needs and strategies in the social software system domain.

6. From grounded theory to design Heuristics

The second phase of the work involves translating the social requirements (which were identified through the study) into technical requirements (which are the actual technical structure of a privacy system to support those requirements). The main objective of the grounded theory study was to improve understanding of information sharing phenomenon, in order to identify guidelines that can inform design. To achieve such goal, we next developed a set of heuristics for the design of privacy management mechanism based on the results of the grounded theory, and then consolidated these heuristics into a privacy framework for social software. Here is a description of these heuristics.

Heuristic 1: Privacy control must be available on a fine-grained basis. The first heuristic suggests that control of the privacy of information must be defined in terms of individual artefacts as well as their collections, and is based on the observation that many of our participants in the grounded theory study expressed the need for fine-grained privacy control. Although this is a confirmation of the long-standing model that access rights should be associated with individual objects (e.g. files) and collections (e.g. folders), the higher granularity and incremental object creation model in social software suggests that the way in which these rights are managed to protect privacy and facilitate sharing needs to be different in some essential ways: the diverse nature of content and audiences in the social software domain implies that different artefacts in the same category might have different privacy requirements (i.e. landscape photographs made visible to public, but family photographs restricted to friends). Moreover, often times users need to grant or deny access rights other than just the read action to their artefacts (i.e. colleagues may view, but not modify), which suggests that social software systems need to support fine-grained privacy management not only for resources, but also for target audiences and actions.

Heuristic 2: Privacy preferences must be defined in context. Research shows that while non-technical users seem to have a good idea of what their personal privacy preferences are, often times they have difficulty articulating them in terms of a set of rules (Egelman and Kumaraguru, 2005). Personal preferences are also context-sensitive, which makes it even harder to enumerate specific privacy rules. Enabling privacy preferences at a fine granular level makes this problem even bigger. This is supported by the fact that while the need for managing privacy at a fine-grained level has been recognised by other social software systems as well, it has often found to pose a trade-off with usability. In Facebook, for example, users have to go to a separate privacy setting page and set privacy preferences for each of their various profile and public search visibility items individually. Privacy-related options for individual applications are found with the application and users have to be aware of the features to find the options and visit separate privacy pages for each. Although fine-grained, the result is a completely unintuitive system where non-technical users are highly unlikely to be able to set sensible privacy preferences or understand the ramifications of their choices. Interestingly, all this effort is needed for just regulating visibility of one's various artefacts (i.e. the read action). Facebook currently does not provide any mechanism for regulating other types of action; for example, who can edit an artefact or leave a comment, etc. To the best of our knowledge, neither do any of other existing heavily used social software systems.

Thus, our second heuristic is a direct follow-up to the previous one and suggests that a privacy management

approach that requires users to indicate their privacy preferences to the system *a priori* (i.e. through a privacy setting page) may not work; rather, we propose that any attempt to support fine-grained privacy management must be paired with enabling users to express their privacy preferences in context (e.g. at the time an artefact is created or modified) when they have a better idea of whom they want to share the artefact with.

Heuristic 3: Privacy mechanisms must provide control over ownership. Another deduction that followed from our grounded theory study was that users have a fundamental assumption that when they put something in the tool, they should have control over its ownership as well as its visibility. Our study suggested that one reason behind users' reluctance to share information was the tool's inflexibility in providing them with the ability to control the transactional aspects of knowledge sharing activities (e.g. getting proper credit for their contribution or ensuring reciprocity). This heuristic suggests that in order for social software systems to properly support information sharing needs, they must provide a complete, persistent sense of the degree to which information that an individual creates or consumes is his/her own, the amount of control s/he has over the use of that information, and the ability to properly assess or exploit its value. In other words, in addition to providing the means for users to control access rights at different degrees between the extremes of private and public, tools need to also allow users to maintain personal ownership control over their shared information.

Heuristic 4: Privacy mechanisms must support various group models. From a user's point of view, the primary concern in managing information sharing is in the ability to define and/or understand the audience that will have access to a particular information artefact. Generally, the choice of audience for a particular artefact or personal attribute is expressed in terms of a group of others who one trusts with that particular piece of information, so we suggest that a privacy mechanism must enable users to understandably model their trust into groups in a flexible and dynamic way. We therefore propose that group management in social software must support various group models rather than a generic unified form, and that groups must be defined and controlled by users, rather than the system.

Heuristic 5: Privacy mechanisms must provide control and/or awareness over group dynamics. Privacy in social software is also affected by the semantics of social network relations. For example, membership in a group with public membership visibility may thereby disclose interests, preferences, or other personal information regarding group members. This means that if a group member discloses information about him or groups including himself, he (whether willingly or inadvertently) might also be disclosing information about someone else. In other words, one member's treatment of his/her privacy has a direct effect on another member's privacy. This suggests that awareness of group dynamics is an essential need for a privacy management system; meaning, such dynamics must be both controlled by and clearly articulated to users.

Heuristic 6: Privacy mechanisms must allow definition of groups that reflect interpersonal relationships. This is a follow-up to the previous two heuristics, and suggests that one group model that must be supported in social software is the egocentric group that is defined based on users' interpersonal relationships. This follows from the observation that in social software domain, one's personal and social information are not always shared with identifiable, accountable individuals or groups, and sharing may happen in a variety of contexts, for example, competitive as well as collaborative. Moreover, people may act simultaneously in several contexts, holding multiple potentially conflicting relationships simultaneously. As such, a lot of users' information sharing needs is better described in terms of the relationship that exists between the owner of the artefact and the person or group with whom the information may be shared, specially since new intricacies have blurred the boundaries between public and private. Boyd (2006), for example, points out that US teenagers feel strongly about preserving a certain form of privacy: they want to be visible and searchable for their friends but not their parents. In terms of rights management, these observations strongly imply that the potential audience for some artefacts or attributes is likely defined in user's own terms, based on a variety of kinds of relationships that more closely resemble real-life privacy boundaries (e.g. one-sided and short-term relationships) and not in terms of any organisational 'roles' or groups. This will enable users to control the release of their personal information in the same manner they would control it in the real world, based on their relationship with the data receiver, rather than some externally imposed constraint such as the receiver's organisational role.

Heuristic 7: Privacy mechanisms must easily accommodate changes in preferences. The next heuristic is based on the dynamic nature of users' privacy preferences in the social software domain. While in general, any act of information sharing can be defined as 'a user sharing an artefact with a receiver based on their relationship', in the social software domain the information sharing act is often about establishing and maintaining a dynamic sharing relationship: over time, the nature and state of personal artefacts might change (i.e. research results getting published, patented ideas getting approval, personal opinions reconsidered), the receiving group with whom the information is shared might change (i.e. competitors joining a group or collaborators leaving), and the relationship between the owner and the receivers of information might change (i.e. people

switching to a different project groups or changing affiliations). We thus propose that a privacy model that statically assigns access rights based on these factors at the time of an artefact's creation or modification will be insufficient. Rather, privacy mechanisms need to be flexible enough to accommodate frequent changes in users' privacy preferences in a non-labour-intensive way.

Heuristic 8: Action possibilities and their consequences must be clearly presented to users. Our last design heuristic emphasises the importance of interface clarity. Our study confirmed the intuition that users can be reluctant to share personal information when they are not sure how exactly to do things, or when the consequences of a sharing decision are unclear. A counterintuitive consequence of this is that some users might be more willing to share personal information in a space that affords virtually no privacy control (e.g. blogs or Myspace pages) than one which offers them an unclear set of privacy management tools. In our study, users were made aware that they could have some control of privacy and should manage the audience for their personal information by the promise of an access control system. However, many found it inadequate because either they could not perceived how to do something they wanted to do (i.e. users did not know they could make something visible only to one person, even though such functionality was supported by the tool), or they were not sure what the consequences of a sharing decision were (i.e. even though the tool provided different information sharing models through supporting both groups and communities, users were not clear on how they differed). As a result, they were not able to take advantage of certain aspects of the privacy management mechanism, because of the inability of the tool to convey their existence or consequences.

7. A framework for privacy in social software systems

The overall goal of the proposed design heuristics was to identify a minimal set of requirements (both technical and social) for privacy management in social software systems. While heuristics Heuristic 1 to Heuristic 7 describe what kinds of privacy control are necessary for managing and sharing personal artefacts, Heuristic 8 pertains to how these controls must be built and incorporated in order to be usable. We now consolidate these requirements into a framework for user-centred privacy in social software domain that describes privacy in terms of required controls over artefacts, audiences, relationships, and change; with an emphasis on the clear presentation of those controls to users.

Artefact control. The principle of artefact control is a consolidation of Heuristic 1, Heuristic 2, and Heuristic 3, and essentially reflects the finding that privacy management in social software must be defined on a per-artefact level as opposed to per-category; and that the access rights need to be applied in context (meaning, at the time of artefact creation or modification) as opposed to *a priori* (through a privacy setting page). Furthermore, in addition to control over visibility (the read action), users also need the ability to control other rights over their artefacts, for example, modification or deletion (the write action), and over further delegation of such rights.

Audience control. The principle of audience control is a consolidation of Heuristic 3, Heuristic 4, and Heuristic 5, and reflects the need to restrict both the visibility and ownership of artefacts to certain user-defined groups. Although most existing social software systems support some group functionality, we suggest that social software systems must provide the means not only for creation of these user-defined groups, but also for definition and control over various aspects of these groups (sizes, membership models, and visibility), and for controlling changes in those aspects. Furthermore, these controls need to be in the hands of users, rather than pre-defined by the system.

Relationship control. The principle of relationship control is a reiteration of Heuristic 6, and reflects the need for the ability to define information sharing based on a user's self-defined relationships with others. In essence, this emphasises that users need the ability to define groups of friends or collaborators in their own terms, and to use this model of their relationships with others as the basis for audience control.

Change control. The principle of change control is a reiteration of Heuristic 7 and is something of a cross-cutting concern within the other three controls. This principle reflects the observation that in the social software domain, the artefacts, the audiences, and the relationships used to define privacy and sharing patterns are all dynamic. A privacy and user interaction model must thus take into account that artefact life cycle and categorisations will change, that a user's requirements to share classes of artefacts with certain audiences will change, and that a user's relationships and trust patterns within those relationships will change, and that users come to expect their tools to provide flexible support for these changes in their privacy preferences when the social parameters that define the sharing model change.

Clarity. While the other heuristics focus on what kinds of control of privacy are needed, the last one focusses on how those controls must be presented to users in order to be usable. As such, it must be considered in parallel with the other four controls as presented in Figure 2. We use the term clarity to represent this heuristic; meaning any functionality to incorporate artefact, audience, relationship, or change control must be designed in a clear and understandable way; to ensure that in practice, the average, non-technical users would be able to take

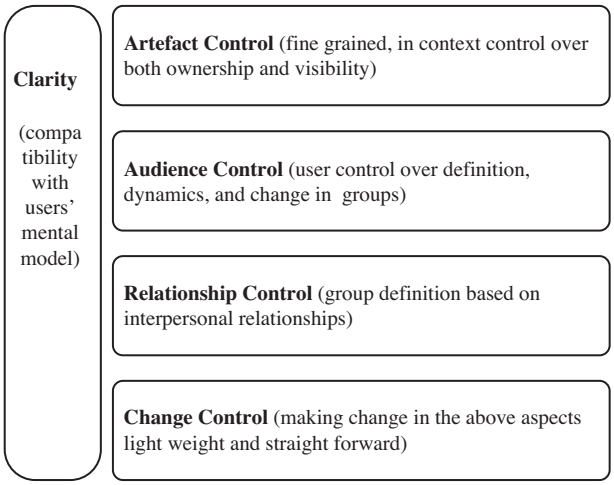

Figure 2. A framework for usable privacy control in social software.

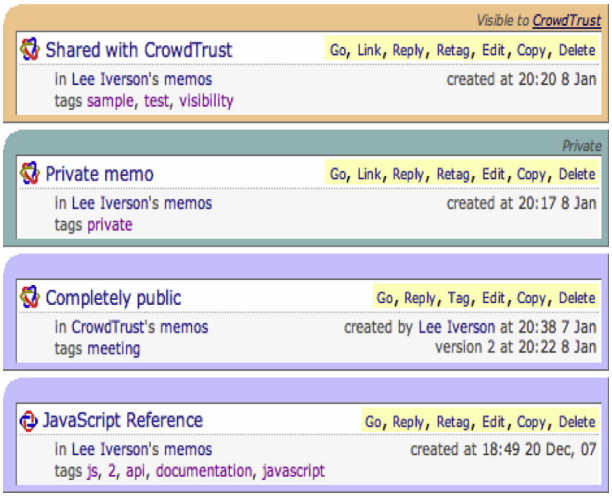

Figure 3. Memos in OpnTag: public, private, and selectively shared in a group.

advantage of the extra control over privacy that these user-centred controls are supposed to provide.

In order to illustrate the suitability of our proposed privacy framework for design, our next step was to design a system that instantiates the framework. As an example, and to provide an environment in which we can test these principles, we have developed an experimental system called OpnTag. We next present the technical structure of OpnTag, along with a discussion of how our framework as embodied in this system supports each of the five user-oriented privacy controls.

8. OpnTag

OpnTag (Iverson *et al.*, 2008) is an open source web application for note taking and bookmarking that we developed to address the information management needs of an individual performing in various social contexts. The fundamental unit of information storage in OpnTag is the memo, a tagged textual annotation that may optionally link to a web resource. Memos can function as bookmarks, notes, or wiki pages and are organised based on their intrinsic metadata (e.g. who owns or created them and when) and the tags applied to them by various users (Figure 3). Each memo has an owner, which presents who creates the memo and thus can edit and delete it, and a potentially restricted audience, which controls who can see that the memo exists and read it. Both the owner and the audience can either be an individual or defined as a group (Figure 4). Also, groups can be defined either by inviting other individuals or by applying people tags to individuals in one's network, thus categorising them into a group. Two types of groups exist in OpnTag: classic groups (Figure 5), with collectively-controlled visibility, size, and membership model; and egocentric groups (Figure 6), created through applying people tags to other

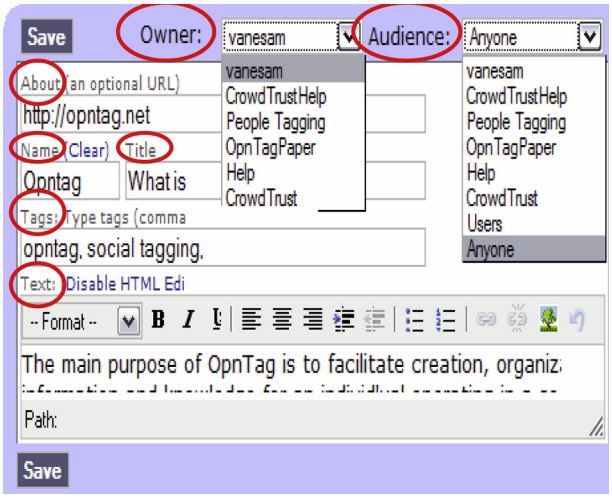

Figure 4. An individual memo in OpnTag with owner and audience list including both the individual and her groups.

individuals in the system, thus resulting groups with completely owner-controlled dynamics. For a more detailed description of the group functionality in OpnTag, see Razavi and Iverson (2008, 2009).

With individuals, classic groups, and egocentric groups, OpnTag's privacy control centres on the joint concepts of ownership and audience management, which ensure individual users retain control and credit over the artefacts they dispose in the system. Although a group can be specified as the designated owner of a memo, each memo is also visibly attributed to its individual creator, thus ensuring that each group member gets proper credit for the contributions s/he makes to the group's shared information space. Only a memo's creator can modify ownership, but any member of the owning group can change a memo's audience. This design choice allows the creator

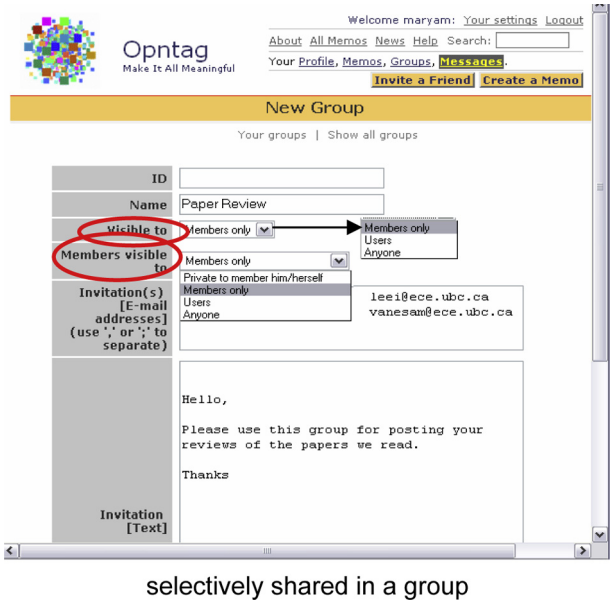

Figure 5. Classic group definition page with group and member list visibility menus.

Figure 6. Egocentric group definition in OpnTag through applying tags to users.

of a memo to decide whether access restrictions of a memo should be controlled collectively or individually. The latter case supports situations when there is a need to ensure that the access restriction policies stay with the shared data; i.e. it is not possible to make a memo visible beyond the intended audience set by its creator.

Audience restriction is the fundamental mechanism for selective sharing in OpnTag: at the time of creating or editing a memo, the creator has access to both his/her classic and egocentric groups and can thus adjust the audience of the memo to be either the owner himself (whether an individual or a group), a classic group that

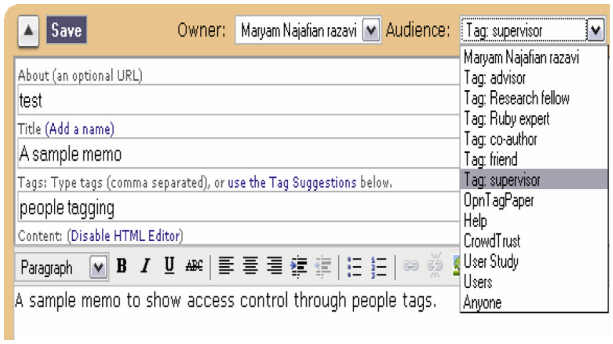

Figure 7. Selective sharing of a memo in OpnTag, with both classic and egocentric groups.

the owner is a member of (with its collective membership dynamics), or one of the owner's egocentric groups (over which s/he has complete control). Figure 7 presents a screenshot of how this choice is made.

9. Satisfying privacy requirements

OpnTag's privacy management mechanism has been designed based on the five user-centred privacy controls. Artefact control in OpnTag is supported by providing ownership and audience management at the level of individual memos. For each memo, the user specifies the memo's owner (who can edit it), with the default owner of a memo set to the creator (either the individual or a group). By enabling context-sensitive ownership management at the memo level, OpnTag's privacy system supports definition of fine-grained privacy policies any time an artefact is created or modified, as opposed to providing a separate privacy page for defining general privacy policies on collections as most current tools do. In addition to fulfilling the fine-grained privacy management requirement, this also allows users to define their own privacy policies when they have a clear idea of their preferences. This is quite important in terms of adhering to the principle of artefact control, as the principle emphasises pairing fine-grained privacy control with in-context policy definition to alleviate difficulties that users often have with administrating such fine-grained privacy control. Moreover, with OpnTag's ownership management, privacy policies stay with the data, meaning, no one other than the memo's owner can change its audience; i.e. it is not possible for someone who is not an owner of a memo to make it visible beyond the intended audience set by its creator.

Audience control in OpnTag is supported through deep visibility management via user-defined groups and relationships. Every memo in Opntag has restricted visibility; meaning, it can only be seen by members of a designated group. OpnTag enables users to control various aspects of the classic groups they define, including size, visibility, and membership. The relationship dynamics of

such groups is then determined by variations of these parameters; i.e. open/limited number of members, public/private visibility, and open/moderated/closed membership. This clear model of group characteristics highlights the potential trade-offs between risks and benefits of information sharing in a particular group and gives users a high-level overview of the effects of their sharing decisions. Knowing how a particular resource could potentially be used in a particular group, users can then tune their sharing decisions accordingly (e.g. one would be less likely to share sensitive information with an 'open' group that anyone could join than in a 'closed' group where new members must be invited). In addition, through people tagging, users can take complete control of the audiences for their artefacts.

Relationship control in OpnTag is supported through people tagging, which enables categorising one's network into user-defined, egocentric groups that may represent various kinds of relationships. Since egocentric groups are controlled entirely by the creator (the acts of creating or deleting a tag on a user, and controlling the visibility of the tag are solely controlled by the tagger), the act of tagging a person via their profile page is equivalent to asserting their membership in a group whose membership is entirely under tagger's control. As such, people tagging provides a lightweight and flexible mechanism for handling volatile relationships that frequently show up and fade out in natural social environments, but are often hard to manage in online world.

The combination of artefact, audience, and relationship control in OpnTag allows fine categorisation of resources, audiences, and actions, provides advanced group functionality, and enables users to define privacy preferences based on their often changing relationships. The fourth principle, change control, is supported by ensuring that ownership and visibility management (modifying owner and/or audience of memos and tags), group management (creating, managing, joining, and leaving user groups), and people tagging are all handled in a flexible and straightforward way and that all the settings are modifiable at any time, making it easy for users to make frequent changes to their information space.

Finally, we have followed some of the known principles of usable design to achieve clarity. One such principle is to make privacy features (as secondary functionalities) highly visible and seamlessly available to users in the context of their primary actions (De Paula *et al.*, 2005). This is derived from the fact that privacy features often act as barriers to action, while usability principles aim to remove such barriers (Dourish *et al.*, 2004). Supporting proper visibility can thus help achieve the right balance between the two seemingly conflicting goals and ensure that privacy management features complement existing actions rather than inhibiting it. In OpnTag, relevant action possibilities for each privacy control are graphically clustered and presented together, while irrelevant or rarely used information are omitted in order to reduce clutter: owner and audience selection for a memo (OpnTag's privacy controls options for artefacts) are visibly presented at the top of the memo edition page, and are the only functionalities presented at this level, separated from other functionalities that deal with the content of the memo (Figure 3). Likewise, group visibility, member list visibility, and people tag visibility (OpnTag's privacy control options regarding audience and relationship) are presented in the same context that classic or egocentric groups are created or modified. While this makes these features visibly and seamlessly available to users in the context of defining information artefacts, audiences, or relationships, it still gives them the option to skip such configuration and accept the defaults, if they would rather focus on their primary task.

Another usability principle is consistency (Nielsen, 1992). In OpnTag, various choices for each privacy control feature (owner and audience for memos, group and member list visibility for classic groups, and visibility of people tags) are consistently presented by drop-down menus, as a learned convention that is successfully and frequently used by people at all skill and experience levels. Designers of privacy systems are also advised to use feedback mechanisms to help users understand the implications of their privacy decisions (Bellotti and Sellen, 1993; Lederer *et al.*, 2003). OpnTag supports this through the combination of providing visualisation and pairing configuration with action. For both memos and people tags, the choice of audience causes the background colour to change accordingly. Such immediate visualisation of visibility and the choice that triggered it will help users understand how to generate rules that reflect their preferences, or to notice when the results of their actions do not correspond with their intended goal.

10. Usability evaluation

The two main features that distinguish our privacy framework from existing models of privacy management in current social software systems are the introduction of per-artefact control (as opposed to per-category control supported by most current tools), and use of people tagging which enables creation of nuanced relationship groups (as opposed to equal-weight, reciprocal relationships supported by the network-of-friends model). The result of these two innovative features is the promise of a more fine-grained, flexible, and dynamic privacy management that provides users with more control over selective sharing of their artefacts. Having developed OpnTag as a test bed, we next administered a small-scale usability evaluation to provide empirical evidence that the overall framework as embedded in OpnTag yields a usable privacy management mechanism. In other words, while as the first iteration in the design/evaluation process the

focus of the evaluation study was mainly on finding out what aspects of the framework work and what aspects do not, an equally important objective was to ensure that the extra control over privacy is not to achieve at the price of sacrificing usability.

10.1. Participants

OpnTag target users are individuals operating in various personal, professional, and social group environments, meaning, our sample pool was practically general public and there were no salient characteristics to look for in participants to assert their suitability for participation in the study. As such, a participation request was distributed via email to mailing lists and social connections of existing OpnTag users. Ten people (six male, four female) who responded to the invitation were recruited for participation in the study. Our participants came from diverse backgrounds, including both technical and non-technical users, which was appropriate since OpnTag is designed for personal and social information management across a wide variety of usage contexts. All participants had some university education, with seven participants having a graduate degree and three an undergraduate degree.

10.2. Procedure

The procedure we employed for the purpose of evaluation was a laboratory study consisted of three stages: first, participants were asked to fill out a survey questionnaire consisting of 10 questions covering general areas of users' demographics (sex, age, and education level), profession, expertise with computers and the Internet, plus specific questions on their familiarity with popular social software systems and the privacy management features in them. The survey also included questions aimed at identifying users' privacy attitudes, i.e. what information they feel comfortable to reveal about themselves in social applications they use and whether they have ever experienced privacy violation problems in the past. All of our participants had some familiarity with social systems (i.e. at least using one other social system for a prolonged period in the past) and some concern for privacy (i.e. at least considering some artefacts private and some sharing uncomfortable). These two criteria ensured that they had an idea of the subject of the study and can relate to it, and that their answers were not affected by lack of knowledge or concern about the topic of investigation.

None of our participants had used OpnTag before. On average, participants reported using various social systems for 6 hours per week. None had any special expertise related to privacy; however, when reflecting on their experiences with privacy management in applications they used, seven participants said they find the feature useful and have used it at some point. Six said although they consider privacy management mechanisms necessary, they often find them too difficult and/or time consuming to use. Five said they do not trust social systems to put their private information there. Only two participants said they think social systems provide adequate privacy management. Seven out of 10 participants reported experiencing privacy violation at some point while using social systems, either as a result of their own action or others.

In the second stage, participants were asked to login to OpnTag as an imaginary persona and perform a set of pre-defined tasks, each involving creating memos of various degrees of sensitivity and sharing them with various people in the imaginary persona's social and/or professional network. For each task, participants needed to decide on the appropriate owner and audience for the memo they created, and on whether to assign their target audience to a classic or egocentric group. Participants were allowed to make any changes to the persona's information space that they felt necessary for the purpose of carrying out the tasks; including creating, modifying, or deleting groups and/or people tags, or changing the visibility of their groups/group members/people tags. The tasks were designed in a way to cover information sharing situations across a variety of privacy-sensitive contexts, ranging from inherently private, to semi-private, to public; and to require a mixture of visibility and ownership control.

In the third stage and upon completion of the tasks, we engaged each participant in a semi-structured interview. We used participants' actions during tasks as a starting point and tried to gather feedback on the reasons behind their actions and why they did things a certain way. We particularly looked for errors and signs of confusion and/or frustration, as well as comments on the strengths and weaknesses of the privacy management mechanism. This combination of methodologies allowed us to make detailed first-hand observations of how first-time users interacted with OpnTag's privacy management scheme and how they reflected on its utility and usability.

10.3. Results

Ease of use and effectiveness. We asked users to rate OpnTag's privacy management mechanism in terms of both ease of use and sense of privacy compared to the social software systems they were familiar with (Orkut, LinkedIn, and Facebook). We used a scale of 1 to 5 for rating, with 1 indicating the worst performance and 5 indicating the best. Ease of use was rated 4.2, with min. = 3 and max. = 5 (users thought there were just too many steps involved to navigate to a user profile for tagging). Users gave their perceived sense of information privacy an average rank of 4.0, with min. = 3 and max. = 5. Although not all of our participants took the optimum path for doing all scenarios, they were all able to navigate their way through the privacy management system to get the tasks done. From the total of 50 tasks that our participants performed

unassisted (10 participants each doing five tasks), we only witnessed five errors. Furthermore, all five errors were results of improper understanding of the task in question; for example, making a memo visible to team members rather than colleagues, as mentioned in task description.

Overall, the concept of setting the owner and the audience for access management seemed to be fairly understandable to users: none of the users showed any signs of confusion or frustration. Also, the majority of our participants seemed to grasp the difference between granting 'write' vs. 'read' access: nine out of 10 users correctly created a group memo for the tasks that involved some form of collective contribution. Since this distinction is not supported by most of the existing tools, it was encouraging to see users quickly picking up and using a fairly new feature.

Usability. Traditionally, usability of a software application is measured based on four quality components: speed, accuracy, success rate, and overall user satisfaction (Nielsen, 2001). However, researchers have often considered different criteria for measuring usability of a security or privacy mechanism. The reason is that privacy management is often a secondary goal in most systems, and therefore does not get the same consideration that many other aspects do (Egelman and Kumaraguru, 2005), which makes it difficult to set particular metrics for usability of privacy aspects (i.e. what exactly should be measured?). Whitten and Tygar (1998) were the first to propose a working definition of usability for security software based on the special characteristics of the usability problem for security, and to suggest several criteria for evaluating usability of a security system. A number of other researchers have also proposed similar and/or complementing guidelines for evaluating usability of security or privacy mechanisms (Chiasson *et al.*, 2006; Clark *et al.*, 2007; Cranor, 2005; Karat *et al.*, 2005). We found these criteria suitable for the purpose of our study. Here we reflect on the usability of OpnTag's privacy framework based on Whitten and Tygar's (1998) four usability criteria, plus the two complementary criteria suggested by Chiasson (2006).

(1) Users must be reliably made aware of the steps they have to take to perform a task
This is a restatement of the first guideline of (Whitten and Tygar, 1998) and suggests that the application must provide user with enough cues as to how to start the process for each task, and to identify the intermediate steps that are required to complete the task. In OpnTag, the acts of setting a memo's owner and audience are fairly straightforward, because those action possibilities are associated and presented with memo creation/modification functionality. Also, the fact that privacy management in OpnTag happens on a per-artefact basis made it easy for participants to figure out that the owner and audience are the two attributes of a memo that they need to set in order to share something with a certain audience.

(2) Users must be able to determine how to successfully perform the steps
Whitten and Tygar's second usability guideline suggests that once the user is made aware of what intermediary steps are necessary for each task, s/he must be able to figure out how to perform these steps. (Wharton *et al.*, 1994) suggest that users develop a mental model of how a system works, and that in order for users to be successful in performing the necessary steps required to complete a task, the model behind the system must match user's mental model.
In our study participants employed different privacy management strategies based on their privacy attitude and concerns (e.g. one participant would consider a memo private, while another one would make it public). Regardless of their privacy attitudes, our subjects were successful in achieving their desired level of privacy, properly disposing the created memo to the right audience. Moreover, OpnTags' privacy management system seemed to have a fast learning curve: after a short, initial training session, our subjects all seemed at ease with creating memos, defining or modifying groups and/or group members, tagging other users, and choosing the appropriate owner and audience for their memos.

(3) Users should not make dangerous errors from which they cannot recover
Since OpnTag is an information management system, the most dangerous error that can happen is exposing a memo to the wrong audience. However, in OpnTag memos are created in the current work space, meaning the default value for both a memo owner and audience is the user himself (if in user space) or the group (if in group's shared space). As such, even if users miss to set the right owner/audience, having rather conservative values as defaults helps decrease the chance of accidentally making a memo visible to a too large audience (e.g. public). Also, the different background colour-codes that reflect various levels of visibility for a memo provide a powerful visual cue to the user as to whether the memo is set to have the right visibility.

(4) User should know when they have completed a task
This is the first complementary criterion to Whitten and Tygar's proposed by (Chiasson *et al.*,

2006) (also mentioned by Cranor (2005)). This criterion suggests that one of the essential usability requirements is enabling users to tell when their task is completed, which implies that the feedback provided by the system to users during a task should be adequate to ensure they are aware of its successful completion.
We asked our subjects several questions in an attempt to gauge their perception of the appropriateness and adequacy of the feedback provided by the system (the owner name, colour-codes, and group name in the resulting memo). After each task, we asked the participant if they believed they performed the task correctly, and if yes, how could they tell if they have been successful in setting the appropriate privacy level for the artefact they created. Eight of our subjects mentioned the use of at least one of the feedback mechanisms to check their results, while the other two participants just assumed they did it right and had not paid attention to any of the feedback information. Overall, the consensus was that the combination of colour-codes and owner/audience in the memo list conveyed enough information to users to form an idea of the current privacy level of their various artefacts at a glance. Six participants mentioned the colour-code as the most useful feedback mechanism, probably because of its visualisation power. Interestingly though, our choice of colour-codes was not popular with the participants. One participant mentioned that blue (OpnTag's colour-code for public) and green (OpnTag's colour-code for private) are quite easy to be mistaken, and suggested we use other colours for the two cases that show the distinction more clearly. Another participant thought green is not the right colour-code for private, since it implies 'green light' in a way. This participant thought a more strong colour like red would be a better choice for the case.

(5) Users must be able to determine the current state of the system at all times
This is the second guideline form (Chiasson *et al.*, 2006), and suggests that it should be visible at a glance what is visible to whom. (Cranor, 2005) advocates the use of 'persistent indicators' that allow the user to see privacy information at a glance. OpnTag does this through its use of colour-codes.

(6) Users must be sufficiently comfortable with the interface to continue using it
This is the fourth principle of usable security of (Whitten and Tygar, 1998), and is an essential part of the principal of psychological acceptability quoted by (Bishop, 2005). After completing the tasks, we asked our subjects if they can relate to the scenarios and whether a tool like OpnTag with advanced privacy management features would be more likely, less likely, or just as likely, to be incorporated in their daily personal and social information management activities. Most participants (9 out of 10) said that although the scenarios do not resemble their information sharing practices exactly, they could think of similar scenarios in their day-to-day activities where similar selective information sharing activities would be useful or necessary. Eight out of 10 participants said that they would try using such a tool for information management and sharing, and six participants said they feel a strong need for a tool with these sort of privacy management in their work place. Table 8 summarises some of participants' comments regarding willingness to adopt the tool.

Table 8. Participants' comments on willingness to adopt the tool.

P#	Example comment
3	It would be nice to have something like this that limits the contribution to a certain group, even though it may be exposed beyond that group. For example, I would like the contributors to a discussion on board level design to be limited to: X [who is a board designer], Y [who is his boss], Z [my boss], and me. That's all people who are knowledgeable enough to contribute to this discussion, although the whole group may read it.
5	I strongly feel that something like this is needed in our workplace, even though personally I am against using something like this because of privacy reasons: I don't think one should put his ideas in a system on the Internet; unless he can control who would have access to it.
6	One thing I like about this [OpnTag] is the control; I like that it can act as an integrated environment; because you can separate your personal and work-related stuff. I think that's very important. The idea of having my desktop online sounds cool.

11. Study limitations

There are a few limitations of the study that must be taken into account when interpreting the results. Ideally, we would have liked to perform a field evaluation with real users using OpnTag and its privacy system in real-life information sharing situations. However, while field studies report on users in their natural environment doing real tasks and as such, can demonstrate feasibility and in-context usefulness, they are time consuming to conduct and require mass adoption. A comprehensive review of literature on privacy evaluation methodologies (e.g. Chiasson *et al.*, 2006; Cranor *et al.*, 2006; DeWitt and Kuljis, 2006; Hawkey and Inkpen, 2007; Iachello and Hong,

2007; James *et al.*, 2007) indicated that many times, privacy researchers opt for controlled laboratory studies, especially in the early stages of the work to examine the viability of an approach before proceeding to more fundamental implementation/evaluation schemes. In the absence of an active OpnTag user community and considering that our goal at this stage was also finding out whether our framework yields a viable approach for improving privacy management and to get feedback on its potential problems to refine the design, we also settled for a controlled laboratory study of a small sample.

While we believe that our laboratory study was appropriate at this initial stage for successfully gathering meaningful feedback from potential users, it also came with some limitations of its own. First of all, even though we tried to recruit participants of varying backgrounds to cover usage scenarios across a variety of contexts, our representative sample was a small one and our representative usage scenarios almost certainly did not cover all privacy concerns and contexts of use for all potential users (although such criticism can be applied to the vast majority of laboratory studies in this area). Different opinions and problems may well be expected for other types of users, which is best to be investigated in follow-up studies.

Second, although laboratory studies provide an appropriate situation for observing users' interactions with the system in a controlled fashion, there are inherent challenges associated with laboratory studies in the privacy domain. One such challenge is that because participants are not dealing with their own data, they might not be as motivated as in real life and as a result, not make the same effort to protect privacy of their data. Another issue is that methodologies for studying privacy may themselves be deemed too privacy-invasive, causing users to deviate from normal practice and/or to withhold revealing sensitive aspects. As a result, relying on self-reported attitudes and behaviour alone may not provide a valid view of normal practices. Many privacy studies suggest that there is a gap between users' stated privacy preferences and their real behaviour (Ackerman *et al.*, 1999; Jensen *et al.*, 2005; Spiekermann *et al.*, 2001); however, such gaps are hard to capture in a laboratory study and are best addressed in field studies.

12. Conclusion and future work

In this article, we presented a grounded theory study of information sharing behaviour in a social software system and suggested eight heuristics for designing privacy management mechanisms in social software domain based on the results of the study. The heuristics included hints on both the kinds of control of privacy that would be required in various dimensions, and how those controls must be presented to users to ensure usability. We then consolidated these heuristics into a framework for designing privacy management mechanisms for social software systems, and presented OpnTag, a social software system we developed as an instantiation of the design guidelines for privacy management in this domain. A preliminary empirical evaluation validated the viability of the proposed framework for building privacy management systems that while usable, provide users with more control over privacy (although it did not directly address its completeness or adequacy). The evaluation process also identified a number of areas where improvements might further increase usefulness and usability.

Since the study focussed on a particular implementation of the framework, we are unable to use it to validate the framework itself. Instead, at this point we can only claim to have demonstrated that exposing this extra expressibility and complexity in the way we have done so in OpnTag does not reduce the usability of the system, and does seem to resonate with our users' own assessment of their privacy needs. Further evaluation will be required (i.e. a longitudinal evaluation in the field) to validate the model in a natural usage environment.

References

Ackerman, M.S., Cranor, L.F. and Reagle, J. (1999) Privacy in e-commerce: examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce* (Denver, CO).

Bellotti, V. and Sellen, A. (1993) Design for privacy in ubiquitous computing environments. In *Proceedings of the Third Conference on European Conference on Computer-Supported Cooperative Work* (Milan, Italy).

Bishop, M. (2005) Psychological acceptability revisited. In Cranor, L. and Garfinkel, S. [eds.] *Security and Usability* (Sebastopol, CA: O'Reilly), 1–12.

Boyd, D. (2006) Identity production in a networked culture: why youth heart myspace. In *Proceedings of the Annual Meeting of the American Association for the Advancement of Science* (St. Louis, MO).

Chiasson, S., van Oorschot, P.C. and Biddle, R. (2006) A usability study and critique of two password managers. In *Proceedings of the 15th USENIX Security Symposium* (Canada: Vancouver).

Clark, J., van Oorschot, P.C. and Adams, C. (2007) Usability of anonymous web browsing: an examination of Tor interfaces and deployability. In *Proceedings of the 3rd Symposium on Usable Privacy and Security* (Pittsburgh, PA).

Cranor, L.F. (2005) Privacy policies and privacy preferences. In Cranor L. and Garfinkel, S. [eds.] *Security and Usability* (Sebastopol, CA: O'Reilly), 447–472.

Cranor, L.F., Guduru, P. and Arjula, M. (2006) User interfaces for privacy agents. *ACM Trans. Comput. Hum. Interact.* **13**(2): 135–178.

Darrah, C., English-Lueck, J. and Freeman, J. (2001) Familes and work: an ethnography of dual career families, available online at http://www2.sjsu.edu/depts/anthropology/svcp/SVCPslnr.html.

De Paula, R., Ding, X., Dourish, P., Nies, K., Pillet, B., Redmiles, D., Ren, J. *et al.* (2005) Two experiences designing for effective security. In *Proceedings of Symposium*

On Usable Privacy and Security (SOUPS) (Pittsburgh, PA), 25–34.

De Witt, A.J. and Kuljis, J. (2006). Aligning usability and security: a usability study of Polaris. In *Proceedings of the Second Symposium on Usable Privacy and Security* (Pittsburgh, PA).

Dourish, P., Grinter, R.E., Delgado de la Flor, J. and Joseph, M. (2004) Security in the wild: user strategies for managing security as an everyday, practical problem. *Pers. Ubiquitous Comput.* **8**: 391–401.

Egelman, S. and Kumaraguru, P. (2005) Report on DIMACS Workshop and Working Group Meeting on Usable Privacy and Security Software. Rutgers University, New Burnswick, NJ, available online at http://dimacs.rutgers.edu/Workshops/Tools/dimacsrpt.pdf.

Erickson, T. (2006). From PIM to GIM: personal information management in group contexts. *Commun. ACM* **49**(1): 74–75.

Glaser, B.G. (1978) *Theoretical Sensitivity* (Mill Valley, CA: Sociology Press).

Glaser, B.G. (1992) *Emergence vs. Forcing: Basics of Grounded Theory Analysis* (Mill Valley, CA: Sociology Press).

Glaser, B.G. (1998) *Doing Grounded Theory: Issues and Discussions* (Mill Valley, CA: Sociology Press).

Glaser, B.G. and Strauss, A. (1967). *The Discovery of Grounded Theory: Strategies for Qualitative Research* (Chicago, IL: Aldine).

Gross, R., Acquisti, A. and Heinz, H.J. III. (2005) Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*, 71–80.

Hawkey, K. and Inkpen, K.M. (2007) PrivateBits: managing visual privacy in web browsers. In *Proceedings of Graphics Interface 2007* (Montreal, Canada).

Iachello, G. and Hong, J. (2007) End-user privacy in human-computer interaction. *Found. Trends Hum. Comput. Interact.* **1**(1): 1–137.

Iverson, L., Najafian Razavi, M. and Mirzaee, V. (2008) Personal and social information management with OpnTag. In *Proceedings of ICEIS 2008* (Barcelona, Spain).

James, R., Kim, W.T., McDonald, A.M. and McGuire, R. (2007) A usability evaluation of a home monitoring system. In *Proceedings of the 3rd Symposium on Usable Privacy and Security* (Pittsburgh, PA).

Jensen, C., Potts, C. and Jensen, C. (2005) Privacy practices of Internet users: self-reports versus observed behavior. *Int. J. Hum. Comput. Stud.* **63**(1–2): 203–227.

Karat, Clare-Marie N., Brodie, Carolyn A. and Karat, J. (2005) Usability design and evaluation for privacy and security solutions. In Cranor, L. and Garfinkle, S. [eds.] *Designing Secure Systems that People Can Use* (Sebastopol, CA: O'Reilly).

Lederer, S., Mankoff, J. and Dey, A.K. (2003) Who wants to know what when? Privacy preference determinants in ubiquitous computing. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems* (Ft. Lauderdale, FL).

Millen, D.R., Feinberg, J. and Kerr, B. (2006) Dogear: social bookmarking in the enterprise. In *Proceedings of CHI 2006* (Montréal, Québec, Canada).

Nielsen, J. (1992) The usability engineering life cycle. *Computer* **25**(3): 12–22.

Nielsen, J. (2001), Usability metrics, Alertbox January 2001.

Olson, J.S., Grudin, J. and Horvitz, E. (2005) A study of preferences for sharing and privacy. In *Proceedings of CHI 2005* (Portland, OR).

Palen, L. and Dourish, P. (2003) Unpacking 'privacy' for a networked world. In *Proceedings of CHI 2003* (Ft. Lauderdale, FL).

Patil, S. and Lai, J. (2005) Who gets to know what, when: configuring privacy permissions in an awareness application. In *Proceedings of CHI 2005* (Portland, OR).

Razavi, M.N. (2009) Towards usable end-user privacy control for social software systems, PhD Thesis, available online at http://circle.ubc.ca/handle/2429/13403.

Razavi, M.N. and Iverson, L. (2008) Supporting selective information sharing with people-tagging. In *Proceedings of the ACM CHI '08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy).

Razavi, M.N. and Iverson, L. (2009). Improving personal privacy in social systems with people-tagging, In *Proceedings of the ACM GROUP '09 on Supporting Group Work* (Sanibel Island, FL).

Spiekermann, S., Grossklags, J. and Berendt, B. (2001). E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proceedings of Electronic Commerce* (Tampa, FL), 38–47.

Tosh, D. and Werdmuller, B. (2004) ePortfolios and Weblogs: One Vision for ePortfolio Development, ePortfolio research and development community (ERADC).

Westin, A.F. (1991) *Harris-Equifax Consumer Privacy Survey* (Atlanta: Equifax).

Westin, A.F. (2003) Social and political dimensions of privacy. *J. Soc. Issues* **59**(2): 431–453.

Whalen, T. and Gates, C. (2005). Private lives: user attitudes towards personal information on the web, poster in SOUPS.

Wharton, C., Rieman, J., Lewis, C. and Polson, P. (1994) The cognitive walkthrough method: A practitioner's guide. In Nielsen, J. and Mack, R.L. [eds.] *Usability Inspection Methods* (Wiley & Sons), 84–89.

Whitten, A. and Tygar, J.D. (1998) Usability of security: a case study. Technical Report CMU-CS-98-155 (Carnegie Mellon University School of Computer Science).