IGI PUBLISHING ITB15288
701 E. Chocolate Avenue, Suite 200, Hershey PA 17033-1240, USA
Tel: 717/533-8845; Fax 717/533-8661; URL-http://www.igi-global.com

This paper appears in the publication, **Advances in Ubiquitous Computing: Future Paradigms and Directions** edited by S. Mostefaoui, Z. Maamar, and G. Giaglis © 2008, IGI Global

Chapter XII

Privacy Threats in Emerging Ubicomp Applications: Analysis and Safeguarding

Elena Vildjiounaite, VTT Technical Research Centre of Finland, Finland

Tapani Rantakokko, Finwe LTD, Finland

Petteri Alahuhta, VTT Technical Research Centre of Finland, Finland

Pasi Ahonen, VTT Technical Research Centre of Finland, Finland

David Wright, Trilateral Research and Consulting, UK

Michael Friedewald, Fraunhofer Institute Systems and Innovation Research, Germany

Abstract

Realisation of the Ubicomp vision in the real world creates significant threats to personal privacy due to constant information collection by numerous tiny sensors, active information exchange over short and long distances, long-term storage of *large quantities of data, and reasoning based on collected and stored data. An analysis of more than 100 Ubicomp scenarios, however, shows that applications are often proposed without considering privacy issues, whereas existing privacy-enhancing technologies mainly have been developed for networked applications and, thus, are not always applicable to emerging applications for smart spaces and personal devices, especially because the users and their data are not spatially separated in such applications. A partial solution to the problem of users' privacy protection could be to allow users to control how their personal data can be used. The authors' experience with mobile phone data collection, nevertheless, suggests that when users give their consent for the data collection, they don't fully understand the possible privacy implications. Thus, application developers should pay attention to privacy protection; otherwise, such problems could result in users not accepting Ubicomp applications. This chapter suggests guidelines for estimating threats to privacy, depending on real world application settings and the choice of technology; and guidelines for the choice and development of technological safeguards against privacy threats.*

Introduction

After having read a large number of scenarios of emerging Ubicomp applications (found in project deliverables and research publications which describe prototypes of smart spaces, smart personal devices, objects and their functionalities) and visionary future Ubicomp scenarios (found mainly in roadmaps), we concluded that most scenarios present a sunny, problem-free vision of our future. With the exception of the surveillance problem in some cases, most scenarios do not consider the privacy issues that the new technologies are likely to raise. For example, they do not discuss possible privacy problems due to conflicts between people's interests or personal curiosity.

The discovery that Ubicomp technologies raise privacy problems is not new; and research into privacy protection is actively going on, but after a state-of-the art review of work on privacy protection, we have come to the conclusion that most of this work deals with privacy protection in such network applications as m-commerce, Web browsing, virtual meetings, location-based services, and so forth, where users can be physically separated from their personal data. Even in these applications, no scalable solutions fully applicable in real life exist, and this lack of protection allows large-scale eavesdropping, as we know from the news (Web site of the American Civil Liberties Union and the ACLU Foundation, 2006).

The work on privacy protection in smart spaces and in connection with personal devices is even less mature than that concerned with network applications, while visionary Ubicomp scenarios suggest many situations in which confidential data and secrets occasionally can be discovered. When reading Ubicomp scenarios, however, we rarely found any discussions about the possible implications of a new technology for privacy, and even fewer descriptions of privacy protection measures. M. Langheinrich has collected a list of excuses why privacy protection is rarely embedded in new applications (Langheinrich, 2006), but such a practice can lead to the danger that problems appear after an application has already been developed and installed, and then either the users are left to suffer from privacy violation problems, or application developers are faced with the negative reactions of the users and the need to update the application. One recent example is a bus ticketing application in Helsinki which was storing data about travellers' routes. The application received bad publicity (criticism in the newspaper *Helsingin Sanomat* (Koponen, 2002)), and updating an already installed application would obviously be a costly operation. In cases where users' criticism is directed against an already installed application, which runs on non-reprogrammable microcontrollers (a common situation in the case of a commercial application), an application update can be very costly. Thus, embedding privacy protection in Ubicomp applications at the development stage would be beneficial for application developers.

The main emphasis in this chapter will be on possible problems rather than the benefits of new technologies and applications, because readers of Ubicomp papers usually encounter descriptions of benefits rather than descriptions of problems. The success of Ubicomp development also requires the understanding of possible problems, however, and safeguarding against them, including safeguarding against possible privacy implications. There is no doubt that the notion of privacy alters with time, so that with the invention of phones (and especially mobile phones), for example, physical distance from other people can no longer guarantee privacy. Similarly, with the development of cameras (especially digital cameras, with their capability for recording more views than their owners can sort through carefully), people have become used to seeing more details of other people's lives than was ever possible before.

There are very important differences between past and future technologies, however, which could change our lives more quickly than we could possibly adapt our understanding of the world, human behaviour, ethics and laws to the new technologies: first, past technologies were largely controlled by a human, whereas future technologies will be capable of automatic actions. Since it is much easier to notice a human observer than a tiny sensor, it will be possible to collect much more data without people being aware of it. Second, large-scale accumulation of data in a digital form will no longer require manual (slow) human work in order to connect information from different sources, so that it may be easier to assemble the full life story of a person in the future than it was to find scattered pieces of information in the past. Third, modern devices are smaller in size, more reliable and move closer to the human body than was the case in the past, and it is proposed that these could be embedded into clothes, watches or jewelry. Consequently, it will become easier to have always-on mobile devices, but more difficult to switch them off. Our perception of the privacy aspect known as the "right to be left alone," for example, has changed significantly with the invention of stationary phones and especially mobile phones, but it has still been preserved by the possibility for switching the phone off or not hearing it ringing when taking a shower or walking in a noisy place (and it is not easy to check whether a person did not hear a phone call or was simply not in the mood to answer it). Will one still be able to avoid undesired conversation in the Ubicomp future of embedded connectivity, or will society change so that people will not be offended or angry when their children, relatives or subordinates do not answer a call that they have evidently heard? How society will adapt to the capabilities of new technologies is an open question, but we think that technology developers should not rely on human nature changing quickly, and the results of deploying new technologies in computer-supported collaborative work (Bellotti, 1993) support this opinion.

This chapter first summarises the views of different researchers on what privacy is, after which it will briefly describe how Ubicomp researchers see the world of the future and what possible implications for users' privacy may not be safeguarded in the scenarios. After that, the chapter will present the authors' experiences of mobile phone data collection and users' opinions regarding their privacy expectations before and after data collection, which suggest that the privacy implications were underestimated before data collection. It will then present the state of the art in privacy-enhancing technologies and highlight the gaps that create privacy risks. After that it will suggest guidelines for estimating the threats to privacy, depending on real world application settings and on the choice of technology, as well as guidelines for the choice and development of technological safeguards against these threats.

Privacy Expectations in the Real World

It is suggested in the work of Lahlou et al. (Lahlou, 2003), that privacy protection requires an understanding of how new technologies change the ways that have developed in the physical world, where personal privacy is protected by the following borders (Bohn, 2005):

- **Natural Borders:** physical borders of observability, such as walls, clothing, darkness, facial expression (a natural border protecting the true feelings of a person)
- **Social Borders:** expectations with regard to confidentiality in certain social groups, such as family members, doctors and lawyers, for example, the ex-

pectation that your colleagues will not read personal fax messages addressed to you

- **Spatial or Temporal Borders:** expectations by people that parts of their lives can exist in isolation from other parts, both temporally and spatially, for example, a previous wild adolescent phase should not have a lasting influence on the current life of a father of four, or a party with friends should not affect relations with colleagues
- **Borders due to Ephemeral or Transitory Effects:** expectations that certain action or spontaneous utterances will soon be forgotten or simply unnoticed because of limitations on people's attention and memory

These borders are bi-directional, that is, people expect these borders not only to protect the person's feelings, appearance, actions, and so forth from the outside world, but also to protect the person from intrusions by the outside world. Physical borders are perhaps perceived as most reliable, as can be illustrated by how poker players control their faces, for instance, or by the custom of knocking on the closed door of somebody's private room or office. People also have a well-developed mental model of the limits of their own or others' ability to notice and remember details of what is going on around them. For example, people in a conference room usually expect that others' attention and memory will be devoted to the content of a presentation rather than to the auditory aspect. Concerning social and spatial borders, people perceive them as not so strong; for example, the likelihood of encountering the same people in different circumstances or of broken confidentiality is not negligible. In general terms, the stronger is the personal belief that a certain border is reliable, the more difficult it will be to adapt to its violation by a new technology. Experiments in the research area of computer-supported collaborative work suggest one example of the adaptation difficulty. In order to facilitate awareness and communication between colleagues, video cameras were installed in the offices of participants. Although this awareness proved to be useful, the experiments showed that people often act according to the "old" mental model of being reliably hidden by office walls (Bellotti, 1993).

Future Vision of the Ubicomp World and Problems with Privacy

A joint vision of Ubicomp researchers regarding the future world was formulated after reading more than 100 roadmap scenarios and research publications. This vision presents a world in which everything is connected and where any activity is possible in any place, supported by applications installed in the environment and in personal devices. Research activities have been devoted to supporting communications between family members and colleagues in different locations (e.g., between workplaces and homes (Aschmoneit, 2002; Dukatel, 2001; Jansson, 2001) and between moving people (Aschmoneit, 2002; Dukatel, 2001; ITEA, 2004), often via video links), and to supporting remote shopping (Dukatel, 2001), learning (Dukatel, 2001), and even remote health care (Bardram, 2004; ITEA, 2004). The future vision also pictures a very safe world, in which technologies ensure safe driving (ITEA, 2004; Masera, 2003) and safe control of home appliances (e.g., locking and unlocking of doors at home (Masera, 2003)), and help in finding keys (Orr, 1994) and toys (Ma, 2005). Technologies are also expected to help people to remember past events, both personal (Gemmel, 2004; Healey, 1998) and work-related (Aschmoneit, 2002), and to give reminders regarding current and future activities and duties (Kim, 2004).

A typical vision of the Ubicomp future involves technology caring for a person, correctly identifying that person's wishes and environment, and reacting to them appropriately. An attractive example of such a vision can be found in the Flying Carpet "Daily Life" visionary scenarios of the Mobile IT forum (Kato, 2004). It is worth noting, however, that the Flying Carpet scenario differs from many others in the sense that its interaction is human-initiated, whereas many other scenarios are more privacy threatening because they suggest that technology will be able to anticipate a person's needs and to act on behalf of that person (e.g., Ducatel, 2001).

There are many positive sides to the Ubicomp vision of future, but it does not seem realistic to assume that Ubicomp technologies will be problem-free. It is thus important to understand the possible problems and to safeguard against them whenever possible. In some cases, achieving safety is more important than protecting privacy. It has been observed, for example, that elderly people can trade their privacy for support and safety (Mynatt, 2000), and that the saving of people's lives through improvements in health care or detection of the location of emergency calls requires

Figure 1. Mobile IT forum, part of a "Daily Life" scenario from FLYING CARPET, Version 2.00 (Kato, 2004), page 4

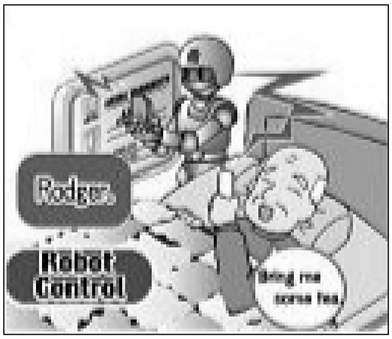

the storage of personal data which could potentially lead to violations of privacy. In such applications, the most important measure of privacy protection is first to store as little data as are needed for the application to function properly, and second, to ensure that the data cannot be easily accessed by unauthorised people. In many other applications, however, it is important not to trade off privacy for convenience in a blind fashion; for example, the personalisation of recommender systems or location-based services can be designed in more or less privacy-protecting ways. The main goal of this chapter is to point to certain important problems and to suggest methods for improving privacy protection in various Ubicomp applications, so that application developers can choose the most suitable methods.

For example, when reading Ubicomp scenarios one can rarely find a description of how access to personal data or to actuators can be controlled, whereas user-friendly access control is one of crucial factors determining the success of the Ubicomp concept. The example of mobile phones shows that personal data in a phone (such as photos, an address book, or a calendar) will in practice be available to anybody who picks up the phone, due to the inconvenience of password authentication, which takes place once, when the phone is being switched on, after which the phone remains in the "on" state for many days or weeks, unprotected. Although it has always been possible to look through somebody's address book, diary or photo albums in order to find the desired information, this has usually required visiting that person's room and searching through the items there, which may be difficult, at least for a person living in another place. Nowadays, personal mobile devices can store as much in the way of information and photos as several old-style address books, diaries, and photo albums (and will store even more when Personal Lifetime Store application scenarios (Gemmel, 2004) become a reality and when mobile payment logs can also be stored), but they are far less well protected, because they are not locked inside a house or a drawer. Personal mobile devices accompany their owners everywhere and can reveal large quantities of stored personal data, because the users often bypass the inconvenient security measures available for data protection (such as entering a password or rolling a finger across a fingerprint sensor). Since no convenient, user-friendly authentication has yet been developed, personal Ubicomp devices and non-personal smart spaces are likely to disclose their users' data and secrets, and we are now obliged to suggest how to reduce this risk.

The threats to privacy presented in this chapter are not really new, because the reasons for their existence (including conflicts of interest between people and organisations, human curiosity, envy, greed, and beliefs in one's own right to control others) are age-old problems. On the other hand, technology has changed the ways in which personal data can be disclosed.

The components of a typical Ubicomp application are shown in Figure 2. Each component can cause problems in its own way.

Figure 2. A generic view of an Ubicomp application: the thin arrows indicate information collection, transmission and storage; the thick arrows indicate information push.

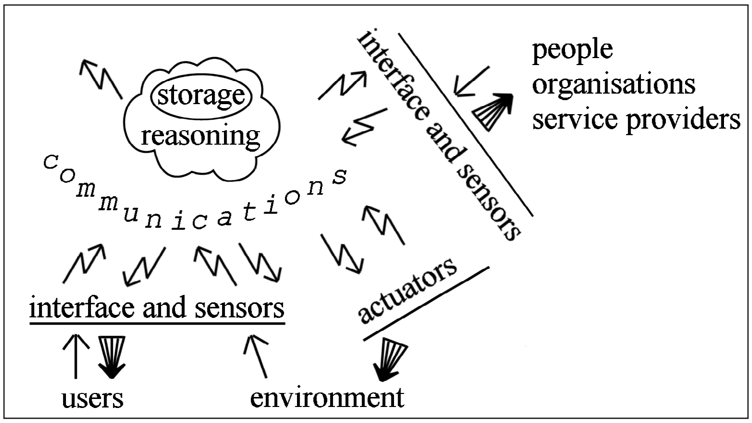

Privacy problems essentially fall into three major groups, the best-known of which concerns problems associated with information flow from the user, that is, due to the acquisition, transmission, and storage of personal data in large quantities. Most privacy-enhancing technologies (PETs) are being developed for the protection of personal data in networked applications, but new Ubicomp applications present new challenges. It has often been proposed, for example, that awareness between family members and colleagues should be supported via the transmission of video data, which violates traditional personal expectations regarding the notion that "if I am hidden behind a wall, I am invisible." Memory aids (personal (Gemmel, 2004; Healey, 1998) and recordings of work meetings (Aschmoneit, 2002)) imply the storage of raw video data, which violates personal expectations regarding the limits of human's attention and memory. There are two reasons for these problems. First, as work in the computer-supported cooperative activity domain has shown (Bellotti, 1993), humans are not accustomed to environments full of sensors, and continue to behave according to their expectations regarding their privacy in the real world. The second reason is the blurring of boundaries between "traditional" application domains. For example, work-related communications from home can intrude into one's personal life, and conversations on private matters from smart workplaces can be recorded automatically along with work-related conversations. In addition, sensors, which were traditionally used only in certain domains (e.g., physiological sensors associated with health care, video cameras for security purposes) have been suggested for use in other domains, such as entertainment. Since the traditional view of the entertainment domain assumes that its data are not very confidential (and consequently do not require strong protection measures), there is a danger of the disclosure of health problems detected by physiological sensors in the entertainment domain.

The second group of privacy problems concerns those caused by linkages between different kinds of data (mainly stored data). For example, it has been proposed that a personal memory aid should not record everything, but instead, it should measure the personal arousal level via skin conductivity sensors and other physiological sensors and record only the exciting scenes (Gemmel, 2004; Healey, 1998). Since none of proposed memory aid prototypes has good access control over stored data, these would allow a young boy's parents, for example, to find out easily which girl their son is most interested in. Physiological sensors also have been proposed for measuring the degree of approval of TV programmes (Nasoz, 2003; Palmas, 2001). In this case, the linking of personal physiological responses to information on TV programmes can facilitate the surveillance of citizens from the point of view of whether they support government decisions or not. The linkability problem is in general acknowledged, and PETs in networked applications aim at protection from such data linkability. In such applications as smart spaces and personal devices, however, the problem of data linkability has received less attention, and privacy problems with memory aids, for example, are usually discussed from two points of view: first, how to achieve agreement with the people recorded; and second, whether the police could search through the recorded data or not. The problem of avoiding the curiosity of family members is usually ignored. The dangers of data linkages are in general underestimated by researchers, as we have observed in the example of our own data collection system (see next chapter).

The third group of privacy problems comprises those caused by information flow towards the users, either because technology-initiated communication intrudes into personal life, because the content of the information can disclose private information, or because actuators fail (e.g., to open or close a door at home). Intrusions of technology into personal life can happen when an application interacts with people (e.g., reminds someone to do something) or when an application does not allow people to escape communication with others. Currently, it is easy for a person to say that he missed a phone call because the battery in his mobile phone was empty, or because of street noise, and so forth, but will it be as easy to avoid undesirable communications in the future, when communication is embedded in clothes and battery life is longer? Most parents have observed how their children miss phone calls or "forget" mobile phones at homes when they want to escape from their parents' control; and although such situations are harmful for the parents' nerves, it seems that in most cases, it is necessary for children to make their own decisions and take risks.

The content of information can disclose personal data in two possible ways: if it is delivered in the presence of other people and they hear (or see) the message (e.g., if a movie recommender application suggests that the users should watch adult videos in the presence of their children), or if the information contains data about people other than the user (as one can notice more details during the playback of a memory aid than during a live conversation).

This group of problems is the least studied of all, and PETs dealing with these problems are almost non-existent. What is also important about this group of problems is that technology-initiated communications can reduce user acceptance (users do not always like it when the technology makes the decisions) or hinder personal development. As the work of Nissenbaum (2004) shows, “the right to be left alone” is very important for personal development because people need relative insularity to develop their goals, values, and self-conceptions. Furthermore, if technology cares about personal safety and comfort and relieves people from many responsibilities (such as remembering to take one’s keys or to close a door), it becomes more difficult to develop responsibility in children. Children traditionally learn to be responsible for not losing keys, for doing their homework, for taking the right books to school, and for other small everyday tasks, but if all these responsibilities are shifted to Ubicomp technologies, what will replace them in growing children? To the best of our knowledge, the scenarios do not suggest any replacement. Instead, the role of children in many Ubicomp scenarios is limited to playing computer games. Research into computer-supported learning is an exception, but even there learning is mainly supported by augmented reality (Price, 2004), which is also a kind of game. One example of Ubicomp scenarios regarding children is the ITEA roadmap (ITEA, 2004) screenplay of “the Rousseaus’ holiday”-- holidays spent by a family consisting of a mother, father, and two children (10 and 13 years old). The screenplay describes how the family goes to a summer cottage “at the seaside on Lonely Island off the Mediterranean coast of France” and that “the kids are unhappy to leave home … because of the high-end, virtual reality video and gaming entertainment equipment, which was recently installed … in their house” (p. 134). If the roadmap leads us to a world in which school children are not interested in Lonely Islands, will we want such a world?

An Example of Unexpected Privacy Problems in Mobile Phone Data Collection

In our case study, the mobile phone usage data were collected with the goal of personalisation and context adaptation of mobile phone applications. The data gathered were a rough location estimate based on phone cell ID (granularity of cell ID-based positioning in our case ranged from several hundred metres to several kilometres), and phone usage data comprised of the duration of incoming and user-made phone calls and usage of different phone applications such as SMS typing, games, a calendar, whether the keyboard was in use or not, and so forth. No phone numbers were logged, nor any SMS contents or application entries, just the start and end of phone calls, the opening and closing of an application, the keyboard being in use, and so on. In addition, logs of Bluetooth activity were collected; each phone that participated in data collection was recording the IDs of all Bluetooth devices in its communication range. All the data items were time stamped with absolute time. Data were collected with respect to five users around the clock for five-seven days.

It is important to note that the users who participated in the data collection were informed about what data were being collected and gave their consent, largely because they did not expect any privacy problems. The users did not want precise location tracking (by GPS, for example) and did not want the content of their actions to be logged, but they allowed logging of the simple facts of actions taking place. It is also important to note that all our users were application developers in the field of context-aware computing, with several years of experience in developing Ubicomp applications. Thus, one would expect them to give their consent to data collection with a much better understanding of the consequences than the average person.

After the data had been collected and analysed, however, we found that when all the seemingly harmless components were linked together, they revealed a lot about the users. They actually allowed us to figure out what kind of person each user was: how communicative they were; whether they usually initiated communications or just received calls and SMS from others and reacted to them; whether or not they had regular routines in their life; whether they were hard-working people; whether they had an active night life; and so on. After discovering that the data told us a lot about the users, we asked them whether they expected such a result, and whether they like this result. For all the users but one, the result was quite surprising, and only one of them told us that he did not care whether other people could gain such information about him. Thus, the power and the unpleasant consequences of information linkage are largely under-estimated even by developers of Ubicomp applications.

We also asked the users to mark which parts of the information (not speculations regarding the user's personality, but lower-level data components) they could make available to their family members, which parts to colleagues and which parts they would not like to become available to a stranger who happened to access the information accidentally. Most of the users did not care much whether their family members or colleagues could access the data or not (although all but one user marked some information as "if family members knew it, it might occasionally be unpleasant" and some other information as "if colleagues know it, the situation might occasionally be unpleasant"). Only one user would allow strangers to access this information, however, and none of them wanted it to appear on the Web.

After that we asked the users' opinion regarding where more personal secrets can be discovered: in public places, at home or at work. Four users told us that a Ubicomp application installed in a home environment had high chances of discovering personal secrets, applications in a work environment and location tracking applications had medium chances, and applications installed in public places had the lowest chances. One user (the only person in our study who had an active night life) named location tracking as the most privacy-threatening, Ubicomp applications in public places (such as streets) and at home as moderately dangerous and those at work as the least dangerous. We realize that five users in a study is not a significant number, but we think that the results are interesting because all the subjects were well acquainted with the Ubicomp concept.

Our own data analysis confirms the users' opinions, because most of our conclusions were made after processing the data acquired in a home environment. We also observed that if the time stamps had not revealed absolute times, it would have been more difficult to analyse the data. For example, if the time stamps and location stamps had been encrypted or relative to certain application-dependent events, it would have been difficult to distinguish between the home and work environments and to make deductions about users' personalities. In our case location was shown only as a cell ID, so it was not very informative, but a very curious person would nevertheless be able to "decode" it.

The experiences with collecting IDs of Bluetooth devices in the communication range were also very interesting, as it was possible to find out from the phone logs when the neighbours of a test subject came home and when they went to sleep and to deduce something about their personalities.

The State of the Art in Privacy Protection

The term "privacy enhancement" has been used for more than a decade to represent technologies concerned with various aspects of Internet security. Privacy protection in Internet applications should be based on the main principles of privacy protection as listed in the Common Criteria for Information Technology Security Evaluation (anonymity, pseudonymity, unlinkability, and unobservability). A survey of privacy-enhancing technologies in the HiSPEC report of 2002 stated, however, that more effort had been invested in protecting user identities than personal data in the previous years (HiSPEC, 2002). Similarly, the PISA project in 2003 concluded that previous research efforts had mainly been concerned with the protection of users' identities, but not very much with users' actions (Blarkom, 2003). Since then, research efforts regarding the protection of personal data and user actions have increased, but they have mainly been concentrated on Internet applications.

Nevertheless, the PRIME study on the state of the art regarding privacy protection in network applications, carried out in 2005, has pointed out many performance problems and security weaknesses, and reached the conclusion that even the most recent techniques and tools are still far from providing a holistic approach to usable and secure anonymizing networks (Camenisch, 2005). It is worth noting that the conclusion refers to the current technology settings, not to future technology settings such as smart environments and personal memory aids. The goal of the PRIME project is to develop a framework for privacy and identity management in electronic information networks given current settings, and the project has made significant efforts in the areas of access control (the term "access control" in PRIME stands mainly for the access / release / processing of data by software methods, unlike the more traditional understanding of the term as the granting of access rights to a person), cryptography, communication infrastructure and user-side (allowing users to specify how their personal data can be used), and service-side (management of obligations) identity management. Access control research is concerned with developing policies for access control and a language for their description, in order to allow users to control the use of their personal information and to allow negotiations between different counterparts without revealing sensitive information.

The PAW project is a continuation of the privacy protection research with regard to the use of software agents and is working on cryptographic techniques and licensing languages (a description of what one is allowed to do with data during processing and what not). Licensing languages and machine-readable privacy policies are an active research area in which most of the research is concerned with the privacy policies of Web sites. A recently developed platform for privacy preferences (P3P) (Cranor, 2003) allows Web sites to convey their policies in machine-readable form, so that they can be checked on the user side and compared with user preferences. P3P does not actually force Web sites to stick to their promises, however.

The goal of the FIDIS project (Bauer, 2005) is to develop privacy-preserving methods of identity management for mobile wireless applications in current technology settings. The project has proposed a privacy diamond model for these settings, the main components in which are user, device, location and action, and has suggested that user privacy should be protected by hiding some of the links between these components of the model. The FIDIS project has also presented a classification of identity management systems (IMS), first as systems for account management (pure IMS, where the main goal is authentication, authorization, and accounting), second as systems for personalized services which need both user identity and profiles or log histories, and third as systems for pseudonym management, for example, in web services. Good practices for these systems were proposed, including separate access to the user authentication data, user account data and personal data (addresses, etc.) in identity management systems of the first type; and an architecture was developed for a mobile device security tool for creating partial identities and using them in wireless and wired networks.

To summarize, the projects listed above, and some others, mainly deal with privacy protection in network applications and, to some extent, with protecting personal data stored in personal devices. It is mainly proposed that the data stored in personal devices should be protected by means of encryption, but the inconvenience of the related security measures creates “large holes in security and privacy” (Caloyannides, 2004, p. 85), which is very dangerous considering the huge increase in the amount of personal data stored in modern mobile phones. Security and privacy problems affecting personal devices constitute a very challenging problem in general terms; on the one hand, the limited computational capabilities, battery life and screen size of mobile devices pose problems for developers of security methods, while on the other hand, the main burden of configuring and updating security settings and anti-virus software is being placed on the owners of these personal devices, who often have neither the necessary special education, nor the time or enthusiasm to do that.

Research into privacy protection in such emerging domains as smart environments and smart cars is in its infancy, and only generic guidelines have been developed. The work of Langheinrich et al. (2001), for example, suggests how the fair information practices (listed in current data protection laws) can be applied to Ubicomp applications, and shows how difficult it might be to apply them. The fair information practices state, for instance, that the user must have access to the data about him that has been stored, and the right to change details that are wrong. In a Ubicomp future, however, it will not be easy for users to find all the items of data about them that are stored in the network and in the personal devices of surrounding people (let alone check them). Moreover, some data processing techniques (such as neural networks) store user models in a form that is difficult to interpret.

The work of Hong et al. (2004) proposes high-level privacy risk models based on two aspects: first, the social and organisational context in which an application is embedded (Who are the data sharers and observers? What kinds of personal information are shared? What is the value proposition for information sharing, its symmetry, etc.?), and second the technological aspect (How is the collection, storage and retention of personal data organized? Who controls the system? Is there any possibility to opt out?). This is close to our understanding of privacy threats, but we suggest that other aspects should also be taken into account, especially the probability of accidental information flow (not intended by the designers). Furthermore, this work mainly suggests guidelines for risk estimation, not for safeguards.

The work of Lahlou et al. (2003) focuses “on the specific issues of the data collection phase” (Forward, p. 2) and proposes high-level guidelines. One of the most important guidelines is to minimize data collection. Such generic design guidelines as “think before doing” and “understand the way in which new technologies change the effects of classic issues” (i.e., existing solutions in the physical world) (p.3) can be applied in other spheres as well as data collection, but these are very generic design guidelines.

To summarize, most of the research into privacy protection is concerned with protection of the information flow from users, whereas other privacy aspects have not received much attention from researchers.

Gaps in Privacy Enhancing Technologies

For most Ubicomp scenarios to work well, advanced privacy-protecting safeguards, which do not yet exist (although research into them has started), will be required. We suggest that the most important safeguards are the following:

- **Intelligent Reasoning Capabilities**: advanced artificial intelligence algorithms capable of recognizing sensitive data in order to avoid recording or publishing it, for example, algorithms capable of intelligent online summarizing of audio recordings (online conversion of a meeting audio stream into a text document, including only working discussions), algorithms capable of detecting that persons in a video or photo are naked or kissing, algorithms capable of adaptation to the user's ethics and culture (a photo of a Muslim woman with her head uncovered is private, while for the majority of Finnish women this would be nothing special) and so on. To some extent these capabilities can be implemented as common-sense rules, such as "if a person is alone, or if there are only two persons in a room, the probability of discovering confidential data is higher than if there are a larger number of people." In addition, algorithms for detecting unusual patterns of copying and processing of personal data are needed (e.g., if a new back-up is made soon after the previous back-up it may indicate data theft, and an alarm should be given), because these would also be of help when a person authorized to work with the data is dishonest, unlike other access control methods, which work mainly against outsiders.
- **User-Friendly Security**: advanced access control and security methods, such as frequent unobtrusive context-aware authentication of users. Different user verification methods should be chosen, for example, depending on the application that a user wants to access, or on the user's location and behaviour; access to a calculator application should not require user effort (Stajano, 2004), whereas access to personal memory aid data should be allowed only to the data owner. Thus, the current "once and forever" password-based user verification on mobile phones, which facilitates unauthorised use when the owner is in another room, for instance, should be replaced with continuous unobtrusive user verification, for example, based on user behaviour or voice recognition, and on stronger authentication methods if unobtrusive authentication fails but access to sensitive data continues to be requested. In general terms, we suggest that security should be a fairly effortless matter for users (e.g., updates of anti-virus software should be system-initiated and happen at convenient times) and should be enforced. We suggest this by analogy with control over technical conditions in personal cars and the security enforced with regard to financial operations, because future Ubicomp scenarios envision personal devices that perform life-critical tasks (health monitoring and health

care (Bardram, 2004; ITEA, 2004), financial tasks, and identity management (Ducatel, 2001). A malfunctioning personal device could fail to notice a health crisis on the part of the device owner, or fail to communicate this to the doctors, for example, and it could also create threats to other people; for instance, if it sends a lot of spam and malware to surrounding personal devices it can significantly slow down their operation and hinder their performing of the tasks required of them. Work on user-friendly authentication is an emerging research area. Current work is mainly concerned with biometrics, which is not a perfect solution, because of possibility of spoofing biometric sensors. Thus, we suggest that biometric modalities which carry a high danger of identity theft (e.g., fingerprint and iris) should be used cautiously and only with aliveness detection, and that a fusion of several not so privacy-threatening biometric modalities (such as voice, gait, or behaviour) should be used as a primary or complementary means of authentication.

- **Communication protocols which do not use Unique Device Identifiers**: it is easy to link a device ID or a smart object ID to a user and to track the user's actions. Communication protocols which hide the very fact of communication would be an ideal case, because a lot of information can be acquired by tracking who communicates with whom. For example, it can be concluded from the fact that a person has started to visit the Web page of a certain bank that that person has opened an account at this bank, and a false request to update a recently created account in this particular bank has higher chances of succeeding than when sent to an established client of the bank, or to a client of another bank. This safeguard has the drawback that it would hinder the discovery of users with malicious intentions or with malfunctioning personal devices sending out spam and viruses, but this problem can be partially solved by means of good firewalls and anti-virus software, which would protect against malware (see "user-friendly security" bullet). In cases where the detection of users' IDs is important, these communication protocols cannot be used (e.g., some applications can require the using of IDs in communication protocols), but their usage should not be a common practice because large-scale logging of everybody's actions is not likely to improve security in society, as the famous security expert Bruce Schneier pointed out (Schneier, 2005, 2007).
- **Secure ad-hoc communications:** if a device owner enables ad-hoc Bluetooth communications, for example, the sending of a large number of requests to this device can slow down its operation or even exhaust the battery. This is not a direct threat to privacy, but it might engender privacy if the encryption of personal data becomes delayed, and it is definitely a violation of the "right to be left alone" in cases where the user cannot simply ignore incoming spam because he is expecting an important message.

- **Encryption for untrustworthy platforms**: not all functions can be executed in encrypted form. Instead, it is common to decrypt the code and the data before execution, which allows spying.
- **Unified, concise user interface methods of maintaining user awareness** about the functionality of the application and its privacy threats, possibly in graphical form (e.g., similar to road signs). A warning about video cameras is currently placed on the doors of shops (although with no information as to whether video data is stored or not, or for how long), but other Ubicomp technologies would require similar icons.
- **More detailed transparency tools for awareness in average non-technical users** regarding security risks, the correct usage of anti-virus and firewall applications, the dangers of data collection and the accepting of ad-hoc messages from unknown devices, and so on. Since users do not want to spend much time on security education, these tools and their user interfaces should be really intelligent and work "just in time." Currently, it is too easy for users to make mistakes (it is often too easy to ignore important questions that all look the same, and click yes without even reading a security-related question). For example, the current practice of asking users whether they agree to "accept temporarily, for this session," a security certificate regarding a certain Web site is not really helpful, because the same question is used for all Web sites, and it is not linked to other data regarding the website in question.
- **Recovery means**: first, if somebody's personal data was compromised (e.g., a fingerprint was forged), it is necessary to switch quickly and easily to a new authentication procedure in all applications (home, work, smart car, banking, etc.), unlike the current situation, in which recovery from an identity theft requires significant efforts on the part of the victim and can harm that person's reputation. Second, if a personal device is lost, the personal data contained in it can be protected from strangers by security measures such as data encryption and strict access control. However, it is important that the user does not need to spend time customising and training a new device (so that denial of service does not occur). Instead, the new device should itself load user preferences, contacts, favourite music, and so forth, from a back-up service, probably a home server. We suggest that ways be developed to synchronize data in personal devices with a back-up server in a way that is secure and requires minimal effort from the user.

Dimensions of Privacy Threats Analysis

Privacy risks fall into two major groups, the first of which is application domain-dependent risks, which depend on the personal or organizational activity being supported. Health data, for example, are considered sensitive, and designers of applications for hospitals are obliged to follow corresponding privacy protection regulations. Second, privacy risks are caused by a mismatch between personal expectations regarding current privacy levels and reality, which do not depend on the application domain. If a person perceives his current situation to be a private one (e.g., being alone at home) but in fact is being monitored, the chances that personal secrets will be discovered are higher than if the person perceives the current situation as public (e.g., giving a talk at a large meeting) and takes care of his own privacy.

Consequently, we suggest the following dimensions for the analysis of privacy threats:

- Real-world dimensions:
  - People's personalities
  - People's activities
  - The environment where an activity takes place
- Dimensions of technology functionality:
  - Information flow
  - Computer control level vs. personal control level
  - Balance between technology aspects (storage and communication vs. reasoning capabilities and control level)

Real World Dimensions

People's personalities are important because the notion of what is considered private and what is not depends on the person and the situation (context) (Nissenbaum, 2004). For example, the chances that a married man's personal data will accidentally be accessed by his wife or children are fairly high, even though secrets withheld from family members are not unusual; for example, parents often prefer to keep children unaware of the existence of adult videos at home, in order to prevent them from watching these while the parents are away. Personal activity is obviously an important dimension for privacy risk analysis because an activity consumes and produces a flow of information; for instance, large quantities of financial data are involved in paying bills, and health and identity data are involved in a call to a doctor. The environment is an important dimension (one which unfortunately is not always considered) because people's mental models of current privacy levels are based on traditional perceptions of their environment (e.g., "now I am alone in my office, so that nobody can see me") and people behave more or less freely depending on their estimation of current privacy levels. We suggest that applications should take the following into account:

- **traditional perceptions of the environment** (e.g., perception of the home as a private environment; perception of a wall as a non-transparent object, perception of a street as a public place)
- **common activities in the environment** (e.g., in an office people usually work)
- **other probable activities in the environment** (e.g., calling a doctor or flirting with a colleague in an office environment). Previous guidelines for the estimation of privacy threats (Hong, 2004) took account of the activity dimension mainly in the sense of the primary user activity supported by the application, but secondary activities are also very important.

Privacy threats coming from real-world settings can be roughly categorized as high, medium, and low in intensity. We suggest that application developers should always consider the privacy risk to be high when the application can run in the presence of children (which concerns most home-domain applications). Furthermore, we suggest that application developers should not give parents unlimited power to check and control what their children are doing. Instead, the children's privacy should be protected carefully, because they need this privacy for their personal development (Nissenbaum, 2004).

We suggest that high-intensity threats exist in connection with activities dealing with health care, finance, and communication between family members and close friends. High threats appear in the home environment, first because people perceive it as private and behave freely, and second because the security of home computers and personal devices is to a large extent the responsibility of their users, whereas many people (elderly people and children especially) do not have the education, skills or in many cases the desire to take care of the security of personal Ubicomp applications, which makes them vulnerable to all kinds of security faults. High-intensity threats also exist in an office environment, because on the one hand people cannot avoid dealing with private issues at work and are highly dependent on their work, and on the other hand, they are not free to decide on the environment in which they have to work, whereas organizations invest a lot of money in the development and installation of Ubicomp applications in workplaces. It is, thus, quite probable that Ubicomp applications will be deployed in workplaces sooner than in homes.

Figure 3. Guidelines for evaluation of real-world privacy threats caused by certain environments and activities

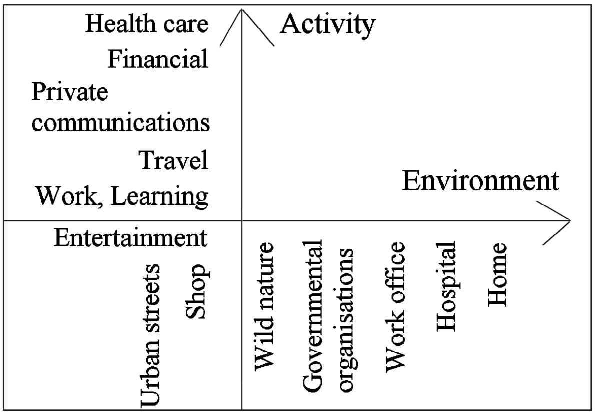

Medium threats to privacy appear in connection with shopping (increasing competition between retailers can lead to advertisements that are targeted at personal preferences and to hunting for personal data), learning and mobility activities (by mobility we mean travelling within a city as well as on holiday or for one's work), and relatively low-level threats are associated with entertainment activities.

Our informal grading of the dimensions of the privacy threats caused by real-life activities and the environment is presented in Figure 3.

Technology Choice Dimensions

Information flows start from data collection performed by sensors. The most popular sensors in Ubicomp scenarios are audio, video, positioning, physiological, safety, and comfort sensors, together with those used for logging human-computer interactions.

Physiological sensors are most dangerous from the privacy point of view, because they detect what is inside a person's body, that is, they "break into" the most private sphere. These sensors are the basis for building health care applications, where strict rules for the protection of health data exist. Ubicomp scenarios, nevertheless, suggest that these sensors could be used for purposes other than health and fitness. Detection of a person's mood and emotions is an active research area (Nasoz, 2003), and suggested applications include the detection of interesting scenes for automatic audio and video capture for lifetime personal stores (Gemmel, 2004; Healey, 1998) and the estimation of a user's preferences for TV programmes (Palmas, 2001). If physiological data are linked to the content of TV programmes and to the presence of other people, however, personal feelings become dangerously “naked” and can reveal to parents such facts as who their child is in love with, or else they can be used by governments for monitoring the loyalty of citizens. Physiological sensors can also detect health problems, but such data will not be properly protected, because the data protection requirements in the domain of TV personalization are not very strict.

Video and audio sensors violate natural privacy-protecting borders such as walls, and video cameras can reveal a lot more than audio sensors. In Ubicomp scenarios, they are suggested for use in real-time communication between people and for helping parents to monitor their children, for instance, by logging potentially dangerous situations (Ma, 2005). Second, such sensors have been suggested for memory augmentation, for example, the recording of work meetings (Aschmoneit, 2002) or personal memory aids (Gemmel, 2004; Healey, 1998). The first type of application “breaks the walls,” while the second type violates people’s belief in the limits of others’ attention and memory.

Biometric sensors have mainly been suggested for access control, and carry a danger of identity theft. Safety and comfort sensors (temperature, light, car acceleration etc.) can reveal users’ personalities and often initiate information push; for example, they may issue reminders to switch the stove off or employ actuators to do it automatically. This is beneficial for people suffering from dementia or for families with babies, but if teenagers are assumed to be as irresponsible in caring about home safety as babies, there may be little opportunity left for them to develop a sense of responsibility.

The application control level denotes how much technology does on behalf of its users. An application that reminds its user to take pills in the event of high blood pressure, for example, has a high control level because it initiates the measuring of blood pressure and a dialogue with the user. Such a dialogue may annoy the individual or reveal personal health details if it happens at the wrong moment or in public. An application which filters shopping advertisements according to user preferences also has a high control level, because the user can never know about certain shopping alternatives if they are filtered out. (An important question for such applications is who sets the filtering rules and how they can be prevented from favouring a particular shop.)

With more extensive information collection, transmission and storage capabilities and higher control levels, technology poses more privacy threats. Most Ubicomp scenarios involve application-dependent information storage and a lot of wireless communication (between objects, people, and organizations). We suggest that significant threats to privacy can arise if technology penetrates walls and the human body, for instance, by using physiological, video and/or audio sensors. Significant threats are also likely to be caused by high control levels (i.e., the capability of a technology to act on behalf of a person, e.g., to call an ambulance in an emergency) or by biometric sensors (due to the possibility of identity theft).

Figure 4. Guidelines for evaluating privacy threats caused by technology choices

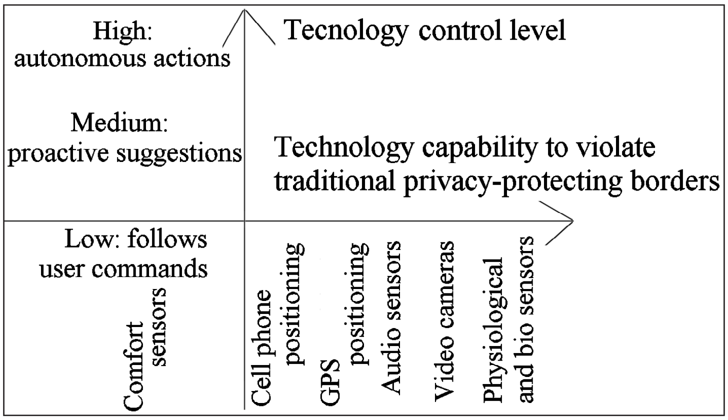

We also suggest that privacy threats should always be regarded as high when the linkage of data from several sources is possible, for example, when either of a lot of data about one person can be aggregated (as in most personal devices), or certain data about a large number of people. We suggest that the dangers of information linkage are often under-estimated, as we have observed in the case of our data collection system.

Medium threats are associated with positioning sensors (without time stamps they provide location data, but not much activity data, whereas location plus time information is a much greater threat to privacy) and with a medium level of technology control (the capability to make proactive suggestions, e.g., to issue reminders). Fairly low threat levels are associated with a low level of control (e.g., ranking advertisements according to criteria explicitly set by the user) and with comfort sensors (lighting, heating, etc.).

We would like to emphasize that threats to personal privacy are very often caused by mismatches between the application control level and application intelligence, and particularly by the fact that the technology is already capable of storing and transmitting a lot of data, but is not capable of detecting which data it should not store or transmit (with the exception of predefined data categories such as health and finance). In order to ensure “the right to be left alone,” however, and to prevent the accidental disclosure of confidential data, for example, via an audio reminder to take medicine when the user is in somebody’s company, it is very important that the intelligence of an application should correspond to its level of control (in other words, to its level of autonomy: what technology can do on its own initiative). Another example can be found in (Truong, 2004), which presents scenarios of Ubicomp applications made by end users, where one of the users suggested automatic recordings of parties in his home. If such an application is deployed in a large home and records two persons discussing personal matters in a room without any other guests, for example, it can lead to privacy problems. These would not appear if the application were intelligent enough not to record such a scene.

Guidelines for Safeguarding Against Privacy Threats

Estimates of the threats to privacy created by the combining of real-world settings and technology choices in certain popular Ubicomp scenarios are presented in Figure 5. Since the scenarios do not describe implementation details, the estimates are only approximate. The threats in the "safe driving" application scenario, for example, depend on data storage (e.g., whether a time-stamped log of speed, acceleration etc. is stored or not) and data exchange (e.g., between cars driving behind the other),

Figure 5. Examples of levels of privacy threats in popular scenarios

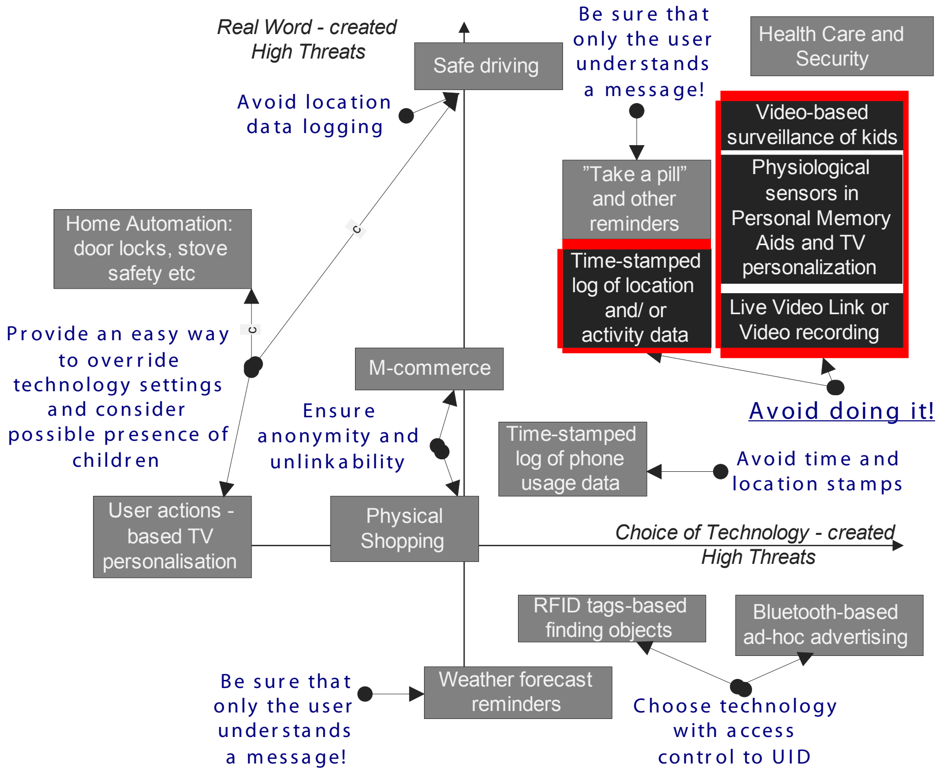
Real Word - created
High Threats
Safe driving
Be sure that only the user understands a message!
Health Care and Security
Avoid location data logging
Video-based surveillance of kids
Physiological sensors in Personal Memory Aids and TV personalization
Live Video Link or Video recording
"Take a pill" and other reminders
Time-stamped log of location and/ or activity data
Home Automation: door locks, stove safety etc
Provide an easy way to override technology settings and consider possible presence of children
M-commerce
Ensure anonymity and unlinkability
Avoid doing it!
Time-stamped log of phone usage data
Avoid time and location stamps
User actions - based TV personalisation
Physical Shopping
Choice of Technology - created
High Threats
RFID tags-based finding objects
Bluetooth-based ad-hoc advertising
Weather forecast reminders
Be sure that only the user understands a message!
Choose technology with access control to UID

but a high application control level is in any case a threat to privacy, because the technology might be wrong, and because the users don't always accept its superiority. Similarly, "issuing reminders about the weather forecast for the destination when on a journey" presents privacy threats because it is a form of technology-initiated interaction. What if the reminder is given when the user is in the company of a person whom he would prefer to be unaware of his journey?

When reading Ubicomp scenarios, we have not found any for applications which do not have either high technology risks, or high real-world risks, or both. In fact, most scenarios fall into the category of high technology risks. We suggest that if an application implies high technology risks, these should be reduced by lowering the control level of the technology, choosing the sensors differently and reducing the linkability of the data and by other applicable methods (see below).

By lowering the control level of technology, we mean that applications should ask the user's permission before taking potentially privacy-threatening actions, for instance, for video and audio recording. By a different choice of sensors, we mean that same kind of data can often be acquired in many ways, each of them presenting different privacy threats. Movie recommendation applications, for example, need user feedback data, and the ways of obtaining it include the use of physiological sensors, the analysis of facial expressions, speech recognition, monitoring of the noise level in a room, and monitoring user actions such as fast forward scrolling (which is the safest in terms of privacy). Even if fast forward scrolling and noise level monitoring might not give as good results as physiological sensors (which have not actually been tested), they should be preferred because they pose less of a threat to privacy. From our data collecting experiences, we would argue that what we can tell about a person through the linkage of different kinds of data it is frequently under-estimated. We regard reducing data linkage as very important, and suggest that absolute time stamps should be avoided; that is, data should be stamped with the time relative to the application and as much real-time data processing should be done as possible.

Furthermore, we suggest that since applications with both high threats due to real-world settings and high threats due to technology settings require advanced safeguards (such as intelligent reasoning capabilities or user-friendly security), which do not yet exist, such applications should be deployed only in domains with strict legal regulations, such as healthcare or banking, and then only with a fairly low level of technology control. In other domains we suggest that the deployment of such applications should be postponed until the technology becomes more intelligent. For example, we suggest that the use of physiological and video sensors and data stamped with absolute times should be avoided unless it is critical for the preservation of life and security. The suggestions made above do not apply to cases where the technology performs its tasks reliably and the users do not perceive the privacy problems as being important; for example, elderly people may be willing to trade off privacy against the gaining of support in time, and babies do not care about privacy at all.

In addition, we suggest the following good practices:

- **Real-time data processing**: select algorithms and hardware capable of processing data immediately in real time (performing real-time feature selection, or finding answers to predefined "pattern exists or not" queries), so that the storage of raw data (even temporarily) is avoided;
- **Encrypted or relative location stamping and time stamping:** For example, instead of investigating the dependence of high blood pressure on absolute time, an application should stamp the data relative to the moment of taking a pill or calculate the average time when the user's blood pressure was above a given threshold;
- **Data deletion or editing after an application-dependent time**: For example, when a user buys clothes, all information about the material, price, designer, and so forth, should be deleted from the clothes' RFID tags. For applications that require active RFID tags (such as finding lost objects (Orr, 1999), the RFID tag should be changed so that no links are left between the shop database and the personal clothes. Similarly, the location of an emergency call does not require the storage of long-term location data, so that this should be avoided;
- **Data processing in a personal device instead of sending data to the environment**: Instead of submitting a query with personal financial preferences to a shop in order to find suitable products, for example, the application should submit a more generic query, even at the cost of an increase in data filtering in personal devices, and anonymous payment procedures should be used whenever possible.
- **Choice of communication technologies which do not use permanent hardware IDs in their protocols, or at least have control over access to these IDs**, and which allow the communication range to be controlled. The current situation with Bluetooth communication, for example, is that if a device owner enables ad-hoc communication (in order to use the full range of possible applications), the device responds to each request with its ID, allowing user tracking even over walls, due to the fairly large communication range that is beyond user control.
- **Detection of hardware removals and replacements**: Users are currently not warned about replacements/ removal of attached sensors or memory cards when devices are in the "off" state, thus making physical tampering easier (Becher, 2006). Since personal devices will be monitoring a user's health in the future (Bardram, 2004; ITEA, 2004), unauthorized replacement of sensors could result in a death if they failed to detect a health crisis.

- **Transparency tools**: These are user-friendly ways to warn users about possible privacy violation problems which might result from the technologies deployed around him/her and ways to configure technology settings easily. For example, users might prefer to sacrifice some of the benefits of an application for the sake of anonymity, to reduce the level of control of applications or adjust the way in which incoming advertisements are filtered (if advertisements which are considered uninteresting by the application are completely removed, this carries a danger that the user will never hear about some options). One solution could be to have several "privacy profiles" in devices, so that each profile defines which groups of applications and means of communication are enabled and which not in different settings. Users would then just need to switch between profiles instead of dealing with a bundle of options with the risk of forgetting some of them. Our own experiences with data collection have shown that since even Ubicomp application developers do not fully understand the possible consequences of their data collection, transparency tools should be really carefully designed.
- **Means of disconnecting** gracefully: Users should be able to switch an application or device off completely, or to switch off some of its functionalities in such a way that other people do not take it as a desire by the user to hide, and in such a way that the device is still usable (e.g., users should be able to check calendar data while having the communication functionality switched off).

Conclusion

We have presented here an analysis of Ubicomp scenarios from the point of view of possible implications regarding privacy and have considered the state of the art in research into privacy protection, which does not allow safeguards to be provided against all possible problems. Recent news reports suggest that large-scale surveillance by means of ubiquitous technologies (the Internet and phones) has already started (Web site of the American Civil Liberties Union and the ACLU Foundation, 2006). The analysis of Ubicomp scenarios does show, however, that privacy protection is not yet considered a necessary design requirement, which can lead to a lack of user acceptance.

A typical approach to privacy threat analysis is to estimate the sensitivity of data that have been collected and stored, which depends on the application domain (e.g., health care data are considered sensitive) and on the consumers of the information (Hong, 2004). We suggest that privacy protection should also depend on which borders of real-life privacy are violated by the technology, because the likelihood of acquiring sensitive data accidentally is high if the technology penetrates through supposedly reliable physical borders. Furthermore, we suggest that privacy protection should consider not only information flows from users, but also information flows towards users.

The design guidelines for the estimation of privacy threats and for privacy protection in emerging Ubicomp applications have been proposed after a thorough analysis of Ubicomp scenarios, observations made during long-term runs with Ubicomp applications in a work environment (Bellotti, 1993) and our own experiences. Our guidelines are intended to protect individuals both from regular leakage of confidential data (such as location tracking data) and from the accidental discovery of sensitive data, for example, the discovery that two guests at a party had a heated discussion on a balcony. The effectiveness of such guidelines is very difficult to evaluate, due to the rare occasions on which such events happen and the fact that attempts to "capture" them would be unethical. We are not aware of any work presenting results on how certain privacy-protecting guidelines actually protect or disclose real secrets.

Our experiment with phone data collection, nevertheless, convinced us that since it is difficult to over-estimate what kind of discoveries an application can make, developers should be very cautious and take care to protect users against infringement of their privacy by various categories of interested persons and organizations ranging from the limited number of experienced hackers up to the large numbers of curious family members, relatives, colleagues, neighbours and so on, who luckily are most probably not endowed with such advanced computer skills. The guidelines and safeguards are proposed in order to help application developers decide which problems they should pay attention to and choose the most appropriate safeguards in relation to the application and the device capabilities. Implementation of some of the proposed safeguards would require a significant increase in the computational capabilities of personal devices, but such notable hardware improvements have been achieved recently, that it is likely to become possible in the near future to dedicate more memory to data processing algorithms instead of only to data storage. In cases where the capabilities of personal devices are insufficient for the desired safeguards, we suggest that users should be made aware of the possible problems (proper transparency tools should be developed for non-technical users) and allowed to choose a trade-off between the benefits and problems of the applications.

Ubicomp technologies can help to make life better if they are accepted by users, but this acceptance will be jeopardized if the problems created by the new technologies are not analysed and minimized. One of the important benefits of Ubicomp technologies will be to increase the security of individuals and society as a whole, for instance, making it possible to locate an emergency phone call, which could help to save the users' lives, or to access descriptions of crimes in remote locations and to compare them, which could help to find criminals. Similarly, access to a patient's lifelong health record could help to reveal an allergy and save the person's life. In general, new technologies provide support for a safer and more convenient life and for communications, so that it would be possible to access everybody and everything (family members, doctors, services, etc.) from any place and any time. The current situation is, nevertheless, such that the benefits are emphasized more than the possible problems, and thus we would like to emphasize the problems in this paper. New technologies can have implications for privacy with respect to the surveillance of citizens by governments and surveillance between people, for example, control exercised by parents or spouses over the activities of their family members. Although in some cases surveillance is clearly undesirable (e.g., parents do not want their children to be able to discover the prices of their purchases easily, to know which videos they watch or to see all their photos), it is an open question whether the surveillance of citizens by a government and the surveillance of children by their parents can increase the safety of society as a whole. The surveillance of children can help to save them from abuse, traumas or drug addiction, but in many cases such surveillance is not likely to do any better than the old-style trust and love in a family. Thus, it may be better when developing new technologies to aim at detecting when children are in real danger rather than at simply providing their parents with means of control all their actions? Although it is easier to develop technology by which parents can monitor their children, it might lead to their growing into irresponsible or helpless people.

Regarding whether the surveillance of citizens by a government can increase safety in society as a whole, we would like to cite the opinion of the famous security expert Bruce Schneier, who says in his essay “Why Data Mining Won’t Stop Terror” (2005) that “we’re not trading privacy for security; we’re giving up privacy and getting no security in return.” Schneier continues to debate over the idea of trading off privacy for extra security in later articles. For example, in “On Police Security Cameras: Wholesale Surveillance” (2007), he says that “the effects of wholesale surveillance on privacy and civil liberties is profound; but unfortunately, the debate often gets mischaracterized as a question about how much privacy we need to give up in order to be secure. This is wrong.” Schneier suggests that, although the police should be allowed to use new technologies to track suspects, data on people who are not currently under suspicion should not be stored. The decision of the society regarding the bus ticketing application in Helsinki was in line with this opinion, that is, society decided against trading off privacy for security. In many cases, however, new technologies can help to increase safety without threatening privacy. The possibility to locate an emergency phone call can help to save the users’ lives, for example, but the threats to users’ privacy can be minimized first by keeping only short-term, recent location data, and second by strict control over access to this data.

We suggest that the most important safeguards are an appropriate balance between the level of technology control and its level of artificial intelligence (how advanced the reasoning is and how the access control methods are implemented), an appropriate choice of sensors (sensors with powerful capabilities for violating natural privacy-protecting borders should not be used wantonly), and other hardware (such as communication chips with a configurable communication range and access control to their ID), the prevention of data linkability by avoiding absolute time stamps and location stamps, especially in applications which cannot provide user anonymity (such as smart spaces and personal devices), user-friendly security and user-friendly system configuration methods. The list of the proposed safeguards is not exhaustive and could well change with the development of new technologies. Novel application scenarios or the unpredictable use of new technologies, for example, could introduce new threats to privacy, which could require more safeguards. On the other hand, if methods of reliable unobtrusive biometric recognition with aliveness detection can be developed for mobile devices in the near future, this will significantly improve the protection of personal data and make some of our recommendations outdated. However, since our analysis is based on application scenarios and roadmaps for Ubicomp technology development and for the development of privacy-enhancing technologies, we believe that our recommendations for the evaluation of privacy threats and safeguarding against them will be valid for as long as the scenarios analysed here are valid, and for as long as gaps in privacy-enhancing technologies pointed out here continue to exist. Since one fairly common reason for privacy problems in these scenarios is an insufficient level of system intelligence for the complexity of the tasks, and since computer capabilities for data collection, storage and transmission are growing faster than the intelligence of data processing algorithms, protection against privacy violations is likely to remain an important problem in the future.

Acknowledgment

This Article is based on research supported by the EU project SWAMI: Safeguards in a World of Ambient Intelligence (IST-2004-006507)

References

Aschmoneit, P., & Höbig, M. (Ed.) (2002). *Context-aware collaborative environments for next generation business networks: scenario document* (COCONET deliverable D2.2). Telematica Institute.

Bardram, J. E. (2004). The personal medical unit - a Ubiquitous Computing infrastructure for personal pervasive healthcare. In T. Adlam, H. Wactlar, I. Korhonen, (Ed.), *UbiHealth 2004 - The 3rd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications*.

Bauer, M., Meints, M., & Hansen, M. (2005). *Structured overview on prototypes and concepts of identity management systems* (FIDIS Deliverable D3.1). Retrieved on March 15, 2006, from http://www.fidis.net/fileadmin/fidis/deliverables/fidis-wp3-del3.1.overview_on_IMS.final.pdf

Becher, A., Benenson, Z., & Dornseif, M. (2006). Tampering with motes: real-world physical attacks on wireless sensor networks. In J. A. Clark, R. F. Paige et al. (Ed.), *The Third International Conference* on *Security in Pervasive Computing* (pp. 104-118).

Bellotti, V., & Sellen, A. (1993). Design for privacy in Ubiquitous Computing environments. *Proceedings of the The Third European Conference on Computer Supported Cooperative Work (ECSCW'93)* (pp. 77-92). Kluwer.

Blarkom, G. W. van, Borking, J. J., & Olk, J. G. E. (Ed.). (2003). *Handbook of privacy and privacy-enhancing technologies: the case of intelligent software agents*, TNO-FEL, The Hague.

Bohn, J., Coroama, V., Langheinrich, M., Mattern, F., & Rohs, M. (2005). Social, economic, and ethical implications of Ambient Intelligence and Ubiquitous Computing. In W. Weber, J. Rabaey, E. Aarts (Eds.), *Ambient Intelligence* (pp. 5-29). London: Springer-Verlag.

Caloyannides, M.A. (2004). The cost of convenience: a Faustian deal, *IEEE Security & Privacy, 2*(2), 84 – 87.

Camenisch, J. (Ed.). (2005). *First annual research report* (PRIME deliverable D16.1). Retrieved on March 12, 2006, from http://www.prime-project.eu.org/public/prime_products/deliverables/rsch/pub_del_D16.1.a_ec_wp16.1_V1_final.pdf

Cranor, L. F. (2003). P3P: making privacy policies more useful. *IEEE Security and Privacy, 1*(6), 50-55.

Ducatel, K., Bogdanowicz, M., Scapolo, F., Leijten, J., & Burgelman, J.-C. (2001). *Scenarios for Ambient Intelligence in 2010*. Institute for Prospective Technological Studies (IPTS), EC-JRC, Sevilla.

Gemmel, J., Williams, L., Wood, K., Lueder, R., & Bell, G. (2004). Passive capture and ensuing issues for a personal lifetime store, In *Proceedings of the First ACM Workshop on Continuous Archival and Retrieval of Personal Experiences* (pp. 48-55).

Healey, J. & Picard, R.W. (1998). StartleCam: a cybernetic wearable camera. In *The Second International Symposium on Wearable Computing* (pp. 42-49).

HiSPEC project (2002), *Privacy enhancing technologies: state of the art review, version 1.* HiSPEC Report. Retrieved on March 1, 2006, from http://www.hispec.org.uk/public_documents/7_1PETreview3.pdf

Hong, J., Ng, J, Lederer, S., & Landay, J. (2004). Privacy risk models for designing privacy-sensitive Ubiquitous Computing systems. In *Proceedings of the Conference on Designing Interactive Systems* (pp. 91-100).

Information Technology for European Advancement (2004), *ITEA Technology Roadmap for Software-Intensive Systems* (2nd ed.). Retrieved on March 1, 2006, from www.itea-office.org

Jansson, C. G., Jonsson, M., Kilander, F. et al. (2001). *Intrusion scenarios in meeting contexts* (FEEL Deliverable D5.1). Royal Technical University. Retrieved on March 1, 2006, from http://dsv.su.se/FEEL/zurich/Item_3-Intrusion_scenarios_in_meeting_contexts.pdf

Kato, U., Hayashi, T., Umeda, N. et al. (Ed.). (2004). *Flying Carpet: Towards the 4th Generation Mobile Communications Systems*, Version 2.00. 4th Generation Mobile Communications Committee. Retrieved on March 2, 2006, from http://www.mitf.org/public_e/archives/index.html

Kim, S. W., Kim, M. C., Park, S. H. et al. (2004). Gate reminder: a design case of a smart reminder. In D. Benyon, P. Moody et al. (Ed.) *Conference on Designing Interactive Systems* (pp. 81-90).

Koponen, K., Matkakorttien käytöstä syntyy valtava tietokanta matkustajista, *Helsingin Sanomat* (Finnish newspaper), 19.9.2002.

Lahlou, S. & Jegou, F. (2003). *European disappearing computer privacy design guidelines v1* (Ambient Agora Deliverable D15.4), Retrieved on March 2, 2006, from http://www.ambientagoras.org/downloads/D15%5B1%5D.4_-_Privacy_Design_Guidelines.pdf

Langheinrich, M. (2001). Privacy by design – principles of privacy-aware Ubiquitous Systems. In G. D. Abowd, B. Brumitt et al. (Ed.) *Proceedings of the Third International Conference on Ubiquitous Computing (UbiComp 2001)* (pp. 273-291). Springer-Verlag (Lecture Notes in Computer Science).

Langheinrich, M. (2006). *Personal privacy in Ubiquitous Computing*. Presentation in UK-Ubinet Summer School 2004. Retrieved on March 2, 2006, from http://www.vs.inf.ethz.ch/publ/slides/ukubinet2004-langhein.pdf

Ma, J., Yang, L. T., Apduhan, B. O. et al. (2005). Towards a smart world and ubiquitous intelligence: a walkthrough from smart things to smart hyperspaces and UbicKids. *International Journal of Pervasive Computing and Communications 1*(1), 53-68.

Masera, M., & Bloomfeld, R. (2003). *A Dependability Roadmap for the Information Society in Europe* (AMSD Deliverable D1.1). Retrieved on March 2, 2006, from https://rami.jrc.it/roadmaps/amsd

Mynatt, E., Essa, I., & Rogers, W. (2000). Increasing the opportunities for aging in place, In *Proceedings of the ACM Conference on Universal Usability* (pp. 65-71).

Nasoz, F., Alvarez, K., Lisetti, C., & Finkelstein, N. (2003). Emotion recognition from physiological signals for user modelling of affect. In *Proceedings of the Third Workshop on Affective and Attitude User Modelling*. Retrieved on March 3, 2006, from http://www.cs.ubc.ca/~conati/um03-affect/nasoz-final.pdf

Nissenbaum, H. (2004). Privacy as Contextual Integrity. *Washington Law Review, 79*(1), 101-139.

Orr, R. J., Raymond, R., Berman, J., & Seay, F. (1999). *A system for finding frequently lost objects in the home* (Tech. Rep. 99-24), Graphics, Visualization, and Usability Center, Georgia Tech.

Palmas, G., Tsapatsoulis, N., Apolloni, B. et al. (2001). *Generic Artefacts Specification and Acceptance Criteria*. (Oresteia Deliverable D01). Retrieved on March 2, 2006, from http://www.image.ntua.gr/oresteia/deliverables/ORESTEIA-IST-2000-26091-D01.pdf

Price, S., & Rogers, Y. (2004). Let's get physical: the learning benefits of interacting in digitally augmented physical spaces. *Computers and Education 43*(1-2), 137-151.

Schneier, B. (2005). Why Data Mining Won't Stop Terror. *Wired News*, March 9, 2005. Retrieved on April 19, 2007, from http://www.schneier.com/essay-108.html

Schneier, B. (2007). On Police Security Cameras: Wholesale Surveillance. *San Francisco Chronicle*, January 2007. Retrieved on April 19, 2007, from http://www.schneier.com/essay-147.html

Stajano, F. (2004). One user, many hats; and, sometimes, no hat – towards a secure yet usable PDA. In *Proceedings of the Security Protocols Workshop 2004* (pp. 51-64).

Truong, K. N., Huang, E. M., Stevens, M. M., & Abowd, G. D. (2004). How do users think about Ubiquitous Computing. In *Proceedings of CHI '04 extended abstracts on Human factors in computing systems* (pp. 1317 – 1320).

Web site of the American Civil Liberties Union and the ACLU Foundation (2006). *Eavesdropping 101: What Can The NSA Do?* 31.01.2006, Retrieved on March 2, 2006, from http://www.aclu.org/safefree/nsaspying/23989res20060131.html

Wright, D., & Gutwirth, S. (2008). *Safeguards in a world of ambient intelligence*. Springer