DISCLOSURE TEMPLATES FOR SELECTIVE INFORMATION DISCLOSURE

Tom Gross and Mirko Fetter
Bauhaus-University Weimar

ABSTRACT

Cooperative environments often capture data about users and mutually inform users in order to facilitate coordination in distributed workgroups. For the users this entails benefits, but also challenges regarding privacy. In this paper we introduce the concept of *Disclosure Templates* to help users configure environments according to their privacy needs. They provide powerful configuration while keeping users' effort low. These Disclosure Templates have been derived from a literature study and an empirical study, and they have been implemented in the PRIMIFaces environment.

KEYWORDS

Cooperative Environments; Awareness; Privacy; Selective Information Disclosure.

1. INTRODUCTION

Cooperative environments are often based on computing technology that provides functionality to capture data from the situation, process them, and distribute them (Gross 2008). These data can be used for many cooperative purposes—for instance, they can notify users about each others' current state of the environment in order to facilitate coordination in distributed workgroups (e.g., Sens-ation (Gross et al. 2006), Khronika (Loevstrand 1991), Elvin (Fitzpatrick et al. 1999), Siena (Carzaniga et al. 2001)). They can comprise facts on the technological environment, but also about the present users' state and activities entailing challenges for the present users' privacy.

Already in the early 1990s Bellotti and Sellen claimed that users should get *feedback* about the data that are collected on them, and to have *control* to influence the disclosure of information about themselves (Bellotti & Sellen 1993). This feedback and control empowers users to configure their cooperative environments according to their preferences.

However, it can also cause considerable extra *effort* to users. This is well known in literature. For instance, Boyle and Greenberg (2005) point out that users should be able to control privacy by easily balancing the amount of information to be disclosed. The authors claim for fine-grained (i.e., allowing for precise control of who sees what and how detailed) and lightweight (i.e., reducing the effort for the users) mechanisms. Also, Richards and Christensen (2004) emphasise the complexity and effort that underlies the control of permissions when disclosing information and argue for understandable and easy control mechanisms.

We have developed an innovative concept for *Disclosure Templates* that empowers users to manage their privacy by easily configuring their preferences for selective information disclosure in cooperative environments. In this paper we present this concept of Disclosure Templates—templates for the efficient selective disclosure of information—that are based on a thorough research of literature and a multilevel user study. We illustrate the conceptual background as well as the underlying concept and describe the seven templates. We explain in detail how these templates were derived from literature and a user study and how they are integrated into the PRIMIFaces environment. We discuss the related work and summarise the paper.

2. BACKGROUND

The need for and the underlying mechanisms to control the sharing of specific personal information to groups of recipients, here referred to as selective information disclosure, is a well known phenomenon in social science (e.g., Altman 1975; Goffman 1959). Altman (1975) defines privacy as 'the selective control of access to self or to one's group' characterising the level of openness and closeness to others as the two facets of this concept. According to Altman, maintaining privacy is a constant optimisation process balancing between the adequate amount of being left alone and having social contact. In Computer-Supported Cooperative Work the notion of privacy plays an important role for users and their social interaction. It is reflected in the two central aspects confidentiality (i.e., control over data a person gives away), and solitude (i.e., control over data a person receives) (Davis & Gutwin 2005).

Despite these conceptual findings, practical support for privacy is still lacking in several systems (Gross et al. 2005). As Palen and Dourish (2003) point out, it is necessary for instant messaging, for group calendar systems, and for ubiquitous computing to support selective information disclosure. Yet, Ackerman (2000) states that current systems rely on simple models for disclosing information, contrasting with the complex understandings 'of people's views of themselves, the current situation, and the effects of disclosure'. Olson et al. (2005) identified with their study the actual demand for selective information disclosure in information systems and also depict dependencies in-between the disclosure of different kinds of information and in-between the addressed groups of persons. Their study also shows that clusters of the information recipients as well as of the information kinds can be identified. These similarities in disclosure were also identified in other studies (e.g., Davis & Gutwin 2005; Lederer et al. 2004; Patil & Kobsa 2004; Patil & Kobsa 2005; Patil & Lai 2005). Accordingly, we can distinguish three fundamental elements that are central to the process of selective information disclosure:

- *recipients* (persons the disclosed information is made accessible for)
- *information* (personal data that is made accessible for others)
- *disclosure precision levels* (amount of disclosed information according to completeness and exactness)

In order to configure recipients, information, and disclosure precision levels, users need indepth feedback about the type of information and the level of detail at which information is revealed about them. Based on this feedback users are able to choose their desired disclosure precision levels. Various studies (e.g., Patil & Kobsa 2005) point out that mechanisms that provide detailed feedback over the disclosed information, influence the users in a way that they averagely disclose more detailed information. Patil and Kobsa (2005), therefore, suggest to make this feedback available to the users with the assumption that this will raise the overall average precision levels of disclosed information.

3. CONCEPT OF DISCLOSURE TEMPLATES

The Disclosure Templates we developed are based on a thorough literature study of selective information disclosure, and particularly on the work of Erwin Goffman (1959). Goffman's *Faces* are a central concept on how humans try to continuously manage the impression they make on others in face-to-face interaction. This act of impression management includes the disclosed information itself as well as the presentation of this information in form of mimic, gesture, and language. In this way, people easily adapt to different social situations by showing specific faces according to the audience, the context, and so on.

For our Disclosure Templates we extend the concept of Faces and apply it as a mechanism to manage privacy by controlling selective information disclosure in cooperative environments. We first describe our concepts for information sources and disclosure precision levels, and then our concept for seven Disclosure Templates—all derived from literature.

3.1 Information Sources and Precision Levels

In order to support selective information disclosure based on Faces we identified *ten information types* for the presentation of self from literature. We distinguish information types that are easy to capture from those that need to be inferred. The information types that are *easy* to capture are:

- *Personal information* (Lederer et al. 2004): data like the name, postal address, email address, and telephone number that can be valuable for choosing a right mode for contacting.
- *Location* (Lederer et al. 2004; Patil & Lai 2005): data on the whereabouts of the user that can give insights on the activity or be helpful to select a meeting place.
- *People in proximity* (Lederer et al. 2004): names of other users in a short geographical distance that allow drawing conclusions about the social context.
- *Calendar* (Patil & Lai 2005): data on upcoming appointments that can be valuable for scheduling meetings.
- *Phone status* (Patil & Lai 2005): data on the telephone use, to estimate the chance of success for contacting a person via telephone or to evaluate the context of the current communication of the user.
- *Applications* (Patil & Lai 2005): data on running software that can provide insights about current tasks.

Furthermore, there are information types that are *inferred* from single or combined data derived from the above information types:

- *Computer activity* (Patil & Lai 2005): inferred data on the user interaction with the computer, to estimate which type of computer mediated contacting is adequate.
- *Activity* (Lederer et al. 2004; Patil & Lai 2005): inferred data on activities that allow an estimation of the current situation.
- *Person in conversation* (Patil & Lai 2005): inferred data on social interaction to avoid unintentional disruptions.
- *Availability* (Patil & Lai 2005): inferred data on the availability based on several information types, whereby the possibility to contact the person is made better assessable.

For each information types we define *four precision levels* that identify the exactness and completeness of the disclosed information (also see: Lederer et al. 2004; Patil & Lai 2005):

- *Precise:* the information is disclosed unchanged—that is, exactly and completely.
- *Approximate:* the selected information is disclosed exactly, but incompletely.
- *Vague:* the selected information is disclosed rounded and incompletely.
- *Undisclosed:* no information is disclosed.

3.2 Seven Disclosure Templates

We define seven Disclosure Templates, each is an archetype consisting of permutations of information types and precision levels. These are based on our studies from literature on the ability to classify information disclosure behaviour (Davis & Gutwin 2005; Lederer et al. 2004; Patil & Kobsa 2005; Patil & Lai 2005) and have been verified in our own study (see next section). The Disclosure Templates represent different *categories of trust*; those are labelled after groups of persons that are typically recipients of such information. The seven categories in descending order of the overall disclosure are:

- *Family:* information disclosure to persons standing in a close family relationship with the person (e.g., parents, siblings, grand parents)
- *Friends:* information disclosure to persons with a close, personal relationship with high sympathy and a high degree of trust (e.g., good friends, confidante)
- *Partner:* information disclosure to the significant other of the person (e.g., friend, fiancé, husband)
- *Team:* information disclosure for groups of persons that are united by a common goal (e.g., project group, sport team)
- *Superiors:* information disclosure to persons that the person reports to (e.g., boss, trainer).
- *Subordinates:* information disclosure to persons that report to the person (e.g., employees, trainee)
- *Public:* information disclosure to other person, independent of the level of familiarity (e.g., fellow passenger in a train, passers-by in a park)

These are typical categories of trust in respect to the disclosure of information—that is, they represent types of canonical relationships that often occur. Yet, there can be different use (e.g., a user A can have big trust towards a colleague B, and therefore put colleague B in the category Friend rather than Team).

4. SURVEY FOR DISCLOSURE TEMPLATES

In our user study we analysed if the social contacts in a cooperative environment can be mapped adequately to the seven categories of trust in terms of information disclosure and how appropriate permutations of information types and precision levels for each of these categories should look like.

4.1 Structure of the Survey

Our empirical study was conducted in the form of an online survey. In an introduction a scenario on cooperative environments with computer-mediated communication as well as contact management including disclosure of information captured from different sensors was described as the setting. The task for the subjects was to specify the disclosure precision level for each social context. 38 persons provided complete information; this information is analysed in the following. The subjects (12 female, 26 male) predominantly had a university background and an average age of 23,6 years.

The survey consisted of four parts: the first two parts refer to the mapping, the last two parts to the permutations. In the *first part* the subjects were asked to divide their social environment into seven categories of trust. Hereupon, in the *second part,* the test persons specified for each category their privacy needs in terms of the disclosure precision level for all of the ten information types mentioned above. Therefore, they were presented with a description and an example for each information type at each precision level. In the *third part* the subjects were invited to specify again the most suitable disclosure precision levels for all information types, but this time for the seven categories of Disclosure Templates that where developed by us. In the *fourth part* the effects of transparency of the communication system in form of feedback about the actual disclosed information on the configuration of each of the Disclosure Templates was analysed.

4.2 Results

Table 1 depicts the final configuration of the seven templates after part four; Figure 1 below shows the mean of the selected precision levels after part three.

Table 1. Seven disclosure templates (*Precise*=1, *Approximate*=2, *Vague*=3, *Undisclosed* =4).

	Family	Friends	Partner	Team	Superiors	Subordinates	Public
Personal information	1	1	1	2	2	3	4
Location	2	2	2	2	3	3	4
People in proximity	2	2	2	2	3	4	4
Calendar	2	2	2	2	3	3	4
Phone status	2	2	2	2	3	4	4
Applications	3	3	3	3	3	4	4
Computer activity	3	3	3	2	3	4	4
Activity	2	2	2	2	3	4	4
Person in conversation	2	2	2	2	3	3	4
Availability	2	2	2	2	3	3	4

The seven categories named by the subjects were compared with the Disclosure Templates defined by us. Thereby synonyms were taken into account when applicable. For instance, the categories work, colleagues, job were all named by different users, but had the same meaning and were therefore assigned to the

Disclosure Template Team. This resulted in the following occurrences for mentioned categories of trust that comply with Disclosure Templates (multiple entries possible): Family: 73,7%, Friends: 100%, Partner: 10,5%, Team (university related 65,8%, work related 55,3%), Superior: 5,3%, Subordinates: 0%, and Public: 28,9%.

For the best matching entries Family, Friends, and Team we compared if there were also matches in the characteristics of the permutation. For all three categories the average match was nine out of ten information types—that is, in each case only for one information type there was a difference in the specified precision level (please note that the magnitude of the difference was not considered; e.g., there was no difference made between the difference of precise and vague versus precise and undisclosed). This suggests that there is a strong coherency between our Disclosure Templates and the need for selective information disclosure stated by our subjects—both by name, but also by content.

The overall distribution of the precision levels, independent of the Disclosure Templates and the information types was the following: Approximate (44,3%), Vague (30%), Undisclosed (21,4%), and Precise (4,3%). Finally, the average chosen precision level in part four was lower than in part three.

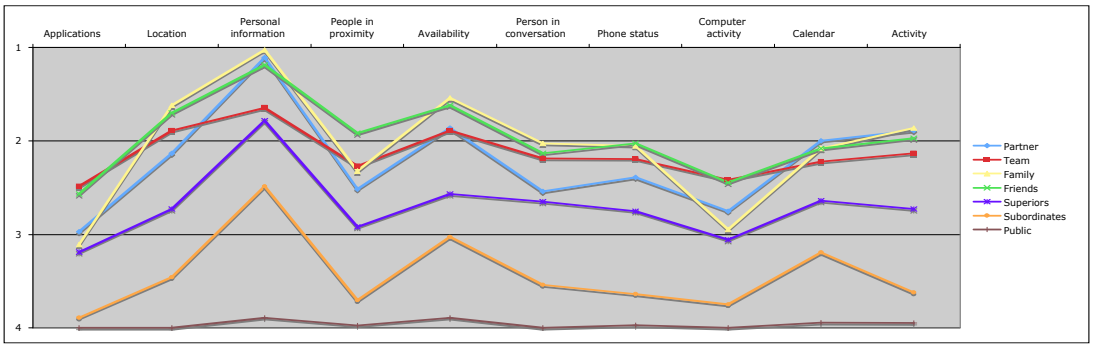

Figure 1. Mean of selected precision levels after the third part of the study.

4.3 Discussion

Our survey confirmed the three Disclosure Templates Family, Friends, and Team to be very relevant for the majority of test persons. The categories Superior, Subordinate, and Partner were mentioned less often.

Both the test persons' actual behaviour as presented in the facts above and their comments show that they had privacy needs and thus were reluctant to disclose of personal information. Hardly any information was disclosed with the highest precision level. Most categories of trust were configured with an approximate disclosure precision level.

The assumption that increased transparency has a positive effect on the disclosure of information as stated in (Patil & Kobsa 2005; Patil & Lai 2005) could not be confirmed. The verification in part four rather resulted in a tendency towards users lowering the precision levels, when they were aware which concrete personal information they reveal to other persons.

5. DISCLOSURE TEMPLATES IN PRIMIFACES

PRIMIFaces is an environment allowing users to adapt the disclosure of information captured in their cooperative situation to different social contexts (Gross & Oemig 2006). For this purpose PRIMIFaces is based on three concepts: *Faces* are typical facades of self-presentation (disclosed information sources and the recipients of this information), and information needs (preferences for the content and presentation of notifications about other users with the aim to minimise disruptions). *Contacts* are the recipients of the outgoing information. *Information sources* define the origin and type of the information that can be disclosed. Information captured from static sources (email address, cell phone number) and dynamic sources that are captured by software sensors (calendar, mouse movement) or hardware sensors (GPS position).

In the following we describe how the concepts of the Disclosure Templates are mapped to PRIMIFaces. The Faces in PRIMIFaces are used to bundle recipients along with some of the ten types of information each with its particular level of precision into a specific instance. The contacts in PRIMIFaces relate to the recipients of the Disclosure Templates. Each information source is assigned to a concrete information type from the Disclosure Templates and for each information source the four precision levels are operationalised (e.g., the information source Application Sensor in PRIMIFaces is assigned to the information type running applications and operationalised with the precision level approximate as, for instance, mail application running (instead of the concrete application Thunderbird)).

In order to simplify the configuration of faces, the interaction with PRIMIFaces is based on the FamilyFaces card-game-metaphor (cf. Figure 2). Each Face is mapped to a playing card divided into three areas on which the icons for contacts (centre), information source (top), and notification preferences (bottom) can be dragged in order to configure a face. The green playing field contains all cards of the current faces-configuration and offers the possibilities to manage these faces. The three pools on the right provide all available information sources, contacts and notification preferences of the user. So, handling disclosure is easier than other access control mechanisms such as in operating systems. More details on PRIMIFaces and the FamilyFaces card-game-metaphor can be found in (Fetter & Gross 2008; Gross & Oemig 2006).

In order to create new faces, users can select from a list of available templates, the Disclosure Template that fits their current needs best. Accordingly a new face is created with all information types and their corresponding precision level as specified in the Disclosure Template. The users then can customise this face by adding contacts and further information sources respectively. Likewise the users can delete a face.

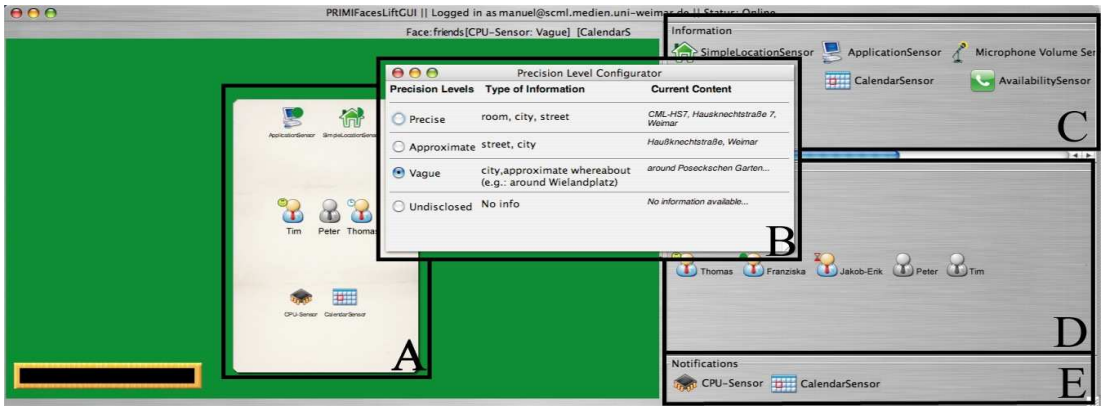

Figure 2. Scenario with the user Manuel. A: face Friends with the contacts Tim, Peter, Thomas in the centre, disclosed information sources like application sensor and location sensor at the top and incoming notification preferences CPU sensor and calendar at the bottom; B: configuration of the precision level for the location sensor; C: pool of information sources; D; pool of contacts; E: pool of incoming notifications.

As a remark we would like to mention that trust does not need to be symmetrical and accordingly the disclosure does not need to be reciprocal (e.g., Person A can categorise B as a friend, but B can have A in the face colleagues). In order to preserve privacy the users do not know how they are categorised by others. Furthermore, there is a possibility that persons are connected over multiple faces. (e.g., Person A and B can interact in a work as well as in private context). Finally, Faces can manually be activated and deactivated in different situations temporally.

By double-clicking on an information source, the users open a pop-up window that allows the modification of the precision level for this information source (cf. Figure 2). This way, the disclosure of this information for a face can be altered. The template on which the face is based on is not affected by this modification.

Furthermore, users can modify existing templates or to create their own templates. Therefore, functions to create new or delete existing faces, to adjust the precision levels for the individual information types, and to add and remove information types, are available. Changes to a Disclosure Template do not affect existing faces that are based on these templates—this way the consequences of adjusting a template can be better comprehend by the user.

PRIMIFaces provides advanced feedback and control concepts. In order to give users feedback on their information that is disclosed to others, we introduce the *GatheringLog*. Each entry in this log represents a single information disclosure item consisting of the contents of the information disclosed as well as the recipient, the face, the time, the information type, and the precision level. Filters allow users to explore and filter log data: they can select all entries for one face, for an information type on a face, for a person on a face, for an information type on all faces, for a person on all faces, or for all information types and all persons. If users discover that a disclosure does not fit their personal needs, the GatheringLog provides control and offers the possibility to directly jump to the corresponding disclosure in order to adjust the precision level.

The concepts for the Disclosure Templates as well as the GatheringLog have been integrated in PRIMIFaces and implemented in Java SE 1.5.0_13 on Mac OS X 10.4.11. The persistence of the Disclosure Templates and of the GatheringLogs is realised with our own XML format with means of the Java Architecture for XML Binding (JAXB). As backend the SensBase infrastructure is utilised for the management of face-configurations and sensor data (Gross et al. 2006).

6. RELATED WORK

The related work lies primarily in the field of privacy and information disclosure as well as in selective information disclosure. Kapadia (2007) developed a concept for policies building on the metaphor of a virtual wall. These policies allow users to protect or disclose sensor information ('digital footprints') at three different levels of transparency: transparent, translucent, and opaque. However, the concept is bound to places and does not allow users to specify disclosure levels on a per sensor basis. Furthermore, there are no templates that help users in the initial configuration phase. They admit that it might 'be cumbersome for users' to deploy walls at every place they visit. IBM's *Grapevine* (Richards & Christiansen 2004) allows specifying permissions based on relationship, groups, or situations. However, it does not take into account different disclosure precision levels and does not provide templates. Langheinrich (2005) developed and implemented *paws*—a concept for the exchange of privacy policies. paws focuses on an automated use of machine-readable policy files and therewith ignores aspects of supporting users when configuring disclosure. Lederer el al. (2004) developed a concept that is called *faces*. It allows users to assign what information they want to disclose for each contact depending on the current situation. These faces are quite powerful, but come along with a high configuration effort for users as there are only single contacts, no grouped contacts and there are no templates.

The major strength of our approach of Disclosure Templates—in comparison to the discussed related work—lies in providing the user with a very simple mechanism for handling complex and elaborated configurations for the disclosure of personal information in cooperative environments.

7. CONCLUSIONS

We presented the concept of Disclosure Templates for the selective information disclosure in cooperative environments and its implementation in PRIMIFaces. It provides fine-grained feedback and control of privacy and helps to minimise the effort for users making it a lightweight solution and so satisfies the requirements introduced by (Boyle & Greenberg 2005). Informal evaluations of PRIMIFaces during our open house showed that people instantly could transfer the notion of faces from their social lives to the metaphor we used for privacy control in cooperative environments.

A limit is that Disclosure Templates only capture a specific point in time—that is, the current concept does not provide continuous adaptation to social relations that change over time. Furthermore, users currently need to manually activate and deactivate faces; in future we plan to provide mechanisms based on machine learning. Besides, a long-term study that delivers insights about the acceptance of Disclosure Templates by the users would be desirable. Finally, a user study on measuring the increase of efficiency when specifying the disclosure of information with and without Disclosure Templates is intended. In order to cope with a larger number of sensors we currently work towards taxonomy of sensor information and heuristics for the definition of precision levels.

ACKNOWLEDGMENTS

We thank the members of the CML; esp. Benjamin Zeller and Thilo Paul-Stueve. Part of the work has been funded by the Federal Ministry of Transport, Building, and Urban Affairs (TransKoop FKZ 03WWTH018).

REFERENCES

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

Altman, I. *The Environment and Social Behaviour: Privacy, Personal Space, Territory, Crowding*. Brooks/Cole Publishing Company, Monterey, CA, 1975.

Bellotti, V. and Sellen, A. Design for Privacy in Ubiquitous Computing Environments. In *Third European Conference on Computer-Supported Cooperative Work - ECSCW'93*. Kluwer Academic Publishers, Dortrecht, NL, 1993. pp. 77-92.

Boyle, M. and Greenberg, S. The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM Transactions on Computer-Human Interaction 12*, 2 (June 2005). pp. 328-370.

Carzaniga, A., Rosenblum, D.S. and Wolf, A.L. Design and Evaluation of a Wide-Area Event Notification Service. *ACM Transactions on Computer Systems 19*, 3 (Aug. 2001). pp. 332-383.

Davis, S. and Gutwin, C. Using Relationship to Control Disclosure in Awareness Servers. In *Conference on Graphics Interface - GI 2005*. Canadian Human-Computer Communications Society, Waterloo, Canada, 2005. pp. 145-152.

Fetter, M. and Gross, T. Contact Management on the Wall: A Card-Game Metaphor for Large Displays. In *Second International Conference on Tangible Embedded Interaction - TEI 2008*. ACM, N.Y., 2008. pp. 247-250.

Fitzpatrick, G., Mansfield, T., Kaplan, S., Arnold, D., Phelps, T. and Segall, B. Augmenting the Workaday World with Elvin. In *Sixth European Conference on Computer-Supported Cooperative Work - ECSCW'99*. Kluwer Academic Publishers, Dortrecht, NL, 1999. pp. 431-450.

Goffman, E. *The Presentation of Self in Everyday Life*. Doubleday Anchor Books, N.Y., 1959.

Gross, T. Cooperative Ambient Intelligence: Towards Autonomous and Adaptive Cooperative Ubiquitous Environments. *International Journal of Autonomous and Adaptive Communications Systems (IJAACS) 1*, 2 (2008). pp. 270-278.

Gross, T., Egla, T. and Marquardt, N. Sens-ation: A Service-Oriented Platform for Developing Sensor-Based Infrastructures. *International Journal of Internet Protocol Technology (IJIPT) 1*, 3 (2006). pp. 159-167.

Gross, T. and Oemig, C. From PRIMI to PRIMIFaces: Technical Concepts for Selective Information Disclosure. In *32nd EUROMICRO Conference on Software Engineering and Advanced Applications - SEAA 2006*. IEEE Computer Society Press, Los Alamitos, 2006. pp. 480-487.

Gross, T., Stary, C. and Totter, A. User-Centered Awareness in Computer-Supported Cooperative Work-Systems: Structured Embedding of Findings from Social Sciences. *International Journal of Human-Computer Interaction 18*, 3 (June 2005). pp. 323-360.

Kapadia, A., Henderson, T., Fielding, J.J. and Kotz, D. Virtual Walls: Protecting Digital Privacy in Pervasive Environments. In *5th International Conference on Pervasive Computing - Pervasive 2007*. London, UK, 2007.

Langheinrich, M. *Personal Privacy in Ubiquitous Computing: Tools and System Support*. Ph.D. thesis, Swiss Federal Institute of Technology Zurich, Zurich, Switzerland, 2005.

Lederer, S., Hong, J.I., Dey, A.K. and Landay, J.A. Personal Privacy Through Understanding and Action: Five Pitfalls for Designers. *Personal and Ubiquitous Computing 8*, 6 (Nov. 2004). pp. 440-454.

Loevstrand, L. Being Selectively Aware with the Khronika System. In *Second European Conference on Computer-Supported Cooperative Work - ECSCW'91*. Kluwer Academic Publishers, Dortrecht, NL, 1991. pp. 265-278.

Olson, J.S., Grudin, J. and Horvitz, E. Late Breaking Result: A Study of Preferences for Sharing and Privacy. In *Conference on Human Factors in Computing Systems - CHI 2005*. ACM, N.Y., 2005. pp. 1958-1988.

Palen, L. and Dourish, P. Unpacking 'Privacy' for a Networked World. In *Conference on Human Factors in Computing Systems - CHI 2003*. ACM, N.Y., 2003. pp. 129-136.

Patil, S. and Kobsa, A. Instant Messaging and Privacy. In *18th British HCI GroupAnnual Conference - HCI 2004*. Springer-Verlag, London, UK, 2004. pp. 85-88.

Patil, S. and Kobsa, A. Uncovering Privacy Attitudes and Practices in Instant Messaging. In *International ACM SIGGROUP Conference on Supporting Group Work - Group 2005*. ACM, N.Y., 2005. pp. 109-112.

Patil, S. and Lai, J. Who Gets to Know What When: Configuring Privacy Permissions in an Awareness Application. In *Conference on Human Factors in Computing Systems - CHI 2005*. ACM, N.Y., 2005. pp. 101-110.

Richards, J.T. and Christiansen, J. People in Our Software. *ACM Queue 1*, 10 (Feb. 2004). pp. 80-86.