SUBMISSION TO IE6

Title:

SPIROS - A System For Privacy-Enhanced Information Representation In Smart Home Environments

Authors:

Carsten Röcker, Steve Hinske, Carsten Magerkurth

Adresses:

Carsten Röcker
Fraunhofer IPSI
Dolivostrasse 15
64293 Darmstadt
Germany

Steve Hinske
ETH Zurich
IFW D48.2
Haldeneggsteig 4
8092 Zurich
Switzerland

Carsten Magerkurth
Fraunhofer IPSI
Dolivostrasse 15
64293 Darmstadt
Germany

SPIROS - A SYSTEM FOR PRIVACY-ENHANCED INFORMATION REPRESENTATION IN SMART HOME ENVIRONMENTS

C. Röcker[(1)], S. Hinske[(2)], C. Magerkurth[(1)]

[(1)] Fraunhofer IPSI, Germany
[(2)] ETH Zürich, Schwitzerland

Abstract: This paper presents a novel concept for personalized privacy support on large public displays in intelligent home environments. In order to validate the conceptual approach a system called SPIROS was developed. The SPIROS system automatically adapts the information visible on public displays according to the current social situation and the individual privacy preferences of the user working at the display.

Keywords: Large Public Displays, Active Privacy Support, Privacy-Enhancing Technologies, Context-Adapted Information Representation.

INTRODUCTION

The concept of Ambient Intelligence propagates a vision of future environments where people are supported and assisted in their everyday activities by information technology that is very different from the computer as we know it today (1). The envisioned technologies "will weave themselves into the fabric of everyday life until they are indistinguishable from it" (2). By making many computers available throughout the physical environment, people are enabled to move around and interact with computers more naturally than they currently do. Instead of using traditional personal computers, users can access information using computational devices integrated into the environment.

As ubiquitously available displays are an integral part of smart home environments, several projects provide special 'walk-up-and-use' applications for the home domain. Most of these applications employ large-screen display technologies to provide context-dependent services and thereby offer access to personal information. For example, the Amigo project (http://www.amigo-project.org) is currently developing several applications that use large-screen displays to provide users with personalized services in common areas. The intended applications range from multi-user gaming applications to personalized awareness displays situated in public areas.

In general, the use of shared and public devices to access information in home environments is neither unusual nor new. Televisions, radios or stereos are just a few examples of shared devices surrounding us (3). But while traditional public devices are only used to access non-personal information, smart home environments provide users with personalized services and information. Many of these services are intended to provide walk-up-and-use functionality, like quickly accessing emails or the Internet.

When using such applications on large public displays, the possibility of other people being able to see confidential information inevitably causes privacy concerns. And empirical evidence shows that these concerns are justified. Exploring the influence of the display size on privacy infringements, Tan and Czerwinski (4) found that, even given constant visual angles and similar legibility, individuals are more likely to read text on a large display than on a small one. Several other studies reported similar results and emphasized the importance of informational privacy, both in the office and in the home domain.

Generally, users are quite concerned to share personal information with other users, even if they are closely acquainted. A recent evaluation (5) demonstrated that, for example, only 16% of the participants were willing to share their internet history with friends and family members. Less than 4% would provide this information to everyone and 42% of the participants regarded the information as

totally private and refused to share it at all. Similar feedback regarding privacy requirements was gained in a cross-national focus group study on smart home environments (1). The results confirm the important role of privacy in ubiquitous computing environments (6) and at the same time show that personal privacy protection is a major issue, even in familiar surroundings, like the own home or towards family members and friends.

EXISTING APPROACHES TO PROTECT PRIVACY ON PUBLIC DISPLAYS

Although privacy guidelines have been available for quite some time (e.g., 7, 8) most developers still rely on social protocols or do not address privacy questions at all, when designing applications for large public displays. Until today, there are very few approaches that help users in preserving their privacy while working on large displays in public places.
In most cases (see, e.g., 9, 10) additional private displays are used to generate and present personal information, while public information is displayed on a shared large display. A different approach using a stereographic display and special shutter glasses is described in (11). Personal privacy is maintained through the filtering of the information on the shared display. Users wearing shutter glasses will see the public information as well as their own private information, while other people's private data is not visible to them. A similar system was developed by Yerazunis and Carbone (12), which uses time-masked images to ensure privacy. As the system requires CRT displays with higher-than-usual screen refresh rates and special eyewear, the usage of the system is likewise restricted. Comparable systems were developed by Eaddy et al. (13), Needham and Koizumi (14) and Berger et al. (15).
Although all approaches support individual privacy in an adequate way, they always require additional personal devices, like PDAs or shutter glasses. But as most public displays are intended for walk-up-and-use applications, the existing solutions are not very suitable.

GOAL AND APPROACH

The goal of this paper is to give users the freedom to spontaneously work on large public displays, without the fear of privacy infringements through passers-by. It is aimed to provide users with a system, which automatically controls the information that is visible to others, without requiring users to employ any additional equipment.
This goal will be achieved by providing users with a 'private space' in a shared home environment. Within that personal space, the information that is visible to others is automatically controlled according to the user's individual preferences. In order to adapt the information representation to the current context, people entering the private space around a public display are automatically detected and identified. Based on the identity of the person(s) entering the private space and the privacy preferences of the user working on the public display, the information currently visible is automatically adapted.

CONCEPTUAL DESIGN

Before the technical realization of the system is presented, the key ideas which guided the conceptual design process are illustrated.

Personal Space Concept

The desire for a personal space that is not penetrated by others is one of the most basic human rights and it has become even more important within the last decades. This tendency is reflected in our society by increasing privacy concerns regarding current and future information and communication technologies (see, e.g., 6, 16, 17). Clarke (18) defines this desire as "the interest that individuals have in sustaining a personal space, free from interference by other people".
In this paper, this real-world concept is adapted to the virtual domain. The term 'personal space' refers to the space in which all interactions are only visible to the interacting user. Or, from another perspective, being outside this specific private space does not allow another person seeing the content and nature of the interaction taking place.
In the following figures the personal space concept is illustrated. Figure 1 shows a user working on a large display in a public space. The semicircle in front of the display shows the user's personal space. In this example, the personal space represents exactly the physical space that is necessary to remain free of interruptions or intrusions. As this space is currently not invaded by other users, all content is being displayed. The black areas on the display represent applications or documents with confidential content, the grey areas indicate applications with public information.

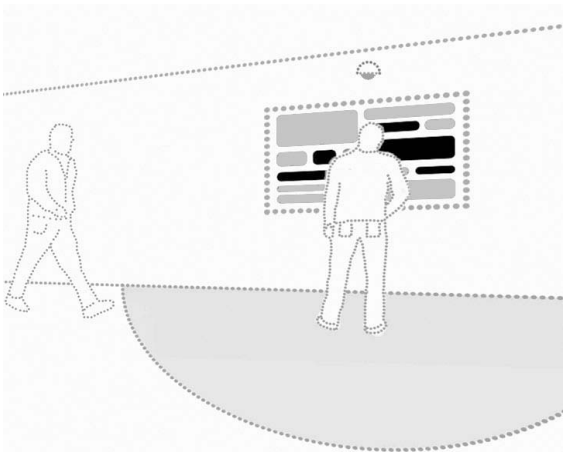

Fig. 1. Personal space around a public display.

In Figure 2 another person is entering the personal space. As this person might threaten the information privacy of the user working at the display, the confidential content (black outline) is temporarily concealed, while the public information (grey) is still visible.

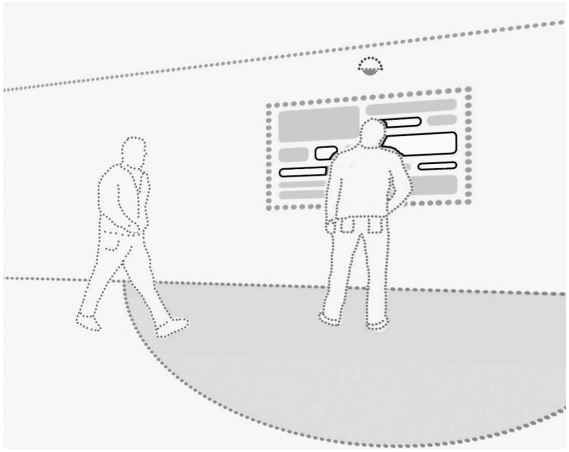

Fig. 2. Confidential information is hidden as soon as another person invades the personal space.

Context-Adapted Privacy Protection

To guarantee satisfactory privacy protection and at the same time offer maximum usability, it is essential to know which information should be hidden from whom and when. Protecting family-related information from other family members passing-by might not be necessary and doing so would most likely result in an unintended interruption of the ongoing activity. But the situation is fundamentally different, when accessing private information while other people approach the display.

The problem encountered here is that information is not generally public or private. Rather, it depends on users, how confidential they regard certain types of information. Palen (19) found, that information regarded totally innocuous by some users, were considered personally private to others. In the same way, Zhao and Stasko (20) argue, that individuals usually have different comfort zones in the level of personal information being broadcasted and that these comfort zones change over time. They conclude that individuals should be allowed to select the level of information about him or her being transmitted. But the privacy settings are not only dependent on the sender's preferences; they are also determined by the information receiver. The behavior regarding the disclosure of personal information in multi-user situations was investigated in various studies (see, e.g., 5, 21). All studies came to the result, that the willingness to provide information varies widely with the type of information and the information receiver.

Privacy Levels to Control Information

In real-world situations, the disclosure of personal information is usually done on an ad-hoc basis and mostly does not follow any strict rules. But when using an automated privacy protection system, the disclosure of personal information has to follow a basic concept with specific „rules".

As illustrated before, users have different preferences depending on what they are doing and to whom they are providing certain information. Hence, the adaptation of the displayed information must depend on the content of the application or document as well as on the person(s) invading the personal space. But using individual privacy settings for each situation would require that all persons, documents and applications have to be classified beforehand.

Group-Based Privacy Control. Experiences with existing systems showed that individual classifications are not necessary to manage privacy in multi-user situations. In a study with N=36 users, Patil and Lai (22) found a significant preference for defining privacy permissions at group level. Around 70% of the participants chose to configure permissions in 'group mode', with significantly different permissions granted to the various groups. Based on the feedback gained by the participants, they conclude that utilizing grouping mechanisms provides the flexibility needed to appropriately manage the balance between privacy control and configuration burden. Results leading to similar design recommendations were found in studies by Olson et al. (21) and Lederer et al. (23).

Classification of Users. These findings led to the approach of using a group-based classification scheme, in which each individual is assigned to a specific 'privacy level'. Table 1 gives an example how privacy levels can be defined.

Table 1: Example for Different Privacy Levels.

Privacy Setting	Description
Level 1 (Private)	The highest privacy level covers the most private content and is meant for eyes of its owner only. This includes personal emails and private documents.
Level 2 (Partner)	The second privacy level includes all information which is meant for intimate circle of persons only. In this example, the user's partner is assumed to be trustworthier than everyone else.
Level 3 (Family)	This level contains all information which is family-internal and potentially accessible by all family members.
Level 4 (Friends)	Hides family-internal information (e.g., banking information), but still allows access to other personal information (e.g., pictures from the last holidays).
Level 5 (Public)	All applications containing personal or confidential information are hidden.

Classification of Documents and Applications. The classification of documents and applications is realized through a keyword system, which allows a high degree of flexibility. Every document and application can be assigned to one of the five privacy levels shown above.

A keyword is a string that is searched for in the window title of a document or application. Fig. 3 shows part of a screenshot taken from an application with an open document. The window title contains the name of the document ('Bank of America') and the name of the application ('Microsoft Internet Explorer').

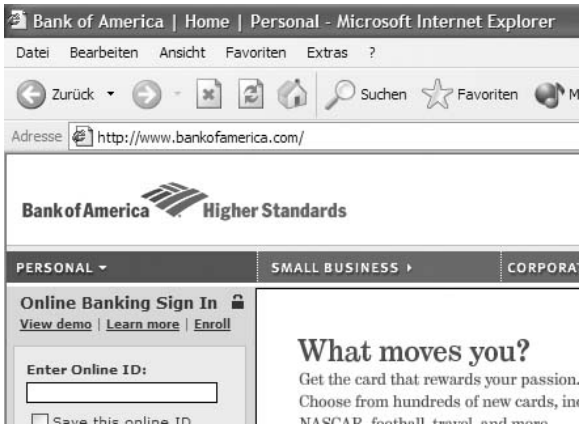

Fig. 3. Screenshot of an open application with several keywords in the title bar.

Using the keyword system, user can classify all websites accessed via the Internet Explorer to the highest privacy level (level 1) by adding the keyword 'Internet Explorer' to the keyword list of the first level. In addition, it is also possible to hide only specific websites. For example, adding the keywords 'Bank of America' to the keyword list of the second level, results in all documents containing the keywords 'Bank of America' being hidden from all users except for the persons assigned to the second level. In this case it is not important which application is used to access the document.

Measure to Preserve Privacy

As there is a natural trade-off between usability and the level of privacy protection, several protection measures were implemented. Currently, there are six possible actions available to users, each having its specific advantages and disadvantages (see Table 2).

Table 2: Advantages and disadvantages of the different measures for privacy protection.

Method	Advantages	Disadvantages
Privacy Protection De-Activated	No interference of ongoing activities	No privacy protection at all
Display Message	No interference of ongoing activities, but users are still aware of people entering their personal space	No automated privacy protection by the system, but user can manually hide specific content
Open Cover Window	All open windows remain in their position, adequate privacy protection since confidential data is visually blocked	Interference of current activities, all windows are still visible in the taskbar
Minimize Window	All open windows will appear in the same position as before, good privacy protection as confidential information is temporarily hidden	Interference of current activities, all windows are still visible in the taskbar
Hide Window	All open windows will appear in the same position as before, high privacy protection since they are temporarily hidden (even in the taskbar)	Interference of current activities
Close Window	Highest level of privacy, since all confidential applications and documents are closed	Interference of current activities, work is lost unless saved before

The settings for the different privacy levels are controlled using a graphical user interface, which enables users to adjust the levels to their individual preferences. The interface consists of two parts, the *Keyword Manager*, and the *Action Manager*.

The *Keyword Manager* comprises five keyword lists (corresponding to the five privacy levels), which can be individually adapted by the user. Furthermore, users can choose in which order the lists are checked, when a keyword is looked up. This concept enables users to express their general attitude towards information sharing. Users with a strong desire for privacy, who regard specific information as sensitive, unless expressed otherwise (e.g., by creating an exception rule), would choose to start the keyword search in the keyword list that represents privacy level 1 (private).

The *Action Manager* allows users to choose the action (as described in Table 2), which is taken in case of a keyword match. In addition, users can overrule the automatic comparison as well as the action taken for affected windows (e.g., windows that have one or more keywords in their title bar). By doing this, they can choose between "All Windows" and "Active Window". The option "All Windows" executes the selected action for every open window, regardless of the keywords in the titlebar. "Active Window" affects only the currently active window (i.e., the window the user is currently working on).

TECHNICAL REALIZATION

In order to fulfill the requirements outlined in the previous sections, a privacy protection system called SPIROS (System for Privacy-Enhanced Information Representation in Open Spaces) was developed. Following the conceptual design approach, privacy protection is achieved in a three-step process:

- First, people entering the personal space of the user are detected and if possible identified.
- Second, the system determines and subsequently compares the privacy levels of the identified person(s) with the privacy levels of the currently open applications and documents.
- Third, the system initiates privacy preserving measures according to the result of the comparison and the user's preferences.

The remainder of this chapter provides a detailed description of the underlying sensing infrastructure as well as the software architecture.

Sensing Infrastructure

The necessary information about nearby individuals is collected via a hybrid sensing infrastructure consisting of infrared sensors and RFID readers. People entering a personal space are detected by the infrared sensors and are simultaneously identified through the RFID system. All sensors are individually adjustable regarding their detection range and angle. This allows an adaptation of the user's personal space to the specific requirements of the environment.

RFID Readers. Radio frequency identification (RFID) is already widely used today (24) and belongs to the basic sensor types available in most smart home environments. RFID technology enables contactless identification of people and objects, which are equipped with a transponder. As RFID transponders are currently being integrated into a vast variety of everyday objects (like, e.g., smart jewelry or smart fashion), it is assumed that users in smart home environments are not required to carry any additional RFID transponders in order to be identified.
The SPIROS system was realized using an active RFID, which allows adjusting the reading ranges of the system to the spatial extension of the personal space. Hence, people carrying an RFID transponder are identified as soon as they enter the personal space around a public display.

IR Sensors. To detect the presence of persons who are not equipped with an RFID transponder, a special motion detection system was developed. The system consists of several infrared (IR) detectors and a central communication unit, which can handle up to ten IR detectors simultaneously. The IR sensors are distributed in the environment and detect motion within a defined area. The communication unit aggregates the signals received from all IR sensors, and transmits them via a USB link to a host computer. Hence, if people approach the display, who can not be identified, it would be still possible to hide all personal information currently being displayed.

SPIROS Architecture

The SPIROS architecture consists of three main components (see Fig. 4):

- the *SPIROS Scanner Manager* (SSM), which identifies persons entering the personal space around the public display,
- the *SPIROS Privacy Manager* (SPM), which adapts the displayed information to the pre-defined user preferences and identified persons, and
- a database for storage and communication.

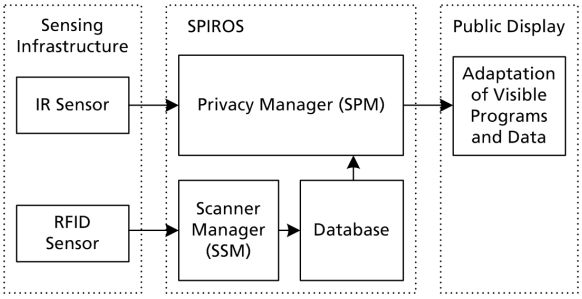

Fig. 4. General overview of the system.

The separation of the SSM and SPM was done due to performance reasons. In addition, the separation allows operating several SSMs simultaneously, with each entity being responsible for one or more large public displays. This is of particular importance if several public displays are spatially distributed in an intelligent environment.
The main function of the SSM component is the identification of people, who are within the reading ranges of the connected RFID readers. The IDs of all read RFID transponders are stored in the database. This data is then used by the SPM component. The SPM is responsible for adapting the displayed information to the user's privacy preferences and the persons currently present around the display. Whenever one of the IR sensors next to the display is triggered, the SPM retrieves the ID of the RFID transponders within the personal space and determines the privacy level of the corresponding person(s) as well as the privacy levels of the currently open windows. Based on this information, the content which is not meant to be seen by the passers-by is temporarily hidden until the persons have left the personal space.
The following sections provide a more detailed description of the different SPIROS components and illustrate their general operating principles.

SPIROS Scanner Manager (SSM). As mentioned before, the SSM operates the RFID readers and inserts all scanned tags within reading ranges of the system into the database (see below).
Once the SSM has successfully connected to the RFID reader (input) and the database (output), it starts polling the reader for RFID transponders in range. Whenever transponders are read by the reader, the SSM checks whether the transponders' IDs have already been inserted into the database. If this is not the case, the newly discovered ID is added to the database.
The main problems encountered in the test phase were malfunctions due to missing or erroneous identification data. This happened, for example, if RFID transponders were not detected although they were within the reading range. To eliminate such problems, the SSM uses an internal counter for each transponder. As explained above, the ID of each detected transponder is inserted into the database. This entry is only removed if the transponder is not read for three consecutive reading cycles.

SPIROS Privacy Manager (SPM). When the personal space of a user is invaded by other persons, the SPM hides confidential information according to the user's preferences.
On starting, the SPM retrieves the user's profile from the database, creates a temporary table, and initiates the connection to the IR sensors using DirectX. The temporary table contains the user's current relationships to all other users. This is done by assigning the user's partner(s) to the '*partner level'*, family members to the '*family level'*, friends to the '*friends level'* and all remaining users to the '*public level'*.
Once the initialization phase has been completed, the SPM switches to the operating phase. Figure 5 shows an UML diagram of the operating cycle, which is executed every couple of seconds.

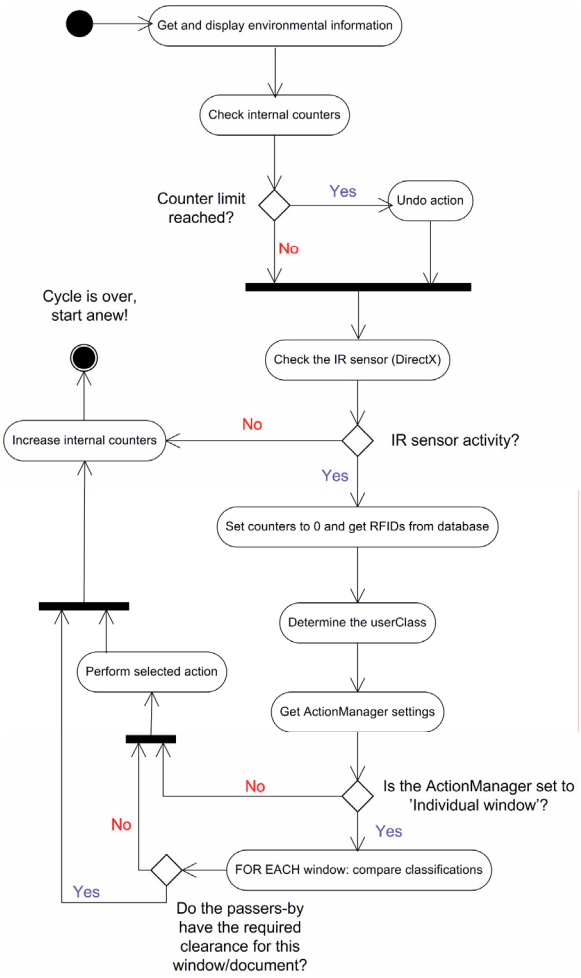

Fig. 5. UML activity diagram of the SPM component during the operating phase.

The SPM starts by checking an internal counter, which that is meant for undoing actions (e.g., windows that are minimized or hidden need to reappear). After each cycle, the counter is increased by one. If this counter reaches a certain limit, all actions taken before are reversed. In a second step, the SPM checks the IR sensors for activity. If any motion is detected, the counter is set back to zero. This is necessary in order to guarantee that all confidential information is hidden as long as other users are nearby the display.

Based on the IDs of the detected transponders, the highest privacy level of all users within the personal space is calculated. In the same way, the privacy levels of the currently open windows are determined by scanning all window titles. To determine which windows should be hidden, the privacy levels of both, the passers-by and the active windows, are compared. Depending on the user's pre-defined preferences, a specific action is initiated. See Table 2 for an overview of the possible privacy protection measures.

SPIROS Database. All information required by the SSM and SPM component is stored in a central database. The main table (*users*) contains the users' IDs as well as the individual privacy preferences. The privacy levels described in Table 1 are represented through the three tables *partners*, *family_members* and *friends*, which contain information about the social relationships among users. If users are not connected by one of these three tables, they are assigned to privacy level five ('public'). To guarantee a quick look-up, a temporary table is created, that contains the user's current relationships to all other users. The keywords used for determining the privacy level of an active window are stored in five individual keyword tables. The IDs of detected RFID transponders are stored in the table *scanned_rfids*, and are then mapped to the table *users* by the table *temp*.

CONCLUSION AND FUTURE WORK

The previous sections presented a novel concept for personalized privacy support in intelligent home environments. The goal of this paper was to give users the freedom to spontaneously work on large public displays, without the fear of privacy infringements through passers-by. Based on the conceptual approach a system was presented, which automatically adapts the information visible to others, without requiring users to carry or wear any additional technology.
After receiving promising results in a first series of tests with a small group of users, the existing implementation is currently being improved according to the received feedback. It is planned to evaluate the revised version in a second, more detailed sequence of tests with a larger group of users. This evaluation will also compare the different protection measures described in this paper. It is expected to get valuable information for usability requirements of automated privacy protection systems, which will help to guide the developments process of future systems.

REFERENCES

1. Röcker, C., Janse, M., Portolan, N., Streitz, N. A. (2005) User Requirements for Intelligent Home Environments: A Scenario-Driven Approach and Empirical Cross-Cultural Study. In: *Proceedings of Smart Objects & Ambient Intelligence* (sOcEUSAI'05), October 12th - 14th 2005, Grenoble, France, pp. 111 – 116.
2. Weiser, M. (1991). The Computer for the 21st Century. In: *Scientific American,* 265(3), pp. 66 – 75.
3. Hilbert, D., Trevor, J. (2004) Personalizing Shared Ubiquitous Computing. In: *Interactions,* 11(3), pp. 34 – 43.
4. Tan, D. S., Czerwinski, M. (2003) Information Voyeurism: Social Impact of Physically Large Displays on Information Privacy. In: *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems* (CHI'03), pp. 748 – 749.
5. Röcker, C. (2005) Providing Personalized Privacy Support in Public Places. In: *Proceedings of the Third Annual Conference on Privacy, Security and Trust* (PST'05), pp. 217 – 220.
6. Lahlou, S., Langheinrich, M., Röcker, C. (2005) Privacy and Trust Issues with Invisible Computers. In: *Communications of the ACM,* 48(3), pp. 59 – 60.
7. Langheinrich, M. (2001) Privacy by Design – Principles of Privacy Aware Ubiquitous Systems. In: *Proceedings of the 3rd International Conference on Ubiquitous Computing* (Ubicomp'01), September 2001, Atlanta, USA, pp. 273 – 291.
8. Bellotti, V., Edwards, K. (2001) Intelligibility and Accountability: Human Considerations in Context Aware Systems. In: *Human-Computer Interaction,* Special Issue on Context-Aware Computing, 16(2, 3 & 4), pp. 193 – 212
9. Rekimoto, J. (1998) A Multiple Device Approach for Supporting Whiteboard-based Interactions. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI'98), pp. 18 – 23.
10. Greenberg, S., Boyle, M., LaBerge, J. (1999) PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal. In: *Personal Technologies*, 3(1), pp. 54 – 64.
11. Shoemaker, G. B. D, Inkpen, K. M. (2001) Single Display Privacyware: Augmenting public displays with private information. In: *Proeedings of the ACM Conference on Human Factors in Computing Systems* (CHI'01), pp. 522 – 529.
12. Yerazunis, W. S., Carbone, M. (2002) Privacy-Enhanced Displays by Time-Masking Images, *Technical Report TR2002-011*, Mitsubishi Electric Research Laboratories.
13. Eaddy, M., Blasko, G., Babcock, J., Feiner, S. (2004) My Own Private Kiosk: Privacy-Preserving Public Displays. In: *Proceedings of the 8th IEEE International Symposium on Wearable Computers* (ISWC'04), pp. 132 – 135.
14. Needham, B. H., Koizumi, D. H. (1998) *Method of Displaying Private Data to Colocated Users*, US Patent 5,963,371.
15. Berger, S., Kjeldsen, R., Narayanaswami, C., Pinhanez, C., Podlaseck, M., Raghunath, M. (2005) Using Symbiotic Displays to View Sensitive Information in Public). In: *Proceedings of the 3rd IEEE International Conference on Pervasive Computing and Communications* (PerCom 2005), Kauai Island, HI, USA, pp. 139 – 148.
16. McCarthy, J. F., McDonald, D. W., Soroczak, S., Nguyen, D. H., Rashid, A. M. (2004) Augmenting the Social Space of an Academic Conference. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW'04), Chicago, Illinois, USA, pp. 39 – 48.
17. Palen, L., Dourish, P. (2003) Unpacking "Privacy" for a Networked World. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI'03), Ft. Lauderdale, Florida, USA, pp. 129 – 136.
18. Clarke, R. (1999) *Introduction to Dataveillance and Information Privacy, and Definitions of Terms* [Online], Available: http://www.anu.edu.au/people/Roger.Clarke/DV/Intro.html
19. Palen, L. (1999) Social, Individual and Technological Issues for Groupware Calendar Systems. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI'99), pp. 17 – 24.
20. Zhao, Q. A., Stasko, J. T. (2002) What's Happening?: Promoting Community Awareness Through Opportunistic Peripheral Interfaces. In: *Proceedings of the Conference on Advanced Visual Interfaces* (AVI'02), pp. 69 – 74.
21. Olson, J. S., Grudin, J., Horvitz, E. (2005) A study of Preferences for Sharing and Privacy. In: *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems* (CHI'05), pp. 1985 – 1988.
22. Patil, S., Lai, J. (2005) Who Gets to Know What When: Configuring Privacy Permissions in an Awareness Application. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI'05), pp. 101 – 110.
23. Lederer, S., Dey, A., Mankoff, J. (2003) Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems (*CHI'03), Ft. Lauderdale, USA, pp. 724 – 725.
24. Fleisch, E., Mattern, F. (2005) *Das Internet der Dinge. Ubiquitous Computing und RFID in der Praxis: Visionen, Technologien, Anwendungen, Handlungsanleitungen*. Springer-Verlag, Heidelberg.