A FRAMEWORK FOR THE DESIGN OF PRIVACY PRESERVING PERVASIVE HEALTHCARE

Simon Moncrieff, Svetha Venkatesh, Geoff West

Department of Computing, Curtin University of Technology
GPO Box U1987, Perth, 6845, W. Australia

ABSTRACT

Privacy is an important aspect of pervasive and ubiquitous computing systems, and, in particular, pervasive healthcare. With reference to previous approaches on developing privacy sensitive pervasive healthcare applications, we detail a framework for the design of such systems that aims to minimise the impact of privacy on such systems. In reviewing previous approaches, we extract and combine common elements in order to unify the approaches and create a more formal methodology for designing privacy mechanisms in pervasive healthcare applications. In doing so we also consider the manner in which ubiquitous technologies impact on privacy and methods for reducing this impact. We demonstrate how the framework can be applied by using examples from the previous approaches. In addressing privacy issues, the framework aims to remove a large obstacle to deployment of pervasive healthcare systems, acceptance of the technology.

Index Terms— Privacy, Pervasive healthcare.

1. INTRODUCTION

Privacy is a vital component in the acceptance, and thus adoption of pervasive healthcare [1]. Failure to implement privacy measure in ubiquitous computing environments will lead to either to a failure to adopt, or the complete rejection, of such systems. Privacy is particularly salient to pervasive healthcare due to the private nature of the environments under observation, particularly in comparison with surveillance, and due to the sensitivity of information related to healthcare. Further, privacy sensitivity needs to be integrated into the system at the design stage as imposing privacy restrictions on an already developed system has the potential to reduce the functionality, or restrict the purpose of the system [2].

Examples of pervasive healthcare applications include assisted living [3, 4], and ubiquitous hospital communication systems [5]. To illustrate the need for privacy in pervasive health care, we consider the case of assisted living ubiquitous environments, which consist of a number of sensors, such as video cameras, and infrared and pressure sensors, that collect, process and interpret information from the environment in a manner that is transparent (ideally) to the occupant. The aim of assisted living environments is to monitor the occupant to ensure their safety, enabling the aged and invalid population to remain in their homes longer, increasing their quality of life, and reducing the financial burden of aged care [6]. However, the monitoring of home and living environments raises serious privacy implications due to the private nature of the home.A lack of privacy will result in a lack of trust, which will in turn impede the deployment of the technology.

To design privacy sensitive ubiquitous systems, it is first necessary to consider the properties of privacy, and how ubiquitous computing impacts upon these properties. The are numerous papers in the field of ubiquitous computing that discuss privacy, see [2, 7, 8] for examples. From examining such discussions, it becomes evident that no single definition of privacy is possible [7], due to its highly subjective nature, both with respect to people and context, encapsulated by concepts such as location, culture, and situation. This difficulty in defining privacy results in corresponding difficulty in designing and implementing privacy in ubiquitous computing.

Consequently, rather than attempting to define privacy with respect to ubiquitous computing, we examine the properties of ubiquitous computing that impact on privacy and subsequently identify important properties necessary for the design of privacy sensitive ubiquitous computing. We review previous approaches to implementing privacy in pervasive, with reference to these properties, in order to produce a framework that unifies these approaches. In doing so, we aim to present a design framework that can be applied to the implementation of privacy sensitive pervasive healthcare. We then demonstrate the properties of the framework with reference to the previous approaches. Our goal in this paper is to demonstrate core elements necessary for introducing privacy measures to pervasive healthcare applications.

2. INFORMATION FLOW AND PRIVACY

In ubiquitous computing, privacy can be thought of in terms of the flow of information, or, more specifically, the control of the flow of information. This view of privacy is perhaps best encapsulated by the *Privacy Regulation* theory proposed by social psychologist Irwin Altman [9]. Altman proposed that privacy is a dynamic and subjective process in which an

individual obtains an optimal level of privacy by controlling interactions with others. This optimal, or desired, level of privacy is dynamic, and changes with respect to both the individual and the context; that is an individual attempts to regulate privacy in accordance with their desire for social interaction (open), or social isolation (closed). The desired privacy is communicated using a set of social tools, including verbal and non-verbal conversation. If an individual fails, or is unable, to attain the desired privacy level, the result is either *social isolation* resulting from an excess of privacy, or *crowding*, which occurs when an individual receives more input than is desired due to a low level of privacy.

Thus, Altman viewed privacy as the control of the communication of information between two or more parties. Consequently, privacy regulation theory can be extended to ubiquitous computing environments by extending this communication to include sensors within the environment [8]. However, due to the unique properties of ubiquitous computing environments, i.e. the collection of information from the environment in a manner transparent to, and not requiring interaction from, the occupant. The occupant no longer controls all the information that is being communicated about themselves. That is, the occupant has no control of the information they reveal to the environment. Traditional Privacy Regulation theory considers the negative impact of too much information input, crowding. However, the theory does not consider the case of excessive information *output*. This was addressed by Lehikoinen [8], who introduced the concept of *Leaking*, which is the state of privacy that that occurs due to the *information leak* caused by the uncontrolled, or unintentional, flow of information from an individual within the environment to the sensors, i.e. the inadvertent disclosure of information. This impacts on the privacy of an occupant of the environment due to the sensors storing and potentially *sharing* information with other users, comprising either observers, in the case of assisted living and surveillance for example, or other users within the environment, such as in the case of media spaces, in which information is shared.

Employing the concepts of *information flow* and *information leak* in conjunction with Privacy Regulation theory enables us to examine the impact of ubiquitous computing environments on privacy. The asymmetric flow of information [10] causes an intrusion into the users experienced privacy as they are no longer able to control their own information flow. Therefore, to maximise the privacy, the unknown, or inadvertent information flow from the user to the sensor should be minimised. Minimal information leak corresponds to collecting no data from the environment (maximum privacy). However, this would invalidate the functionality of the ubiquitous environment, i.e. the purpose of the observer. Consequently, we can view privacy in ubiquitous computing as an optimisation problem, balancing the privacy of the user, with the functionality of the system.

Control over what information is captured, and *feedback* on this information have previously been identified as important to the implementation of privacy in ubiquitous computing, particularly within private environments [7]. These concepts can be included in the Privacy Regulation theory approach to privacy in ubiquitous computing. Control enables a user to control the information flow, reducing information leak and enabling the user to attain a closer approximation to the desired level of privacy. Feedback reduces inadvertent information flow by providing the user with details of the information that is communicated to the environment, which in turn decreases the asymmetric information flow.

For complex ubiquitous environments, the optimisation between the privacy and functionality will require a dynamic approach due to the multiple, and changing situations that are present in complicated, real world situations. That is, a single privacy policy would not suffice as the situation, or context within the environment, does not influence what an observer is able to view. This approach is inflexible, and particularly unsuited to to the real-time, active monitoring of complex environments due to the different contexts that can occur. Consequently, such environments require a more dynamic approach to privacy to achieve the trade off between minimising the intrusion into the privacy of those monitored, and retaining the purpose of the system. A single privacy policy would be either too invasive for the occupant, or too restrictive for those monitoring the environment. For example, in an assisted living smart home, a carer observer may not be able to verify an occupant status in the case of an injury (e.g. a fall), if they are not given access to sufficient information. However, revealing too much information to an observer is likely to cause embarrassment, and a subsequent rejection, or resentment, of the monitoring. Consider the example of the bathroom, which is associated both with sensitive and private activities, but also potential hazards that require monitoring.

3. APPROACHES TO PRIVACY IN PERVASIVE HEALTHCARE

While privacy is important to the long term success of pervasive healthcare, addressing privacy remains relatively unexplored. There are potentially three perspectives on privacy;
1. Approaches that recognise the need for privacy, but do not deal with it explicitly, e.g. Vemuri and Bender [11] (a personal memory for recording conversations), and de Silva *et al.* [12] (an assisted living environment), both recognised the need for privacy measures in order for the proposed technologies to be accepted, but did not address privacy.
2. Approaches that avoid the issue of privacy by limiting the information collected, e.g. [13] limit the information gathered by using simpler sensors, such as binary infrared motion sensors. The main drawback of this approach results from limiting information, which in turn reduces the functionality of the system. For example, verifying whether an alert is a false positive in an assisted living smart home environment is

costly without the presence of cameras.
3. Approaches that actively address privacy. These works form the major part of our review and are discussed in the next section.

4. PRIVACY SENSITIVE COMPUTING

4.1. Framework

To put the previous approaches in perspective, it is necessary to construct a common framework in which the approaches can be analysed. The design of such a privacy sensitive system has to account for both the observer, the observed, and the relationship between the observer and the observed. A design framework to achieve this is proposed in Figure 1 and is built with reference to the information process flow discussed in Section 2.[1] In broad terms, the context or user preferences accounts for the user of the environment (observed), while the rule set represents the influence of the observer on the privacy filter. The context incorporates the observed into the privacy system by enabling the privacy filter to be changed according to the situation within the environment. The feedback and control incorporate the observed/observer relationship, and reduce information leak. It should be noted that there are a number of design factors that can be considered when implementing privacy in ubiquitous computing, such as the balancing of risk with benefit [7], and trust. These factors are not mutually exclusive from design framework, but are factors that can be considered in implementing instances of framework.

[1] An early version of this framework was introduced in [3]

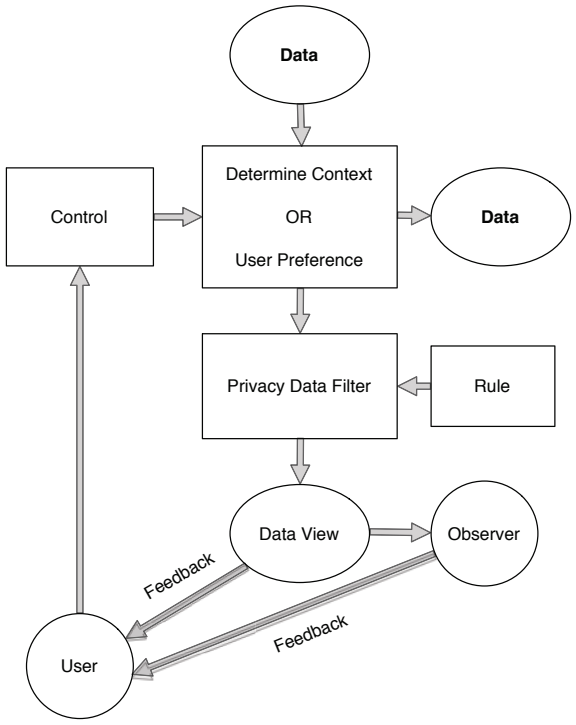

Fig. 1. Design framework for privacy sensitive ubiquitous computing applications in pervasive healthcare.

4.2. Active Approaches to Privacy

A number of methods have been proposed that actively address privacy in a number of pervasive healthcare applications. These applications include; assisted living, both within a nursing home [4] and a smart home environment [3], a ubiquitous mobile hospital information system [5], hospital security [14], the privacy preserving sharing of videos containing medical information [15], and access to medical databases for surveillance [16]. Privacy policies are generally implemented using data hiding techniques on the sensor data in order to obscure, or remove, privacy sensitive information. The usefulness of the proposed framework presented in Section 4.1, is now demonstrated through a survey of the above methods.
The input to the framework is the data collected from the environment, corresponding to both contextual data (e.g. time, date), and sensor data. Examples of sensor input include video [3, 14, 4], binary sensors and audio [3].
The next module interprets the environmental contextual information present in the sensor data, *or* determines the user preference that should apply, in certain cases the meta-data will influence the preference. For example, Moncrieff *et al.* [3] determined the environmental context using indicators of the activity present within the environment. Wickramasuriya *et al* [14] and Chen *et al.* [4] used identity of the context, either determining whether or not an individual was authorised to enter the environment [14], or the identity of residents in a nursing home [4]. Examples of using preferences include [15], in which the level of data hiding applied was based on the video owner's assessment of the risk versus benefit of sharing the video, and Tentori *et al.* [5], who adjusted user preferences according to contextual information such as location.
The data filter then determines the correct privacy policy to apply given the preference, or environmental context, and the applied rule. The rule represents the influence of the observer on the privacy, while, conversely, the preference and context adjusts the privacy with respect to the user (observed). The sensor data is then filtered according to the privacy policy. This is achieved by applying the data hiding technique corresponding to the privacy policy to the data, resulting in the *Data View*, which represents the input data, transformed by the privacy filter. Fan *et al* [15] applied the data hiding technique corresponding the video owner's set preference level. A number of data hiding techniques techniques were used to obscure data at different levels, for example blurring the video, or replacing people within the video with virtual objects, or avatars. Moncrieff *et al* [3] used multiple data hiding methods representing different levels of privacy, determining the appropriate data hiding method according to the context and the

rule input, which was encapsulated by the role of the observer. For example, for a carer observer, privacy measures were reduced if an abnormal activity was detected, enabling the carer to determine if an alarm needed to be raised. In [4], images of residents of the nursing home were replaced with an edge motion history image, revealing details of the residents activities, but not identifying information. Sweeney [16] used the detection of unusual activity in a medical database as a mechanism to lower the level of anonymity in the visualisation applied to the database, i.e. data was presented at a higher resolution if an unusual set of data occurred. The rule imposed can vary in complexity. For example, Moncrieff *et al* [3] used multiple levels of data hiding, and thus privacy, for each sensor present (audio, video and binary sensors), using a decision tree was to generate the rules mapping the context to the appropriate data hiding level. In [14], the rule is encapsulated in authorisation level required to enter the environment.

The data view is then presented to the observer, and to the user (observed), representing feedback, who can then *control* the privacy level by adjusting either the context or the applied preference accordingly. Although not always implemented, there is a provision within the framework providing feedback on the observer to the observed. While not viable in surveillance applications, in applications such as assisted living smart homes, giving the occupant access to information concerning both who is observing and what they are observing reduces asymmetric information flow, and will in turn increase trust in the system [3]. The feedback and control can be implicit within the environment. For example, in the method proposed by Wickramasuriya *et al* [14], feedback and control is provided to people with authorised access *a priori*, as they are informed of the surveillance, and that authorised personal are removed from the video and given the opportunity to select the data hiding method applied to them. Feedback to unauthorised people within the environment can be given in the form of warning signs. In private environments, such as an assisted living smart home, a more robust approach to feedback is possible, and indeed necessary given the private nature of the home environment. Moncrieff *et al.* [3] proposed a number of feedback mechanisms suitable for a smart home. Detailed feedback was given in the form of logs indicating what information each observer accessed, and when; while instantaneous low level feedback was given to enable the occupant to adjust the context, i.e. the occupant was given the ability to influence the privacy by controlling the context. For example, if a false abnormal event is detected, the occupant can indicate that the actual context within the house should register as *normal*, this interaction in turn verifies the status of the occupant. Consequently, the privacy can then be adjusted to the appropriate level for a normal context. This limited control was necessary due to the reduction in functionality of the system that would result if the occupant was given control to adjust preferences on the fly (i.e. the option of turning the monitoring off).

5. CONCLUSION

In this paper we have presented a design framework for implementing privacy measures in ubiquitous computing environments, and demonstrated its application to pervasive healthcare. Given the sensitivity of healthcare environments, and the associated data, addressing privacy issues will play a large part in the adoption of pervasive healthcare applications.

6. REFERENCES

[1] P. A. Nixon, W. Wagealla, C. English, and S. Terzis, *Smart Environments: Technology, Protocols, and Applications*, chapter Security, Privacy and Trust Issues in Smart Environments, pp. 249–270, Wiley, 2004.

[2] J. I. Hong and J. A. Landay, "An architecture for privacy-sensitive ubiquitous computing," in *2nd international Conference on Mobile Systems, Applications, and Services, MobiSys '04*, Boston, MA, USA, June 2004, pp. 177–189, ACM Press, New York, NY.

[3] Simon Moncrieff, Svetha Venkatesh, and Geoff West, "Dynamic privacy assessment in a smart house environment using multimodal sensing," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 5, no. 2, pp. 1–29, 2008.

[4] Datong Chen, Yi Chang, Rong Yan, and Jie Yang, "Tools for protecting the privacy of specific individuals in video," *EURASIP Journal on Applied Signal Processing*, vol. 2007, pp. 107–115, 2007.

[5] Monica Tentori, Jesus Favela, Marcela D. Rodriguez, and Victor M. Gonzalez, "Supporting quality of privacy (qop) in pervasive computing," in *ENC '05*, Washington, DC, USA, 2005, pp. 58–67, IEEE Computer Society.

[6] S. Helal, B. Winkler, Choonhwa Lee, L. Kaddoura, Y. Ran, C. Giraldo, S. Kuchibhotla, and W. Mann, "Enabling location-aware pervasive computing applications for the elderly," in *PerCom 2003*, March 2003, pp. 531– 536.

[7] Jason I. Hong, Jennifer D. Ng, Scott Lederer, and James A. Landay, "Privacy risk models for designing privacy-sensitive ubiquitous computing systems," in *Proceedings of the 5th conference on Designing interactive systems*, New York, NY, USA, 2004, pp. 91–100, ACM.

[8] J. Lehikoinen, J. Lehikoinen, and P. Huuskonen, "Understanding privacy regulation in ubicomp interactions," *Personal and Ubiquitous Computing*, 2007.

[9] I. Altman, *The Environment and Social Behavior: Privacy, Personal Space, Territory and Crowding*, Brooks/Cole Pub. Co., Inc., Monterey, CA, 1975.

[10] Xiaodong Jiang, Jason I. Hong, and James A. L, "Approximate information flows: Socially-based modeling of privacy in ubiquitous computing," in *In Proceedings of Ubicomp 2002*, 2002, pp. 176–193.

[11] S. Vemuri and W. Bender, "Next-generation personal memory aids," *BT Technology Journal*, vol. 22, no. 4, pp. 125–138, 2004.

[12] Gamhewage C. de Silva, Byoungjun Oh, Toshihiko Yamasaki, and Kiyoharu Aizawa, "Experience retrieval in a ubiquitous home," in *CARPE '05*, New York, NY, USA, 2005, pp. 35–44, ACM.

[13] D. H. Wilson, *Assistive Intelligent Environments for Automatic Health Monitoring*, Ph.D. thesis, Robotics Institute, Carnegie Mellon University, September 2005.

[14] J. Wickramasuriya, M. Alhazzazi, M. Datt, S. Mehrotra, and N. Venkatasubramanian, "Privacy-protecting data collection in media spaces," in *ACM Multimedia*, New York, NY, October 2004.

[15] Jianping Fan, Hangzai Luo, Mohand-Said Hacid, and Elisa Bertino, "A novel approach for privacy-preserving video sharing," in *ACM CIKM '05*, New York, NY, USA, 2005, pp. 609–616, ACM.

[16] L. Sweeney, "Privacy-preserving surveillance using selective revelation," *IEEE Intelligent Systems*, vol. Sept-Oct, 2005.