MyWorkPlace: Personalised information about a Ubiquitous Computing enabled building

David Carmichael, Judy Kay, Bob Kummerfeld and William Niu
University of Sydney
School of Information Technology
NSW 2006 Australia
{dcarmich, judy, bob, niu}@it.usyd.edu.au

ABSTRACT

This paper describes MyWorkPlace, which uses personalisation of automatically generated ontologies to provide users with personalised information about new and invisible items within a ubiquitous computing environment.

A key component of this system is the automatic generation of the ontologies, and models used to drive it. This data is being gathered from a number of sources including: building maps, the build manual, staff directory, student timetables, the departmental calendar and room bookings.

We describe planned evaluation of the system in a deployment to a new building and its new inhabitants.

Keywords

User modelling, Invisibility Problem, Automatically Generated Ontologies

1. INTRODUCTION

Ubiquitous computing aims to embed our everyday environment with devices, sensors and services in such a way that they are as unobtrusive as possible, to the point of becoming "invisible to common awareness" [1]. When achieved this invisibility creates its own problems. Users may be unable to discover what services are available to them, what sensors are detecting them or why the system has reacted in a particular way. We call this the *Invisibility Problem* [2]

To motivate this paper we consider a real life example of the invisibility problem, a University Department moving to a new building instrumented with a number of ubiquitous computing features. We examine the interactions and information needs of Fred, an academic.

Initially all facilities (ubiquitous computing and otherwise) within the building are unknown to Fred. He may have an idea of some types of facilities which are available, but is unlikely to know the full extent of the facilities available and details about them. A list of all the information about the building, its sensors, services, devices and occupants would have hundreds or thousands of items and clearly overwhelm the user.

MyWorkPlace solves this problem by modelling users, places, devices, sensors, services and objects to provide personalised views of the items within a user's environment. An important part of the system is the use of an automatically generated ontology to assist with the selection of items to display for the user.

The ability for ontologies to facilitate human-machine and machines-machine communications has gained wide recognition in the development of the UbiComp. It has been used in middleware to facilitate context management and reasoning [3, 4, 5], and user modelling [6].

A novel aspect of our work is the use of ontological data generated using different sources which have different levels of reliability to personalise the information given to a user based on their context.

The effort in creating a comprehensive ontology is substantial. Partial or completely automated generation has the possibility to greatly reduce this effort. Depending on the degree and type of automation, the reliability of the ontology can vary greatly. To cope with this, we are examining multiple levels of ontologies.

Our automatically generated ontology in being built from a number of sources including: building maps, the build manual, staff directory, student timetables, the departmental calendars and room bookings.

The rest of this paper is organised as follows: We first describe some related work in Section 2 to set the scene. Section 3 describes MyWorkPlace when used in the scenario described of Fred. The methods we use for automatically generating the ontologies are described in section 4. We conclude with a discussion of our proposed evaluation and future work in section 5.

2. RELATED WORK

Weiser predicted that ubiquitous computing would become a technology which disappeared and became invisible [1]. Others such as Heer and Khooshabeh have examined the nature of this invisibility [7]. They note that an invisible interface does not imply literal physical invisibility. Edwards notes some of the problems associated with Invisibility while examining the challenges of putting ubiquitous computing into the home [8].

There has been some work which addresses the issue of informing users of ubiquitous computing systems what devices and services are available to them. The AFAIK system is a multimodal help system for an intelligent room [9]. Help

ubi-PCMM at UbiComp'06, September 2006, Orange County, CA, USA.
.

content must be entered in XML, but the help system is not personalised to the user or their context. The Digiscope [10] is a system for viewing attributes of objects within an intelligent environment. It consists a large semitransparent display mounted on an movable arm. Information is retrieved from a database about objects which are identified using RFID and visual tagging. The NearMe system [11] provides users with a list of nearby devices by examining the signatures of nearby wifi access points, and making a request to a server of all known nearby devices. This is different from our work as it does not seek to deliver the same level of detail and is not personalised.

The CONON system [12] is an OWL encoded context ontology (CONON) for modelling context in pervasive computing environments. Its context model is split into into an upper ontology and other more specific ontologies. The upper ontology describes high-level features of basic contextual entities, of which, the most fundamental ones are location, person, activity and computational entity. Then each subdomain has a more specific ontology with additional details. It is implemented using Jena2 Semantic Web Toolkit and OWL-Lite. Reasoning is either: ontology reasoning using description logic or user-defined reasoning using first-order logic. COBRA-ONT [3] is an ontology used in the Context Broker Architecture (CoBrA) to facilitate knowledge sharing and context reasoning in ubiquitous computing. The system tries to determine location and status of agents (human or software) with in it. COBRA-ONT is expressed in OWL and models places, agents, and events. The ontology is categorised into 4 themes: 1) physical places, 2) agents (humans and software agents), 3) location text of the agents, and 4) activity context of the agents.

Outside of Ubiquitous Computing there has been work on extracting ontologies from existing text sources. To do this both concepts and relationships between them need to be learned. ConceptNet [13] is a massive ontology of commonsense knowledge. The concepts and relationships are extracted by processing the 70000 sentences of the Open Mind Common Sense Project. The sentences are elicited from the user in a semi-structured way in order to make the information easier to extract. Khan and Luo [14] focus on concept learning from a text corpus. Concepts are "placed" in a hierarchy. However, the type of relation between them is ignored; that is, it is only possible to tell two concepts are related, but not how they are related. Some projects, such as MindNet [15], Mecureo [16] and Janninik and Wiederhold's approach [17], focus on extracting relations between terms from dictionaries.

3. SYSTEM DESCRIPTION

The MyWorkPlace system provides users with personalised views of places, sensors, devices, services, objects and people in their environment. This advances on our earlier work, MyPlace, described in [2]. The key difference from earlier work is the inclusion of an automatically generated ontology to assist with the selection of items to show the user.

We now return to the scenario from the Introduction for interactions of Fred, a member of staff interacting with MyWorkPlace shortly after moving into the newly constructed School of IT smart building.

Fred is an academic who does not know very much about the building. He knows the location of his own office, and those of the students he supervises. He knows approximately where the the front counter is, but has not been there so is unaware of what facilities are available. He knows nothing of the seminar room, staff common room, pervasive computing laboratory, or undergraduate computer laboratories.

Figure 1 shows a screenshot from MyWorkPlace personalised for Fred, as it would be shown on a PDA while he is standing in the Foyer of the building.

Figure 1: The view Fred is presented by MyWorkPlace, as he is standing in the foyer (room 100), after inhabiting his new office building for a week.

The Status bar at the top tells Fred what the system believes his location and status is. In this case, his location is believed to be the Foyer, because the Mac address of his Bluetooth mobile phone has been detected there. There is a "details" button to allow him to scrutinise and correct the reasoning used for his location and status.

The content panel of the main screen consists of five expandable headings. The headings are *Devices at this location*, *Nearby Devices*, *Nearby Places*, *Services/Events*, and *People*. Clicking a heading shows or hides the contents. The "[Show all items]" button displays all the items the user is allowed to use. It also allows the user to see why an item was included or excluded by MyWorkPlace.

The system must determine which of the myriad of devices, sensors, places, services, events to display to Fred. For each heading, it examines the evidence and places each item into one of a number of *relevance categories*:

- *Already knows* - The user is believed to already know about this, based on either user feedback, or observations such as the use of a device, or being detected in a location.

- *Needs to learn* - The information is thought to be useful to the user and the user is believed not to already know it. Whether information is useful to a user is determined based on manually entered stereotypes, or the generated ontologies.

- *Needs to know now* - This is a special case of *Needs to learn* where the item is believed to be important based some aspect of the users' current context. For example: If the user has stated they are on their way to a seminar in a room they are believed not to know the location of, then location of the the seminar room is very important.

- *Not relevant* (Neutral) - The information is about something for which there is no information suggesting that it is useful to the user.
- *Doesn't want to know* - The user has indicated that they do not wish to be informed about this, or a very similar item.

The main screen shows items in the *Needs to learn* and *Needs to learn now* categories. The user can override this personalised selection of information to see all items and their relevance categories by choosing the "[Show all items]" button.

The *Nearby Places* category in Figure 2 lists a number of the places which Fred does not know about. There are many others which he is not informed about as the system does not believe they are relevant to him. For example with suitable information in the system in the system it might be able to omit details of the Undergraduate Laboratories, as it is semester break so he is not currently teaching classes.

Figure 2 shows the view when Fred returns to his office. Here are more devices he may wish to learn about. Clicking on an items brings up more information about the object. It includes usage instructions and troubleshooting information. In addition to this, it includes links to related items, as suggested from the automatically generated ontologies. Each time Fred clicks a link requesting more information, a piece of evidence is added to his user model suggesting he knows about it.

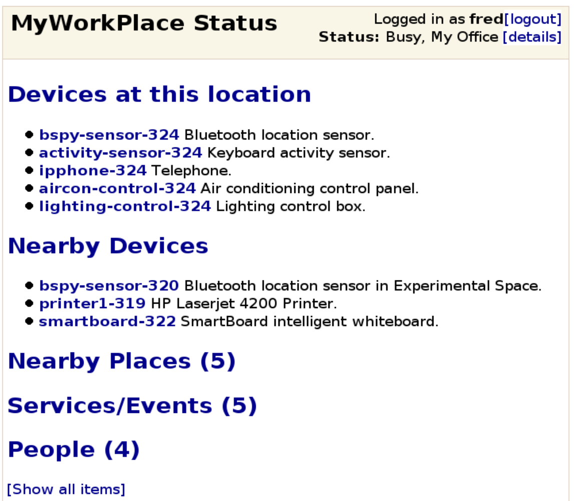

Figure 2: The view Fred is presented by MyWorkPlace, when he returns to his office (room 324).

It is important for the user to be able find out why a certain item has been displayed to them, and others have not. We call this *scrutability*. When a user clicks the "[Show all items]" button the full list of items is colour coded according to which of the five relevance categories defined above it belongs. An explanation of the reasoning used to categorise each item is shown if the user hovers the mouse over it. An example explanation might be "You are teaching Algorithms 101, this room is used for Algorithms 101, and you have not yet been detected there."

4. ONTOLOGY GENERATION

The data in our ontology is being built from a number of sources, such as building plans, staff directory, the building manual, student timetables, the departmental calendar, room bookings, and a relatively small, handcrafted base ontology. The degree of automation used and level of user input required in generating ontological information from these sources varies considerably. The reliability of each source also varies for a number of reasons, such as input errors and frequency of maintenance. MyWorkPlace takes account for the inaccuracy problem with its evidence accretion and delayed resolution approach [2]. This means that it can apply simple, explainable reasoning processes for dealing with conflicting and noisy information.

Our initial source for location relationships were the building plans for each floor of the building. Features on the plans are grouped in the relevant layers. For example all the room number labels are in one layer, the room description texts are in another, another layer holds all the doors, one layer holds all the solid walls while another holds all the glass. There are over 100 different layers in total.

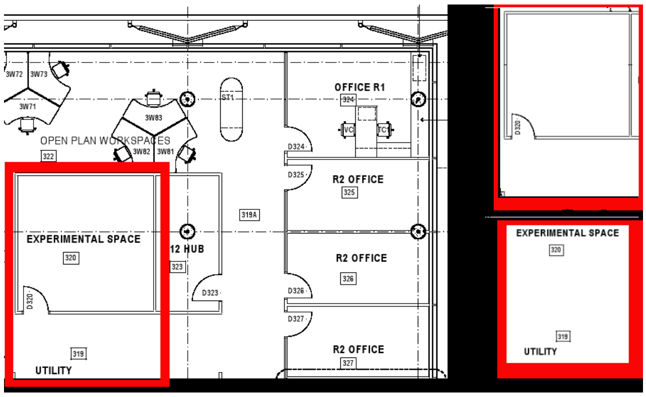

Figure 3: A section of the building plans with only selected layers displayed.

Analysis of the lines and text data in each layer allows us to use relatively simple reasoning to determine rooms and relationships between them (e.g. distance). Figure 3 shows a section of the plans (left-hand side sub-figure), and further extracts showing: a) only wall and door layers (top right-hand sub-figure), and b) only room number and label layers (bottom right-hand sub-figure). The majority of the data generated from these plans is assumed to be very reliable as the building was built according to them.

The departmental staff directory is also an important source for automatic population of an ontology and user models. The staff directory yields a list of all academics, administrative staff and postgraduate students. It also gives relationships between people and their research groups (e.g. undergraduate coordinator, chair of a research group), office or workspace locations, and contact information. A number of the student-supervisor links are also available in this database. This data can be extracted with minimal human interaction. However, it suffers from some inaccuracies as the phone list is not always kept completely up to date, for example, when people move offices or change research groups.

Another source of ontological data is the building manual. This documents various elements about the building in a glossary and human readable format. Example entries in the manual are an entry for "fax" that describes policies for using a fax machine, and an entry for "reception" that gives the opening hours and the location of the reception desk. We believe we will be able to use Sago, a descendant of Mecureo [16], to extract implicit relationships from this manual to populate the base ontology. Mecureo is an ontology learning tool that takes a glossary and mines relationships between the concepts, or glossary terms.

We propose to use other sources, such as: the university timetable that lists times and venues for every class in the university; the departmental calendar and room bookings that give scheduled times and locations for certain activities. Both sources may help generate more evidence that assists the MyWorkPlace system to deduce, say, current location and activity for a person, and suitable activities for certain venues. In future we may link databases of particular class enrollments to generate views of MyWorkPlace for undergraduate students. However, this information might not be as reliable as some of the other sources.

The final source is a small, handcrafted base ontology. This serves as a base for other sources to build on.

5. DISCUSSION

At the time of writing the new building for which the system has been implemented has only just been occupied. We plan to evaluate this system by two methods: firstly we will ask a number of users to perform a task requiring knowledge of the building, and compare the performance of a group of people given access to MyWorkPlace against those who are only given access to the data sources used to populate the ontology (probably building plans, building manual, and phone list).

The other form of evaluation we wish to perform is to make MyWorkPlace available to volunteer academic staff and postgraduate students, and record their use of the system. In addition to this we will request feedback from a sample of users.

One aspect we are particularly interested in is how any inaccuracies in the generated ontological data affect the user experience. As mentioned previously some aspects of this data are known to be very accurate, while others are expected to contain some out of date or incorrect items.

Another piece of future work is to investigate the use of a suggested template and categories for entries in the building manual to better facilitate information extraction. We would also like to include data extracted from existing, maturer ontologies.

6. REFERENCES

[1] Weiser, M.: The computer for the 21st century. Scientific American **265**(3) (1991) 94–104

[2] Carmichael, D., Kay, J., Kummerfeld, R.: Consistent modeling of users, devices and environments in a ubiquitous computing environment. User Modeling and User-Adapted Interaction **15** (2005) 197–234

[3] Chen, H., Finin, T., Joshi, A.: An ontology for context-aware pervasive computing environments. Knowledge Engineering Review **18**(3) (2004) 197–207

[4] Christopoulou, E., Goumopoulos, C., Kameas, A.: An ontology-based context management and reasoning process for ubicomp applications. In: sOc-EUSAI '05: Proceedings of the 2005 joint conference on Smart objects and ambient intelligence, New York, NY, USA, ACM Press (2005) 265–270

[5] Gu, T., Wang, X.H., Pung, H.K., Zhang, D.Q.: An ontology-based context model in intelligent environments. In: Proceedings of Communication Networks and Distributed Systems Modeling and Simulation Conference 2004. (2004)

[6] Heckmann, D., Krueger, A.: A user modeling markup language (userml) for ubiquitous computing. In Brusilovsky, Corbett, d.R.e., ed.: LNAI 2702: Proc. of the Ninth Int. Conf. on User Modeling (UM'03), Springer (2003) 393–397

[7] Heer, J., Khooshabeh, P.: Seeing the invisible. In: Workshop on Invisible and Transparent Interfaces, part of Advanced Visual Interfaces (2004). (2004)

[8] Edwards, W.K., Grinter, R.E.: At home with ubiquitous computing: Seven challenges. In: Proceedings of the Conference on Ubiquitous Computing (Ubicomp 2001). (2001)

[9] Chang, A.: Afaik: A help system for the intelligent room. Master's thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (2001)

[10] Rerscha, A., Keller, M.: Digiscope: An invisible worlds window. In: Video in 5th International Conference on Ubiquitous Computing. (2003)

[11] Krumm, J., Hinckley, K.: The nearme wireless proximity server. In Davies, N., Mynatt, E.D., Siio, I., eds.: UbiComp 2004: Ubiquitous Computing: 6th International Conference, Nottingham, UK, September 7-10, 2004. Proceedings. Volume 3205 of Lecture Notes in Computer Science., Springer (2004) 283–300

[12] Wang, X.H., Zhang, D.Q., Gu, T., Pung, H.K.: Ontology based context modeling and reasoning using owl. In: Proc. of Workshop on Context Modelling and Reasoning, IEEE PerCom '04. (2004)

[13] Liu, H., Singh, P.: Conceptnet - a practical commonsense reasoning tool-kit. BT Technology Journal **22**(4) (2004) 211–226

[14] Khan, L., Luo, F.: Ontology construction for information selection. In: ICTAI '02: Proceedings of the 14th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'02), Washington, DC, USA, IEEE Computer Society (2002) 122

[15] Richardson, S.D., Dolan, W.B., Vanderwende, L.: Mindnet: Acquiring and structuring semantic information from text. In: Proc. of COLING '98. (1998)

[16] Apted, T., Kay, J.: MECUREO: Ontology and modelling tools. WBES of the International Journal of Continuing Engineering Education and Lifelong Learning **14**(3) (2004) 191–211

[17] Jannink, J., Wiederhold, G.: Ontology maintenance with an algebraic methodology: A case study. In: AAAI Workshop on Ontology Management. (1999)