Ambient Documents: Intelligent Prediction for Ubiquitous Content Access

Gregory M.P. O'Hare[1], Michael J. O'Grady[1], Conor Muldoon[1], and Caroline A. Byrne[2]

[1] Adaptive Information Cluster (AIC), University College Dublin (UCD), Belfield, Dublin 4, Ireland
[2] Carlow Institute of T Carlow, Kilkenny Road, Carlow, Ireland
{Gregory.OHare,M.J.OGrady,Conor.Muldoon}@ucd.ie, Caroline.Byrne@itcarlow.ie

Abstract. Ubiquitous service delivery expects that content will be available where, when and how the user needs it. Consumers are becoming ever demanding and the consumers of ubiquitous services are no different in this regard. Their expectations escalate in terms of relevance, ease of access, recency, accuracy and latency of content supply. In addition they expect that the content be supplied proactively in anticipation of their needs together with delivery when they require it. This presupposes that content can be delivered relative to both the consumers location and their technological context. Within this paper we explore how traditional document access can be transformed and introduce *Ambient Documents* a new metaphor for document content access.

Keywords: Ubiquitous computing, Ambient Intelligence, ambient documents, intelligent agents.

1 Introduction

Ubiquitous computing seeks to attain a world of embedded computational infrastructures and artefacts where people can access services in an almost anytime-anywhere basis. This is an attractive vision and has motivated researchers for almost twenty years. Much of this research has focused on the enabling technologies – sensors, wireless networking technologies and on on. Significant progress has been made, and though it cannot be said that all problems have been resolved, it can be assumed that sufficient progress has been made such that practical issues such as data access, usability and so forth can now be explored. One initiative that aims to take accessibility and usability to the next stage is Ambient Intelligence (AmI) [1].

AmI was conceived in response to the realisation that ubiquitous computing systems could easily become unusable. A networked environment of artefacts competing for the user's attention could quickly result in the user being overwhelmed, and as such, being tempted to *switch-off*, both mentally, and, in the case of the ubiquitous computing system, by physically shutting the system down. Thus, AmI seeks to introduce intelligent techniques as a means of mediating between the user

C. Stephanidis (Ed.): Universal Access in HCI, Part II, HCII 2007, LNCS 4555, pp. 971–979, 2007.

and the services supported by the ubiquitous environment, some of which will involve access to a variety of electronic content. In this paper, the concept of an Ambient document as a means of document content access is introduced.

This paper is structured as follows. In the Section II, a description of the ambient document concept is provided. Two practical illustrations of ambient document construction and provision is provided in Section III. In Section IV, the use of intelligent techniques for anticipatory document provision and construction is outlined. Some related research is presented in Section IV, after which the paper is concluded.

2 Ambient Documents

Traditional document structure and access is predicated upon certain assumptions:

- Document retrieval is based on user requests (pull technology);
- Documents are generic;
- Documents are static;
- Documents are primarily comprised of text and images;
- Document content is always regulated;

Ambient documents radically challenge each of these assumptions. Such documents envisage a world where ambient documents exist almost as overlays upon the physical world. These overlays enable documents to be accessed in an ambient nature with content acquired and delivered in an intuitive manner.

Ambient Document Retrieval
The traditional pull model may still be used with Ambient Documents. However, it also incorporates the push/subscribe model to a far greater degree. In particular, the triggering parameters for pushing a document are intrinsically linked with various aspects of the user's context, and may be highly domain dependant. Ambient documents may be also anticipatory in that they can be dynamically constructed in anticipation of a future requirement, which may or may not arise.

Ambient Document Audience
Traditional documents are generic in the sense that they are written for a certain audience. In contrast, Ambient Documents are personalised for their audience, either as a group or individually. Their content is filtered and adapted such that it is of relevance to its readers and, in particular, to the pertinent aspects of the readers' profile that are dominant at the time of invocation.

Ambient Document Content
Ambient Documents are inherently dynamic entities. Their content is highly dependent of the prevailing situation or context of the user. What aspects of the user's context are of most significance depend on the application domain and the task being performed. Location is one obvious element but there are others. For example, if the document should include a forecast of the weather, then this must be gleaned from an appropriate source such that it is relevant to the current time and place.

Ambient Document Media
Traditionally, documents consisted of text and images. The WWW has challenged this assumption, and views documents as comprising various multimedia elements. Ambient Documents also subscribes to this view, though the selection of media elements is highly dependent on the user's context.

Ambient Document Ethos
Traditional documents are inherently regulated in their content and their structure. Ambient documents are unregulated in that their content and structure is highly dynamic. A typical example of this is a blog where the number of people contributing is unpredictable, and the contributed content varies in diverse ways.

Ambient Document Repositories
Ambient documents are inherently distributed and do not use a central repository or digital library. Such documents are frequently deployed at the very edge of networks and on devices of limited computational capabilities such as sensors.

Ambient Document Presentation
Ambient Documents do not presuppose the availability of a standard desktop workstation with a standard WIMP interface. Rather such documents can be accessed and viewed in a multimodal fashion. This may involve voice and gesture recognition for the input modality. For the display modality, audio and video may be used in cooperation with animated or virtual characters, via a range of physical media ranging from a smartphone interface to an Augmented Reality (AR) system.

3 Ambient Document Construction and Delivery

Constructing and delivering ambient documents to individuals and groups raises a number of challenges for designers, software engineers and service providers. The complexity of the process results from a combination of diverse target user groups, including experience and expectations, as well individual attributes such as language, age group and so on; diverse supporting infrastructure, including the capabilities of any embedded or mobile devices and prevailing networks; and other contextual elements, including location, time and so on. As an illustration of how such documents may be constructed and delivered, two examples are provided: one that conforms to a regulated ethos, and a second that is inherently unregulated.

3.1 Tourist Information Systems

Tourists are an example of an end-user group that is heterogeneous and mobile. As tourists explore their surrounding, an opportunity exists to deliver information to them in a timely basis. A number of systems have been described in the literature that aims to this, for example CRUMPET [2]. For the purposes of this discussion, the method adopted for Gulliver's Genie [3] [4] is presented and discussed in light of the ambient document concept.

When a tourist encounters a popular attraction, the Genie makes a presentation on this attraction available to them on their host device. The presentation itself utilises a

combination of multimedia elements, namely, audio, images video and text. However, it is the nature and construction of the presentation that is of most interest. First of all, the context of the presentation is always relevant to the location and orientation of the tourist, that is, their spatial context. Secondly, the content is personalised for the tourist in question in two distinct ways. Initially, content is selected based on the tourist's personal profile. For example, an adult may be presented with material that would be pitched at level that would not be suitable for a child. After this initial pruning of content, the content is further filtered such that it is consistent with the tourist domain-specific model, in this case their cultural interests. Thirdly, and finally, all content is prioritised such that it coincides with the level of interest expressed by the tourist in previous interactions. The practical implication of this is that a dynamic model must be maintained at all times. Thus the presentation incorporates a number of diverse elements including a location-aware element, a content personalisation and customisation element, and a broad context-aware element.

If the ambient document attributes as listed in Section II are considered, it can be seen that the approach adopted by the Genie concurs with these. Gulliver's Genie adopts a data-push approach, although explicitly with the permission of the end-user. In a business model where the end-user pays, this is of critical importance. The Genie is designed to address the needs of a specific audience, namely tourists. However, the content is inherently dynamic, and adapted according to the location and profile of the tourist in question. Furthermore, this profile has a specific domain component, in this case a cultural interest sub-model. From a structural perspective, a range of multimedia elements are supported. Media browsing is nominally via standard navigation-pad manipulation, though activating the availability of the presentation itself can only be achieved by physically going the attraction in question. Hence, the interaction modality is multimodal, and incorporates both an implicit and explicit element. Technically, the content can be hosted on a centralised server, though a distributed approach is preferred. All content must be approved a priori, thus the presentation is regulated. Indeed, the adequate tagging of the content so that it can be matched with interested tourists is of fundamental importance.

However, the tourists themselves are oblivious to the ongoing suite of ongoing processes necessary to realize the service, as indeed they should be. From their perspective, a series of multimedia documents are presented to them for their perusal, at a time and in a location where they are of most interest. In this way, documents assume an ambient nature, and in the tourists' eyes, become perceived as a specific information overlay superimposed on the physical environment which can only be accessed at certain locations, which in this case, coincides with the location of tourist attractions.

3.2 GeoBlogging

Tourist activities are frequently solitary, and in an effort to address this, a blogging application was designed [5]. Blogging is a well known internet activity, but in this case, a geospatial element was added, hence the term geoblogging. The core idea is quite simple. As tourists explore their environment, they can access individual blog entries that have been made in their vicinity, or they can construct a blog entry themselves.

A blog entry may contain a combination of media elements, and it is at the discretion of the tourist as to which elements to use. Geospatial data and tourist profile are the two key parameters used for tagging the blog entry. Both will be used in the retrieval process: the geospatial data for matching the tourist's current geospatial context, and the tourist profile for filtering and prioritizing blog entries that are likely to be of most interest to the tourist.

Much of what was discussed previously about the Genie and ambient documents applies to the geoblogging applications. However, there is one critical difference. Blogging is an inherently unregulated activity. Where a blog entry is created, and accessed, cannot be predicted. The media elements used cannot be predicted. And the capability of the tourists' devices to render individual blogs must be factored into the retrieval process.

3.3 Intelligent Prediction

Ambient documents are almost invariable constructed on the fly, for example, the Genie presentations, and blog listings just described. This is a time consuming exercise, and one that could easily contribute to a poor end-user experience. One method of mitigating this difficulty concerns the predicting of content requirements in advance, and the subsequent preparation of content in anticipation of short term future access. Fundamental to this is a sophisticated model of the tourist's physical environment as well as an awareness of their spatial context. Based on this, their future route can be estimated and their future information needs anticipated. However, a degree of intelligence is necessary to do this. Gulliver's Genie adopts an approach termed intelligent precaching [6] as a mechanism for anticipating content requirements and the pre-caching of this content in the expectation that it can be pushed to the tourist. In the next section, a brief overview of approaches that can be adopted for using intelligent techniques for the production and dissemination of ambient documents is discussed.

4 Intelligent Techniques for Ambient Documents

There is a potentially serious conundrum at the heart of AmI. Ubiquitous computing is characterised by small lightweight distributed sensors and devices that are inherently limited from a computational perspective. However, intelligent techniques usually demand sophisticated software that will only operate satisfactorily on workstation-type environments. Reconciling these two issues requires significant effort and ingenuity from software engineers. Furthermore, even to make maximum use of the limited computational resources available calls for the facilitation of intelligent decision making. It is essential that these systems operate effectively and use efficient algorithms to minimise computational overhead. There are two broad approaches to achieving this, namely centralised and distributed.

– Centralised Intelligence can be centralised on a fixed network node. In this case, computational resources are not a serious problem. Thus the software designer has the luxury of choosing the AI techniques that are most appropriate. A key difficulty concerns system responsiveness. If, as may frequently be the case, there is a

wireless networking component, this may become a bottleneck in the system, particularly if a significant amount of data must be transferred between a central network node and the host device to generate the predictive behaviour. This may have a detrimental effect on the usability of the system as a whole – something that should be avoided at all costs. Should the application seek to push data to the user, a small window of opportunity may exist to do this which must be taken advantage of. If not, the user may have moved on and an opportunity will be lost.

– Distributed Intelligence represents a second strategy for realizing the necessary intelligence for ubiquitous content access. In this case, the intelligence may be distributed throughout the network. One can easily envisage the sophisticated algorithms being run on fixed network nodes while some of the less computationally expensive algorithms may be deployed on the user's mobile device, or even on an embedded sensor. A practical implementation of the distributed AI approach is that of Multi-agent Systems (MAS). A MAS takes a decentralised approach to the problem of deploying intelligence in a distributed network. Though traditional MASs primarily targeted desktop environments, there have been a number of agent platforms reported in the literature that target mobile devices, and one of these, AFME, is described briefly in the following section. From an economic perspective, the number of cellular digital mobile phones is already greater than the number of desktop and notebook computers [7] and continues to grow at a considerably faster rate, particularly in developing countries. Thus the potential for the deployment of MASs in resource constrained computational environments is greater than that of networked desktop computers and laptops.

4.1 Agent Factory Micro Edition

Agent Factory Micro Edition (AFME) [8] is a minimised footprint agent platform specifically designed for resource constrained mobile devices. It is loosely based on Agent Factory [9], a pre-existing agent platform for personal computers. It uses a

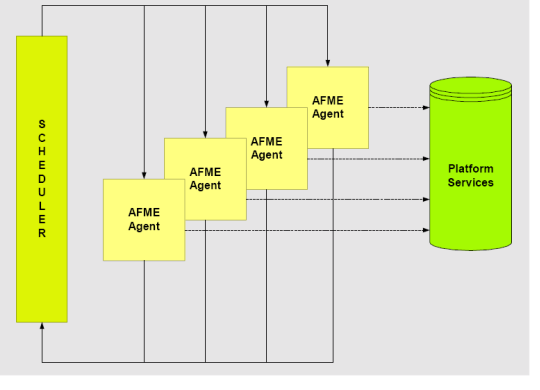

Fig. 1. A rchitecture of AFME

subset of the Agent Factory Agent Programming Language (AFAPL) and augments it with a number of features specific to AFME. AFAPL is founded on a logical formalism of belief and commitment. Rules that define the conditions under which agents should adopt commitments are used to govern and encode agent behaviour.

An AFME platform (Fig. 1) comprises a scheduler, several platform services, and a group of agents. The scheduler is responsible for the scheduling of agents to execute at periodic intervals. Rather than each agent creating a new thread when they begin operating, agents share a thread pool.

AFME delivers support for the creation of Belief, Desire, and Intention (BDI) [10] agents that follow a *sense-deliberate-act* cycle. Beliefs are, perhaps inaccurate, information that an agent has about the world. Desires are states in the world that agents wish to bring about. Intentions are desires that agents have committed resources to achieving. The control algorithm performs four functions. Preceptors are fired and beliefs are resolved within the belief resolution function. The beliefs are used within the deliberation process to identify the agent's desired states. Agents are resource bounded and will be unable to achieve all of their desires even if their desires are consistent. A subset is chosen, in the *intention selection process*, that maximises their self-interest with respect to their finite resources. The final function of the control algorithm concerns commitment management. Depending on the nature of the commitments adopted various actuators are fired.

In AFME agents are rational; sometimes they collaborate other times they compete depending on the context. In certain circumstances the developer will want to ensure that agents always behave in a particular manner. This type of behaviour is modelled, within AFME, through the use of fixed utility values that ensure that under such circumstances it is always in an agent's interest to behave in a certain fashion.

5 Related Research

Ambient documents are closely related to ubiquitous user interfaces, a number of which have been described in the literature. Most of these interfaces are concerned with the nature of the device, while ambient documents adopt a more holistic view, emphasising the construction process, though not at the expense of the rendering and visualisation element.

The Pervasive Adaptive Visualization and Interaction Service (PAVIS) is an approach to performing adaptive visualisation and user-interface generation based on device-specific information, such as resolution, screen size, and networking capability [11]. It aids the user in visualising data and interpreting information. The form visualisation takes has a considerable impact on how effectively the information is interpreted. The goal of PAVIS is to automatically generate user-interfaces that are best suited for the display device.

The User Interface Mark-up Language (UIML) [12] is an approach to simplifying the construction of multi-platform user interfaces. The goal of UIML is to reduce the development time of user interfaces for multiple device families through the construction of a generic interface vocabulary.

The Personal Universal Controller (PUC) [13] system automatically generates an intermediary interface for an appliance, such as stereos, copiers, elevators, and the

non-driving functions of a car. PUC allows users to control appliances in their environment through a remote user interface. The remote control may be a graphical user interface on a PDA or mobile phone or speech recognition system that uses microphones in the room. The system provides a two-way communication channel between the appliance and the controller. This facilitates the transmission of the user's commands and enables the controller to receive feedback.

ICrafter [14] is a framework for the construction of *services* and ubiquitous user interfaces. The term services when used in this context is referring a devices, such as a light or a projector, or an application, such as Microsoft PowerPoint operating on a large screen, that provides useful functionality to users. The framework is concerned with the development of what are known as interactive workspaces. An interactive workspace is a technology rich space of interconnected computers, utility devices, and IO devices, where people come together to collaborate on a particular task, such as brainstorming.

6 Conclusion

As ubiquitous computing systems proliferate, the practical issues of data preparation, dissemination and access in such systems becomes increasingly important. The traditional view of documents as inherently static entities limits the potential and possibilities envisaged by the ubiquitous computing and AmI visions. In this paper, the concept of ambient documents has been introduced. Such documents offer a dynamic and radical view of document access, production and dissemination. Essential to the ambient document construct is the availability of a suite of intelligent techniques for content identification and filtering, as well as the prediction of future document usage. Having illustrated the potential of ambient documents in a mobile tourist information and blogging system, it is the authors' intention to investigate further the use of alternative interaction modalities as means of accessing and retrieving and navigating such documents.

Acknowledgments. This material is based upon works supported by the Science Foundation Ireland (SFI) under Grant No. 03/IN.3/1361.

References

1. Vasilakos, A., Pedrycz, W.: Ambient Intelligence, Wireless Networking, Ubiquitous Computing, Artec House (2006)
2. Poslad, S., Laamanen, H., Malaka, R., Nick, A., Zipf, A.: Crumpet: Creation of user-friendly mobile services personalised for tourism, Proceeding of the Second IEE International Conference on 3G Mobile Communication Technologies, London, UK (2001)
3. O'Hare, G.M.P., O'Grady, M.J., Gulliver's Genie, A.: Multi-Agent System for Ubiquitous and Intelligent Content Delivery, Computer Communications, vol. 26 (11) pp. 1177–1187 (2003)

4. O'Grady, M.J., O'Hare, G.M.P., Sas, C.: Mobile Agents for Mobile Tourists: A User Evaluation of Gulliver's Genie, Interacting with Computers Journal, vol. 17 (4), pp. 343–366 (2005)
5. Chen, J., O'Grady, M.J., O'Hare, G.M.P.: Autonomy and Intelligence - Opportunistic Service Delivery in Mobile Computing. In: Gabrys, B., Howlett, R.J., Jain, L.C. (eds.) KES 2006. LNCS (LNAI), vol. 4253, pp. 1201–1207. Springer, Heidelberg (2006)
6. O'Grady, M.J., O'Hare, G.M.P.: Just-In-Time Multimedia Distribution in a Mobile Computing Environment. IEEE Multimedia 11(4), 62–74 (2004)
7. Bratt, S.: Mobile Web Initiative, World Wide Web Consortium (W3C), MIT Research and Development Conference, Innovation for Industry (2005)
8. Muldoon, C., O'Hare, G.M.P., Collier, R., O'Grady, M.J.: Agent Factory Micro Edition: A Framework for Ambient Applications. In: Conti, R., Ruberti, A. (eds.) Proceedings of Intelligent Agents in Computing Systems (IACS 2) Workshop held in Conjunction with International Conference on Computational Science (ICCS) 2006. LNCS, vol. 3, pp. 727–734. Springer, Heidelberg (1973)
9. Collier, R.W., O'Hare, G.M.P., Lowen, T., Rooney, C.F.B.: Beyond prototyping in the factory of agents. In: Mařík, V., Müller, J.P., Pěchouček, M. (eds.) CEEMAS 2003. LNCS (LNAI), vol. 2691, pp. 383–393. Springer, Heidelberg (2003)
10. Rao, A.S., Georgeff, M.P.B.: Agents: From Theory To Practice. In: Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95), San Francisco, CA, pp. 312–319 (1995)
11. Alimohideen, J., PAVIS,: – Pervasive Adaptive Visualization and Interaction Service, CHI Workshop on Information Visualization and Interaction Techniques for Collaboration Across Multiple Displays, Montreal, Canada (2006)
12. Ali, M.-F., Marc Abrams, M.: Simplyfying Construction of Multi-Platform User Interfaces Using UIML, UIML Europe 2001 Conference (2001)
13. Nichols, J.: Automatically Generating User Interfaces, Pervasive Computing, Vienna, Austria (2004)
14. Ponnekanti, S.R., Lee, B., Fox, A., Hanrahan, P., Winograd, T.I.: A Service Framework for Ubiquitous Computing Environments. In: Abowd, G.D., Brumitt, B., Shafer, S. (eds.) Ubicomp 2001: Ubiquitous Computing. LNCS, vol. 2201, pp. 56–75. Springer, Heidelberg (2001)