Personalisation in Intelligent Environments: Managing the Information Flow

Craig Chatfield[1], David Carmichael[2], René Hexel[1], Judy Kay[2], and Bob Kummerfeld[2]

[1] School of Information and Communication Technology
Griffith University
Nathan Campus, Australia
{r.hexel,c.chatfield}@griffith.edu.au
[2] Smart Internet Technology Research Group
School of Information Technologies,
University of Sydney, Australia
{judy,kummerfeld,dcarmich}@it.usyd.edu.au

Abstract. This paper describes research into the personalised delivery of information about an intelligent environment. The challenges we need to address are the dual forms of the Invisibility Problem. Both follow from the fact that the sensors and services in an intelligent environment are intended to be invisible, blending into the environment naturally.
We describe the elements of the Invisibility Problem in relation to our previous work on one such intelligent environment, MyPlace, and present an initial user study (N=17) which will inform the design of the personalisation that should be done to enable users to effectively negotiate and control an intelligent environment.
This user study finds that trust is a critical aspect of intelligent environment design, and effective user feedback and the management of information flow across physical, social and temporal borders is essential to maintain user trust in an intelligent environment. This study identified that users could readily understand the invisibility problem, and are willing to share personal information to recieve a tangible benefit in the form of the personalisation of an intelligent environment's services.

Keywords: Personalisation, User Modelling, Privacy

1 Introduction

Intelligent Environments (IEs) are electronically augmented environments that provide services based upon user preferences and current environmental conditions. One challenge for IE developers is balance the goal of invisibility [1] with the need for accessibility.

We call this the Invisibility Problem. It has two main dimensions which relate to two facets of invisibility. First, the IE is supposed to include calm technology

that just blends into the environment. So, for example, a person entering a home may not even notice the mirror near the door because that mirror fits naturally into the environment. However, in the IE, that mirror may also be able to deliver personalised information to known users. This may be a very useful service.

However, it brings a collection of challenges. How does a new user learn that this mirror is more than a mirror? If the mirror is only intended for members of the household, does a visitor need to know that it is any more than a mirror? Suppose the person entering this IE home is a relative who is coming to stay for a few months. How can we provide information about this person to the IE so that they are treated like a member of the family and given access to the IE facilities.

There has been some exploration of ways to address aspects of the first form of invisibility. Heer et al. [2] examine invisibility from a psychological perspective, making the point that an invisible interface does not in any way imply physical invisibility. The Digiscope [3] explores the user of augmented reality: it provides a large semitransparent window to view the world. Visual tag recognition and RFID are used to detect items in a suitcase, and the display shows information about items within the suitcase. We have explored an approach based on a mobile phone with a camera, based upon fiducial markers [4]. These systems require the user to point a device at the right part of the IE. Even then they are presented the standard information stored in the database. If the user is unaware of a service in the environment, they would not be able to use this class of interface.

A quite different dimension of invisibility relates to the automatic and natural use of facilities in the IE as envisaged by Weiser [1]. The challenge in this case is that an artefact which is natural to use for the experienced user requires learning for the new user. For example, an experienced user of a microwave oven is barely aware of the steps in making normal use of it: yet, the first time user takes some time to work out just what to do. The same will apply for communication with an IE. In the period from first meeting a new part of an IE, the user may well appreciate personalised information about the device, with tuition and assistance in mastering it. This is a candidate for personalised tuition [5].

MyPlace is an intelligent environment [6]. We have been exploring technical dimensions of the Invisibility Problem in the context of that work. In particular, we have been exploring the modelling of people, places, devices and sensors in a consistent manner [6]. One part of that work provides a user interface to MyPlace, as personalised information about an environment. Having built prototype interfaces for such a system, we are now convinced that it is important to gain greater understanding of how people relate to the many issues involved in providing personalised information delivery about an IE like MyPlace.

Another dimension of our work has explored the problems of building user models for use in IEs. We have defined an architecture for the secure management of identity in such environments [7] and we have implemented a particular form of that environment [8] where a service provider constructs a user model *persona*. For example, a university creates a persona for each student and, as a service, provides that to the student, with authentication. Of course, the univer-

sity provides student information constituting just a part of a user model, which is why we call it a persona [9]. We have also explored a range of the technical issues of managing such personas, including and P3P-inspired approached to defining policies for releasing a persona into an IE [10]. In that work, we also created and evaluated a prototype system that resides on a user's PDA and enables the user to load a persona from a service provider, integrate it into their own user model and release it into an IE. This work identified some real challenges in constructing such user interfaces, especially in the context of current handheld devices. But more important than that, it also highlighted the importance and timeliness of user studies which can inform future research in the management of user models in IEs and the use of those models for personalised information and service. These need to explore users' concern about privacy and ownership of personal data [11].

A initial user study (N=17) was carried out to inform on the MyPlace interface, to examine the invisibilility problem and users perceptions of privacy and information sharing within IEs. User privacy within IEs is diminished when users' information crosses a personal border of privacy (e.g. physical, social or temporal borders) [12]. This study investigates this, examining the effect of information flow across physical borders, within certain social groups and outside of temporal borders. We also consider the influences of trust and reputation [13], [14] and user anonymity on user's privacy [15]. Results from this study will inform on managing user privacy and the personalisation of services within an IE.

In Section 2, we introduce our user interface that attacks the Invisibility Problem in the MyPlace system. We briefly outline the type of customisation it affords. Section 3 provides an outlines of the architecture and Section 4 is the core of this paper, describing the initial user study and its implications for the future design of MyPlace and other intelligent environments.

2 Personalised Views

MyPlace [6] provides personalised intelligent environments. This means that the environment gathers information about users to build user models. It uses these to customise services within the environment. One of those services is MVI the MyPlace Visibility Interface. We will describe the simple personalisation available in a scenario involving two users arriving at a building in a university.

Kate and Sam are postgraduate students who have just joined the SmartLab research group. As each enters the Smart Meeting Room, they see the screens shown in Figure 1.

Kate is really concerned about issues of privacy and always wants to know about any sensors in her environment. Sam is unconcerned about these issues and MVI presents information on the assumption that he does not want to be bothered with the details of the many sensors in this building. Accordingly, MVI presents the MyPlace environment to Kate as shown in Figure 1a and to Sam as in Figure 1b.

Both Kate and Sam are new, attending their first seminar. Their user models both indicate unfamiliarity with the building and, hence, of this Smart Meeting Room. So the MVI interface presents a limited number of the services available and information about nearby locations. This is partly to match the screen size of the PDA that each is using to see what is in the IE and partly to customise the information to be helpful: it should not overwhelm the users and should ensure that they are likely to be able to learn about the most relevant service in the IE.

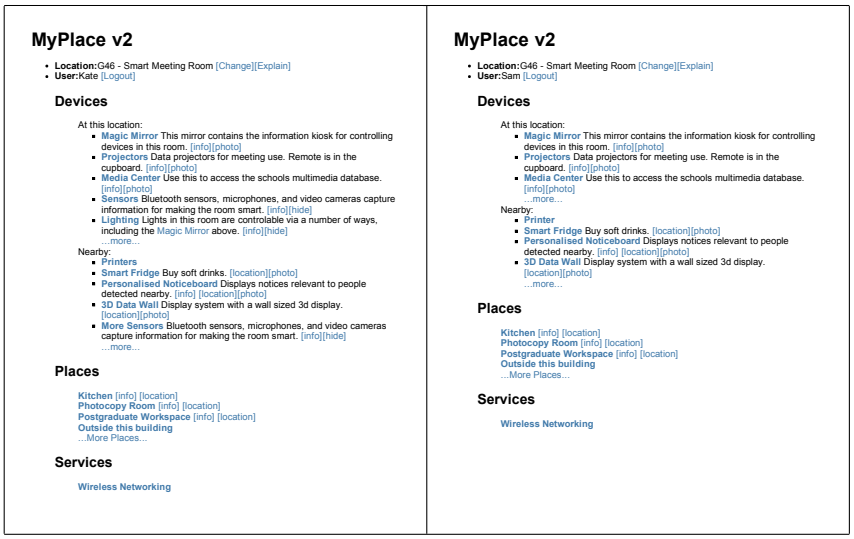

Fig. 1. Example views for Kate (a:left) and Sam (b:right)

As important as hiding information from the user, is allowing the user to scrutinise what has been displayed or hidden from them. This facility is provided when the user clicks the '...more...' link at the bottom of a list of items. Figure 2 shows how Sam can see more of the IE elements, including the sensors, lighting and air-conditioning control. The longer list of devices in the current location is shown with those normally hidden in a light grey.

The display also provides links to additional information about the devices. This too can be personalised. For example, the Media Center has many facilities and most first-time users have difficulties getting the basic parts working. If Kate or Sam click the link for this information, they will be provided details as defined by their user models.

We might suppose that Kate, although new to this research group, has used the same type of Media Center in another room managed by MyPlace. We might further suppose that Kate's user model, built up by MyPlace as she used that other centre, reflects her knowledge of all the basic facilities of that type of Media Center. In that case, it should be possible for MVI to present details of the more sophisticated facilities she has not used and is modelled as not knowing.

Of course, this scenario introduces many user concerns. The underlying assumption is that the IE will hold some form of user model. It is important to learn how people can appreciate the issues involved and how they respond to them. We need the type of exploration conducted to inform work on P3P [16]. This work influenced the design. They learnt that most users do not want completely automated exchange of private data: they want to okay any transfer of private data. Of course, these results relate to the web and very limited interaction that is largely conscious. We need to study the way people perceive such issues in the case of pervasive intelligent environments.

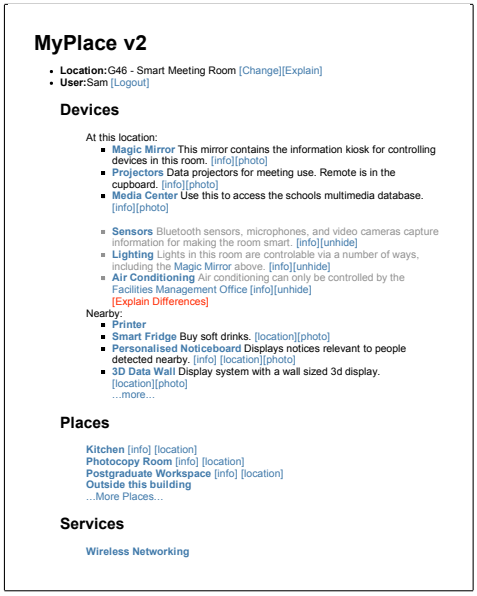

Fig. 2. Sam's world view when scrutinising devices available at the current location. The devices normally hidden are shown in light grey.

In the above example, the scenario was strictly tied to the individual students, Kate and Sam. In a real world scenario, people have different needs and requirements of the system, depending on their current activities. Often, these activities can be so fundamentally different that it makes sense for the system to combine several different profiles within one user model.

This ability to control what information is presented to the IE is essential for users to control their privacy. The selection of different personas is comparable to the multiple *faces* that Lederer, Dey and Mankoff [17] suggest are necessary for users to maintain privacy within ubiquitous computing environments. This ability to control what information is shared is essential in any system designed to respect and support user privacy. The influence of this MyPlace feature on privacy will be investigated in the evaluation described in section 4.

3 Technical System Description

In this section we describe the technical aspects of the MyPlace system. Of note are its delivery architecture, uniform space, device and user modelling, modelling of changing values and interface.

The MVI elements of MyPlace are shown in the partial system architecture in Figure 3. The broader system, with its collection of sensor data and delivery of data across the IE is described elsewhere [6]. The user's personal device connects to the MyPlace server to access the service. Their personal device also holds their personal user model. They can also choose to release a portion of their personal user model, known as a persona, to the MyPlace server. The information in this persona is used to complement the model of the user which the MyPlace server has.

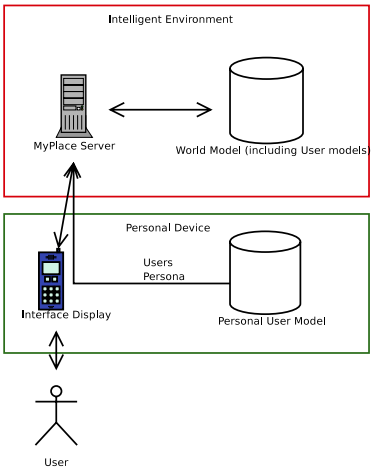

Fig. 3. The MVI elements of the MyPlace architecture.

People, devices, places and activities are modelled in a uniform manner within the MyPlace system (Figure 4). Entities are organised by *Contexts*, which are in a tree structure, with additional pointers between contexts being used to represent semantics. Each context contains components which hold evidence about a particular aspect of an entity. The evidence in components can be resolved to a value, which can either be a simple type such as a string or number, or a pointer to another context.

Changing and uncertain values are modelled using a method taken from the Personis user modelling system [9]. As sensor data about an item or user arrives in the system it *accretes*, simply being stored without interpretation. When a value is required the stored evidence is examined and *resolved* to a value.

To give a concrete example, we model whether users know about a particular room. To do this we collect a range of evidence, such as that provided by a

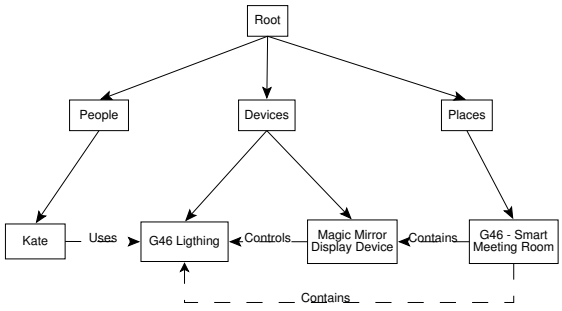

Fig. 4. A small subsection of the world model. The solid arrows show the tree structure, the dashed, labelled lines are example semantic pointers.

sensor which detects that user in the room. Such information is stored without interpretation, and only when we need an answer to 'does the user know this room?' do we resolve to a value. This saves unnecessary computation for cases when evidence is stored, but a value is never needed. It also allows for different interpretations of the evidence by simply changing the resolution algorithm at the point the values are needed.

We have built prototype IEs, with sensors providing evidence for user modelling. We have used this in applications, such as one that locates the user. There are many on-going challenging technical issues still to explore. However, at this point, it is timely to gain a better understanding of how people respond to the many human issues that need to be considered in conjunction with the design of technical solutions.

4 Evaluation

This study examined user interaction within a simulated IE covering a university common area using the MVI. The user study begun with discussion questions to orient users on the nature of the invisibility problem and intelligent environments. Participants then used the MyPlace interface to examine the personalisation process and answered questions on their privacy and information sharing within this IE. Participants were recruited by requesting volunteers around the university campus where the study took place. Seventeen users took part in this study, comprising of eight males and nine females. Most participants were in their 20's (12 users), reported their field as 'IT' (12 users) and owned a mobile phone or PDA (14 users).

To identify the effectiveness of using this interface to explain the difficulties of displaying relevant information within an intelligent environment, users we asked to describe the invisibility problem and the process of personalisation. While four users professed not to understand the invisibility problem, the others described it as a form of personalisation and identification of relevant intelligent environment services. All users seemed happy with the descriptions of these

processes given during the pre-study discussion questions, and the authors feel that users understand the need for this type of invisibility and personalisation. User feedback on the MVI included concerns about the navigation and presentation of such a large amounts of information on such as small screen, including the need for a more graphical interface to better organise information, and the desire to have more effective personalisation of the displayed services. 88.2% of the participants indicated that they considered the services to be useful, while the remaining two users described other location based services they would prefer to see in this intelligent environment. Other users comments indicated that the delivery of these services could be better customised to their needs or that only anonymous services would be useful.

Users were asked to rate how concerned they were with the collection of generic user information (e.g. whether a student or staff member, interests, etc) and more identifiable information (e.g. name or ID number) on a scale of 1-5. As expected, users showed a low level of concern (2.00) in sharing generic information with the environment, but were much more concerned (3.41) with sharing personal information. Twelve users (70.6%) indicated that they would be more likely to share more personal information if the environment was unable to link this information to their real identity. All other respondents indicated that it would not effect their information sharing preferences, with user comments describing their wish to see 'tangible benefits' before sharing personal information.

Users were then asked whether the sharing of information outside of this environment would make them more or less likely to provide personal information to personalise their interaction with the IE. Ten users (58.8%) noted they would be less likely, with half of these users citing the potential for abuse as their reason. Five users indicated it would have no effect, while two respondents declined to answer, commenting that it would completely depend on what information was shared with what IE. To gauge any changing attitudes on the completion of this survey, a final question asked if any shared personal information was restricted for use within this environment, would the user be more or less likely to share personal information. A similar result was found; with nine users (52.9%) indicating this would make them more likely to share information.

Users were then asked to consider the use of a centralised location service that provide users name and location details to other users, in accordance with their personal preferences. User responses on the likelihood of using this service showed reluctance, the average was 2.76, with the only user responding with a '5' commenting that the university setting gave them some trust in the IE. When asked whether the restriction of access to the users current location (i.e. historical searchers we not allowed) would make the more or less likely to use this service, eleven users (64.7%) responded with more likely and six users (35.3%) remained unchanged. Thirteen users (76.5%) identified that control over who could have access to their location information would make them more likely to share information; a sign that control of the usage of personal information reduces privacy concerns.

Users were asked to describe how they would normally share their location information. Most respondents indicated that their friends (15 user, 88.2%) or co-workers (11 user, 64.7%) were the only people they would share this information with, and several user comments indicated that more fine grain control was needed (e.g. only sharing location data with colleagues during work hours). Only two users indicated they would share information with all users or with commercial services, and only one indicated they would allow the sharing of their location history. This result was reflected when asked to consider the sharing of all information. 82.4% of all participants agreed that the ability to control who has access to their information would make them more likely to provide information for the personalisation of IE services. 76.5% of users also indicated that the owner of an IE would effect how much information they would share, and many users commented that this was related to their trust of the IE's owners.

5 Conclusions and Future Work

Personalised information access has had very little attention in ubiquitous computing research. This is unfortunate because there are many aspects where the IE will need to communicate with users. In many of these, there will be a real need to personalise that communication to match the user's knowledge, especially their familiarity with the IE and the services within it. Equally, IEs offer a potentially important context for personalisation research because IEs can collect evidence about users in the IE and use this to build user models.

The results of this user study show users are concerned about their privacy, and want to control information flowing. Users want control over who has access of their information, and want explicit feedback on its use and distribution. When users evaluate an information exchange, they are evaluating their trust of the system, and the expected tangible benefits they receive for providing information. Intelligent environments should seek to make feedback on the information as complete as possible to allow users to make their own determination of the risk to their privacy versus the perceived reward for sharing information. This study supports the limitation of information flow across physical, temporal and social borders [12] to maintain user privacy, and that methods of managing IE trust and reputation should be investigated. Users demonstrated they undestood the invisibility problem, and were willing to provide personal information in exchange for effective personalisation. Further research will expand this work to a wider audience and attempt to develop a more detailed understanding of user privacy within pervasive computing environments.

References

1. Weiser, M.: The computer for the 21st century. Scientific American **265** (1991) 94–104
2. Heer, J., Khooshabeh, P.: Seeing the invisible. In: Workshop on Invisible and Transparent Interfaces, Advanced Visual Interfaces, Gallipoli, Italy (2004)

3. Ferscha, A., Keller, M.: DigiScope: an invisible worlds window. In: Adjunct Proceedings, The Fifth International Conference on Ubiquitous Computing, Seattle, WA, USA (2003) 261–264
4. Assad, M., Carmichael, D.J., Cutting, D., Hudson, A.: AR phone: Accessible Augmented Reality in the Intelligent Environment. In Viller, S., Wyeth, P., eds.: 2003 Australiasian Computer Human Interaction Conference OZCHI 2003, University of Queensland, Australia (2003) 232–235
5. Brusilovsky, P.: Adaptive hypermedia. User modeling and User-Adapted Interaction **11** (2001) 87–110
6. Kay, J., Kummerfeld, B., Carmichael, D.J.: Consistent modelling of users, devices and environments in a ubiquitous computing environment. User modeling and User-Adapted Interaction: the Journal of Personalization Research (to appear)
7. Hitchens, M., Kay, J., Kummerfeld, B.: Secure identity management for pseudo-anonymous service acces. Technical Report 546, School of Information Technologies, University of Sydney (2004)
8. Hitchens, M., Kay, J., Kummerfeld, B., Brar, A.: Secure identity management for pseudo-anonymous service acces. In: Proceedings of SPC2005, International Conference on Security in Pervasive Computing, Boppard, Germany (2005)
9. Kay, J., Kummerfeld, B., Lauder, P.: Personis: a server for user models. In: Proceedings of Adaptive Hypertext 2002, Springer-Verlag (2002) 203–212
10. Brar, A., Kay, J.: Privacy and security in ubiquitous personalized applications. Technical Report 561, School of Information Technologies, University of Sydney (2005)
11. Lahlou, S., Langheinrich, M., Röcker, C.: Privacy and trust issues with invisible computers. Communications of the ACM **48** (2005) 59–60
12. Langheinrich, M.: Privacy Invasions in Ubiquitous Computing. Privacy in Ubiquitous Computing Workshop, In Proceedings of the 4th International Conference on Ubiquitous Computing Gteborg, Sweden (2002)
13. Chatfield, C., Häkkilä, J.: Designing intelligent environments - user perceptions on information sharing. In: Proceedings of 6th Asia-Pacific Conference on Computer-Human Interaction. (2004)
14. Goecks, J., Mynatt, E.: Enabling privacy management in ubiquitous computing environments through trust and reputation systems. In: Computer-Supported Cooperative Work 2002. (2002)
15. Langheinrich, M.: Privacy by design — principles of privacy-aware ubiquitous systems. In: Proceedings of the 3rd International Conference on Ubiquitous Computing, Atlanta, Georgia, USA (2001) 273–291
16. Ackerman, M.S., Cranor, L.: Privacy critics: UI components to safeguard users' privacy and security in ubiquitous personalized applications. In: The CHI is the Limit: Human Factors in Computing Systems, CHI99 Conference Proceedings, Pittsburgh, Pennsylvania, ACM Press (1999) 1–8
17. Lederer, S., Dey, A.K., Mankoff, J.: Conceptual model and a metaphor of everyday privacy in ubiquitous computing environments. Technical Report, UCB-CSD-02-1188 (2002)

Acknowledgements

The authors would like to acknowledge the support of the Smart Internet Technology CRC.