Research Repository UCD

Title	Distributed network intelligence : a prerequisite for adaptive & personalised service delivery
Authors(s)	O'Grady, Michael J., O'Hare, G. M. P. (Greg M. P.), Chen, Jiangyan, Phelan, Donnacha
Publication date	2009-03
Publication information	O'Grady, Michael J., G. M. P. (Greg M. P.) O'Hare, Jiangyan Chen, and Donnacha Phelan. "Distributed Network Intelligence : A Prerequisite for Adaptive & Personalised Service Delivery" 11, no. 1 (March, 2009).
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/1137
Publisher's version (DOI)	10.1007/s10796-008-9140-1

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Distributed Network Intelligence: A Prerequisite for Adaptive & Personalised Service Delivery

M. J. O'Grady, G.M.P. O'Hare, J. Chen, D. Phelan *School of Computer Science &*

Informatics University College Dublin

Belfield

Dublin 4

Ireland

Mobile computing is undoubtedly one of the predominant computer usage paradigms in operation today. The implications of what might be cautiously termed a usage paradigm shift have still not crystallised fully, either for society, or those envisaging a new raft of applications and services for mobile users. However, fundamental to the current and future success of mobile computing are mobile telecommunications networks. Such networks have been a success story in their own right in recent years, both as traditional voice carriers and, increasingly importantly, as a conduit of mobile data. The potential for new mobile data applications is immense, but, crucially, this potential is severely compromised by two factors inherent in mobile computing: limited bandwidth and computationally restricted devices. Hence, the academic and commercial interest in harnessing intelligent techniques as a means of mitigating these concerns, and ensuring the user experience is a satisfactory one. In this paper, the broad area of intelligence in telecommunications networks is examined, and issues relating to the deployment of intelligent technologies are explored. In particular, the potential of intelligent agents is identified as a viable mechanism for realising a full end-to-end deployment of intelligence throughout the network, including possibly the most crucial component: the end user's device. As an illustration of the viability of this approach, a brief description of a mobile blogging application is presented.

Keywords

Mobile computing; mobile telecommunications; intelligent agents; context-aware computing; personalisation.

1 Introduction

A number of approaches have been proposed in an effort to identify the most appropriate method of supporting the average mobile user in the course of their everyday activities. Ubiquitous computing (Vasilakos and Pedrycz, 2006), conceived in the early 1990s, advocates the embedding of computational artefacts in the environment, enabling pervasive access to computational resources. Wearable computing (Bardfield and Caudell, 2001) takes a different approach, observing that the explicit use of computing is rarely an end-goal in itself, and that computing should augment the user in the pursuit of their everyday activities. Between both these extremes of the mobile computing spectrum, there exists the mobile computing paradigm that most people would be familiar with: laptops, PDAs and mobile phones. Inherent in this paradigm is the availability of a wireless data communications facility. Hence, the increasing perception that mobile computing and mobile telecommunications are practically indistinguishable.

Mobile telecommunications, once synonymous with voice traffic, is becoming increasingly perceived as a bearer of wireless data. Naturally, the desire of subscribers for 24/7 internet access has been a major factor in this focus shift. While the market is enormous, network operators are compromised by the limitations of the average mobile computing devices, and the poor data-rates supported by the current generation of mobile telecommunications networks. Though the situation continues to improve, there will invariably be a discernable disparity between the experience on the computationally-rich fixed network, and the relatively computationally-poor mobile network. This raises a particular difficulty for network operators, as subscribers cannot help but notice this difference. Furthermore, subscribers may query why they are paying for what they perceive to be a sub-standard service. Thus, the objective of minimising this performance differential while improving subscribers' perceived experience of the technology is one that all mobile operators share. While there is no obvious answer to this problem, the selective use of intelligent techniques is seen as one way in which the experience of the subscriber may be improved, and the technological limitations partially mitigated.

This paper considers the use of intelligent techniques for mobile computing, and use of intelligent agents in particular, as a viable mechanism for introducing intelligence on the last frontier of the mobile network, namely the mobile device. In the next section, a broad overview of the use of intelligent techniques in mobile networking is presented. From a subscriber's perspective, the dynamic adaptation and personalisation of media content according to their contextual situation is an attractive proposition. It will be argued that appropriate harvesting of intelligent techniques offers a viable vehicle for fulfilling this proposition. To illustrate the pertinent issues, a mobile blogging scenario is considered, and its design and realisation through the intelligent agent paradigm presented.

2 Embedding Intelligence into the Network

Historically, telecommunications network operators have understood the benefits of deploying intelligent techniques *throughout* their networks. A number of benefits were envisaged accruing from such an approach:

- New services could be introduced throughout the network quickly and efficiently without the need for physical intervention by staff.
- Services could be customised to individual customer requirements.
- Vendor independence would be established allowing operators to mix and match their services according to their customers' requirements.

Significantly, this would also allow network operators to differentiate their services and target particular niche markets.

- Open interfaces could be created allowing new entrants into a market that has traditionally been dominated by the major telecommunications manufacturers.

Intelligent Networking (IN) and, in some case, Advanced Intelligent Networking (AIN) became popular terms in telecommunications marketing; so much so that the International Telecommunications Union (ITU) proceeded to define a set of standards (Q.1210 to Q.1219) frequently referred to as Capability Set One (CS-1) (ITU, 1997). In essence, the IN initiative was an effort to separate service creation from the core network switching function, and this initiative gathered pace throughout the late 1980s and 1990s as developments in microprocessor technologies developed. Hence, core network functions could be distributed throughout the network, and particularly at the periphery of the network, that is, nearest the customer. For the customer, the implications of IN technologies were immense. Services such as toll free calling, prepaid calling, and call screening amongst others became possible and their popularity has increased to such an extent that people now take them for granted. However, a momentous development in networking technologies rapidly changed the telecommunications landscape: the Internet.

Developments in internet technologies had a profound effect on those institutions closely connected with the broad computing and telecommunications industries. Indeed, telecommunications companies were deeply affected by an almost parallel development: mobile computing, and implicitly, wireless networking. In short, two paradigm shifts may be identified. The first concerns the increased importance of data traffic as a revenue generator, as distinct from that of voice, typified by the Plain Old Telephone Service (POTS). The second concerns the nature of telephony services changing from one that was inherently static in nature to one that is inherently dynamic.

Computer usage has evolved from being a static activity to one that is increasingly mobile; and mobile computing may well become the de facto computer usage paradigm in the near future. Hence, the increasing demand for mobile data services, and the ongoing augmentation of mobile telecommunications networks with facilities which can meet this demand. As a consequence we are witnessing a blurring of the boundaries between telecommunications and computing is taking place, with computing technologies being increasingly integrated into telecommunications networks. Nowhere is this more evident than in the actual handsets which are in essence, mobile computing devices augmented with a real-time voice transmission component. As the penetration of computing into telecommunications networks continues, possibilities for the use of sophisticated computational technologies arise: intelligent techniques being a particular case in point.

A mobile network typically comprises a number of essential elements. Though individual manufacturers may use their own terms for these components, the interfaces between them are standardised by the ITU or another approved body such as the Third Generation Partnership Projects (3GPPs). A schema of a typical mobile network is illustrated in Figure 1. A detailed description of the various components is beyond the scope of this discussion; however, it is suffice to observe that all components are augmented or can be augmented, with significant computational processing capability. Thus, the possibility exists for harnessing any combination of Artificial Intelligence (AI) techniques on individual components, as the task at hand requires - the one exception to this is the actual mobile device. The implications of this being that AI techniques cannot be deployed, at least not without significant difficulty, where they may potentially have the greatest effect.

Figure 1. A simplified mobile network topology.

2.1 Intelligent Solutions for Mobile Devices

Implementing AI techniques on mobile devices is feasible, particularly as the raw computational power of the average mobile device continues to increase. The key difficulty is the time the reasoning process takes. This will normally be of several orders of magnitude longer than an equivalent operation would take on a normal workstation; the net result being that the user experience will be adversely affected. Users are conditioned to expect almost instantaneous responses, and will interpret a lax system response as a fault in the software, and may be reluctant to use that particular application or service again. Thus software designers in the mobile computing domain have a formidable task in reconciling the conflicting demands of maintaining an adequate response time, or Quality of Service (QoS), within the technological constraints of the average mobile device. This is irrespective of whether or not intelligent techniques are used; and the non-judicious usage of such techniques could exacerbate the situation yet further.

Before reflecting on how an AI solution might be realised on a physical mobile device, there is another strategy that could be adopted in appropriate circumstances. Recall that the mobile device is at the extreme periphery of the network, and that transmission of data over the air interface, that is, between the supporting base station and the device, is an acknowledged bottleneck. However, the introduction of new techniques in 3G networks, and the imminent deployment of are frequently termed 3.5G technologies such as High Speed Download Packet Access (HSDPA), will significantly increase the available data rates, making mobile broadband networking a reality. A situation can be envisaged whereby the AI component resides on the serving base station and the connected devices use it in an as-needed basis. Assuming that a significant data transfer is unnecessary, such an approach is feasible, and may satisfy the needs of certain applications without compromising the perceived response time.

Which AI technique should be adopted is ultimately determined by the nature of the application. AI techniques are not interchangeable; each is oriented towards certain classes of problem. By reflecting on the distributed nature of the mobile network, and the need for intelligence at various nodes, an initial inkling of the most appropriate solution arises. Distributed AI (DAI) arose in response to a need for a solution to distributed problem solving. Delivering a service to a subscriber requires significant collaboration on behalf of the various components of the network. Hence, it can be seen that a DAI approach offers an intuitive mechanism for deploying AI solutions across the network. In practice, DAI

solutions are mostly realised through the Intelligent Agent paradigm or Multi-Agent Systems (MAS).

Agent characteristics are many, but they coalesce around a number of core features. At the centre is a reasoning engine. How sophisticated this is remains at the discretion of the software designer. It may incorporate machine learning principles, Case Based Reasoning (CBR) techniques or neural networks. Again, the task or objective of the software solution is all important. In the case of intelligent agents, one popular implementation of the paradigm is the Belief-Desire-Intention architecture (Rao and Georgeff, 1991). Such agents are modelled around so-called mentalistic precepts where the agent maintains a model of its world via a belief set, a model of its objectives via a set of desires, and a model of those objectives which it can fulfil at a given moment in time via a set of intentions. A second core characteristic is an inherent social ability. Agents can communicate via an Agent Communication Language (ACL), usually one that is standardised by an appropriate organisation. This is an essential characteristic as agents rarely exist in isolation but congregate to form communities, termed Multi-Agent Systems (MAS). Though there are many other characteristics, not all of these may be supported by individual agent implementations.

2.2 Systems for deploying agents on computationally-restrained devices

A number of MAS environments have been successfully ported to mobile devices. It is not envisaged that the full development environment operates on the device, nor is it necessary for it to. In practice, developers will use desktop environments to design and develop their agent-based applications. On the physical device, it is only necessary to deploy an interpreter, or an Agent Virtual Machine (AVM). Prudent software engineers frequently test their application at various stages in the development process on the actual targeted physical devices, as there can be minor differences between the implementation of the emulator logic and that of the software on the actual device. System performance and implications for usability can also be monitored in this way. From a mobile device perspective, a number of emulators have been developed by various manufacturers and these usually seamlessly plug into standard development environments.

Two general characteristics are immediately noticeable when considering MAS environments currently available for mobile devices. Almost all are Java based; and most are extensions of well known systems that were originally developed for workstation environments. In the former case, Java has become the de facto development environment for mobile applications, and most mobile phones come with a Java interpreter pre-loaded. In the latter case, researchers are interested in gaining a deeper understanding of the various issues relating to the design, development and performance of agent-based systems on mobile devices, and the literature is sprinkled with examples of such MAS environments. 3APL-m (3APL-M, 2007) is an extension of the well-known 3APL environment (Dastani et al, 2003). LEAP (Adorni et al, 2001) has emerged from JADE (Bellifemine et al, 2007) and the open-source platform FIPA-OS (FIPA, 2007) has been extended through the development of MicroFIPA-os (Tarkoma and Laukkanen, 2002). In the case of the BDI architecture, JACK (JACK, 2007), a commercial product, and Agent Factory (Muldoon et al, 2006) have also been ported to mobile devices.

2.3 Motivational Issues

Given the computational demand of agent-based systems, a cost-benefit analysis must justify this approach. One reason researchers are interested in the area of AI and computationally restricted devices concerns the Ambient Intelligence (AmI) initiative (Aarts and Marzano, 2003). AmI is closely associated with ubiquitous computing, envisaging a world where significant computing resources are embedded in everyday objects, and where these

resources can be activated by users in an intuitive fashion. This latter issue is of fundamental importance. If we envisage a world populated with smart objects all vying for the user's attention, it would not be a pleasant place to live. Thus the goal of AmI is to embed intelligence in physical objects such that the interaction issues are addressed, and that the lifestyle of the average person is indeed improved. While a mobile phone is currently considered the effective extremity of a mobile network, this may not always be the case. Future scenarios may well include smart objects, of which a mobile phone would be just one instance, as the ultimate consumer of the resources on the mobile network. This vision is achievable, but much research needs to be undertaken in various domains before it can become a reality. At present, the mobile phone may be considered as the next frontier that must be surpassed for the deployment of intelligent solutions. So what might this intelligence be used for? And how might it positively affect the end-user? The answer is of course domain dependant. However, one common scenario can be considered, that is, the retrieval of multimedia content by mobile users. In the next section, issues relating to this objective are reflected on.

3 INTELLIGENT TECHNIQUES FOR ADAPTIVITY

With the myriad of situations in which they can find themselves, context has assumed an important role in determining the information needs of mobile users. Numerous authors have offered competing definitions of context within the ubiquitous and mobile computing literature (Brown et al, 1997; Schmidt et al, 1999); however a general, intuitive description of context has been proposed by Dey and Abowd (2000) who informally define context as follows: *“Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves”*. With the advent of mobile computing, context becomes a rich aggregation of factors. These include the subscriber's location and orientation; individual user profiles that are dynamically updated based on user activities and interactions with such data repositories as diary, travel schedules and so on; device capabilities including screen size, processor and memory restrictions, operating system; network profile including nature, bandwidth, latency, load and so on; and temporal factors including such features as time of day and day of week (Figure 2). Another related factor that often is not considered is that of the service provider's context.

The act of content delivery whilst necessarily being targeted to the perceived needs of the subscriber must also consider service providers' needs. Examples include consideration of different subscription/advertising tariffs thus necessitating preferential content placement or display prominence or display duration.

Figure 2. Some factors that determinine context for the mobile user

By providing computing devices with knowledge of the subscriber's context, this knowledge, or context-awareness, can be harvested so as to improve the interactive experience. Given the limited input capabilities of many mobile devices, the fact that context awareness can reduce the amount of explicit interaction required by an arbitrary application is a significant incentive for its use. The overall objective is, therefore, to exploit contextual information of the what, when and where of the subscriber's task, the subscriber's individual preferences and the devices capabilities, so as to simplify, augment and improve the subscriber's experience. A key challenge that is facing designers is obtaining the information needed to enable an application to function in a context-aware manner (Satyanarayanan, 2001). Very often, the desired information may constitute some of the subscriber's personal data repositories. Information from schedules, diaries, address books, and to-do lists is locally available. It is frequently necessary to sense other information pertaining to the subscriber's environment in real time. Such information may include geographic position, orientation, the identities of people nearby, locally observable objects and ongoing activities. Also, it can even include items that could indicate a subscriber's physiological and emotional state, such as body temperature, heart rate and galvanic skin response (Abowd et al, 2002). The diversity between possible contexts, in particular the varying data input and output capabilities of mobile devices, has ensured that the traditional 'one-size-fits-all' approach to content delivery is not sufficient and it is increasingly desirable to adapt the content presented to subscribers to aid them in the process of quickly identifying information of relevance. This has prompted much research into the personalisation of content to meet individual subscriber's contexts and preferences - research that increasingly incorporates a significant AI component.

3.1 Personalisation

Personalisation is an intriguing concept. The core idea underpinning it concerns the adapting of information and services to individuals' preferences. In this way, it is hoped to maximise user satisfaction, resulting in increased revenue. While both these objectives are laudable,

personalisation offers a further advantage in the mobile services sphere. It can serve as a filter through which only the most appropriate content is streamed to a mobile user. Given the inherent brevity of mobile user interactions, as well as the implicit costs associated with all mobile interactions that encompass a connection to a fixed network, this is an important consideration. Historically, developments in personalisation technologies are closely associated with the internet domain.

One prominent Web-based personalisation technology often used on e - commerce sites is collaborative filtering. The task in collaborative filtering is to predict the utility of items to a particular user (the active user) based on a database of user votes from a sample or population of other users (the user database). The typical application is the recommendation of books, music CDs, or movies. Collaborative filtering systems are often distinguished by whether they operate over implicit versus explicit votes. Explicit voting refers to a user consciously expressing his or her preference for a title, usually on a discrete numerical scale. For example, the GroupLens system of Resnick et al. (Resnik et al, 1994) uses a scale of one (bad) to five (good) for users to rate Usenet news articles, and users explicitly rate each article after reading it. Implicit voting refers to interpreting user behaviour or selections to impute a vote or preference. Implicit votes can be determined using browsing data, purchase history, web hyperlinks, or other types of information access patterns. Of the other approaches to collaborative filtering, two noteworthy alternatives are demographic filtering and content-based filtering (Montaner et al, 2003). Demographic filtering approaches use descriptions of people to learn about a relationship between a single item and the type of people that like that object. Content-based filtering approaches use descriptions of the content of the items to determine a relationship between a single user and the description of the items. Typically this is done by matching the frequency with which keywords associated with the user are present in the item's description. As with collaborative filtering, the data utilized by both these techniques can be gathered either explicitly or implicitly. These techniques should not be regarded as mutually exclusive. Indeed it is frequently the case that a combination of these three techniques is employed by each approach, enhancing the quality of their results (Huang et al, 2004).

Personalisation has also been successfully deployed within the field of mobile computing, most prominently by the mobile phone industry. With the improved data processing capabilities of today's mobile phones, it is now very common for mobile phone operators to offer content to users by means of mobile portals. Given the limited data input capabilities and the small screen size of these devices, a key problem with the usability of mobile portals relates to the amount of time that users spend navigating to content as they browse. Many mobile phone operators have embraced personalisation technology to develop intelligent portals that are automatically personalized based on subscriber usage patterns. The ClixSmart system (Smyth et al, 2005) is a good example of one such solution. It automatically adapts the navigation structure of a mobile portal for individual users by tracking the choices the user makes and proceeds to use probabilistic methods to promote and demote menu options in line with usage patterns. This reduces the amount of scrolling and clicking required by the user to reach desired content and services.

Recently, considerable research has been conducted into developing personalisation techniques for location-aware information services, focusing on how information content, and the presentation of this content, can be adapted to meet the implicit needs and preferences of the mobile user in an efficient manner, thereby reducing information overload. Availing of context sensitive information to manage a user's information needs is a non-trivial task. Deciding how to filter and prioritise a potentially vast information space according to a user's changing context requires significant capacity for analysis and decision making. One approach is to employ intelligent agents to manage user profiles, determine context, and undertake personalisation. Being capable of flexible autonomous action in dynamic, unpredictable and open environments, agents are particularly well suited to profile management in systems where frequently changing context results in user profiles that are in a constant state of flux. The ACCESS architecture (Strahan et al, 2005;Phelan et al, 2004)

seeks to provide support for multi-user environments by offering support for mobile lightweight intentional agents and intelligent prediction of user service needs, as well as personalisation of content through user profiling and context harvesting. ACCESS-compliant systems are capable of exchanging information about user interactions. This unique feature of the ACCESS architecture facilitates the development of cross-domain personalisation techniques, where services can deduce the context in which they are being invoked, allowing ACCESS-compliant services to facilitate and compliment one another.

MyCampus (Sadeh et al, 2003) is another agent-based environment for context-aware mobile services developed at Carnegie Mellon University. The system aims to aid the user with a PDA (normally a student on the university campus) in carrying out different tasks (planning events, sending messages, finding other users, and so on) by discovering and accessing Intranet and Internet services. The information from these services is filtered by use of context. The contexts that the system currently uses include the user's location on the campus, their class schedule, the location of their friends and the weather which is determined from a local weather website. Personal preferences are also used to provide the user only with information that is relevant to them. For example the user could create a context-sensitive preference that, when they are in class, promotional messages should not be delivered to the device. This is an example of rules-based personalization which makes use of embedded logical constructs: if A, then B. Rule-based personalisation engines can serve personalisation content "on the fly," cross-sell relevant products, and offer pertinent advice and/or problem resolution tips. Its most effective use is when the nature of the relationship with users is clearly defined and the procedure for adapting content clearly specified. Unfortunately, the solutions are often not very scalable because the rule sets needs to be constantly maintained, although having a well-defined process for re-setting rules can offset this. What is clear from published research is that personalisation techniques are evolving. Unfortunately, poorly implemented personalisation increases rather than decreases information overload, as it complicates the user experience and renders content difficult to locate. Today, this difficulty is further compounded by the fact that computing, as a discipline, is in the midst of a radical shift from the traditional desktop metaphor to the palmtop and mobile phone. Now, with the advent of mobile computing, there exist additional factors to consider about the user in relation to their information need such as their location, current activities, time of day etc. Therefore the ongoing challenge of personalisation is to design and implement efficient strategies that successfully overcome the added complexities of mobile computing whilst minimising information overload.

In the next section, some practical issues relating to the intelligent harnessing of context and personalisation are illustrated through a discussion on a mobile blogging application.

4 Example: The blogging application domain

Having considered the question of deploying intelligence throughout a mobile network, and illustrated the potential of intelligent techniques for content adaptivity, we will now analyze a blogging application to exercise and focus some of this discussion.

Blogging is one of the most popular activities on the WWW. Weblogs, usually abbreviated to blogs, are in essence online diaries or journals. Entries are usually textual, and arranged in reverse chronological order. The subject matter can relate to any topic. Some blogs are popular, allowing their creators to generate some additional income, usually through advertising. Others less so, and usually exist for the creators own amusement. Though text forms the core media element in most blogs, some focus on photography or images (photoblogs) while others focus on audio or podcasting.

Naturally, the mobile computing community has not been immune to these developments, and the option of maintaining one's blog from a mobile phone, or

moblogging, is becoming increasingly popular - so much so that a number of service providers are including moblogging software with their products. Indeed, the integration of digital cameras and video recorders into mobile devices seems to be paying dividends for operators as subscribers increasingly use these facilities to access dedicated photoblogging sites as well as to augment their own blogs with images.

Another blogging phenomenon of interest is called geoblogging or geotagging. The key idea here is that content is tagged with geographic coordinates, usually latitude and longitude, although place names can also be used. Thus, geobloggers can search for blog entries using coordinates as search terms. The widespread availability of cheap GPS receivers has been fundamental in enabling geoblogging. Indeed, as GPS chips are likely to be embedded in mobile phones in increasing numbers over time, one can expect geoblogging to flourish.

Recently, there has been an upsurge in interest in issues relating to the continuous archiving of personal experiences. Such issues include the capture, indexing, storage, retrieval and sharing of the constituent multimedia data. Blogging encapsulates a subset of those problems and issues necessary to making continuous life experience archiving a feasible endeavour. One exemplar project in this area is MyLifeBits (Gemmell et al, 2006). In essence, this project is concerned with enabling people to maintain a true multimedia record of everything they encounter. In addition, it is also concerned with capturing many things that the average person would be unaware of, for example, their physiological state, or salient aspects of the surrounding environment captured from embedded sensors.

At first sight, blogging might seem a solitary activity. However, experience shows that bloggers also tend to form communities and gravitate towards blogs that are of particular interest to them. For the purposes of this discussion, the needs of one group of mobile users will be considered, namely tourists.

4.1 Moblogging for Tourists

Tourists are an inherently mobile group as they usually want to see the sights and experience the culture of the area that they are visiting. They also want to record the experience, usually through photographs or video. Given the proliferation and increasing sophistication of mobile phones, new opportunities exist for disseminating content to tourists, as well as aiding them maintain rich multimedia records of their visits. In addition, they also have the opportunity of contributing to what we might term the “visit memory” of various locations by recording annotations and attaching them to the blog of the area in question. However, before examining how a service to facilitate this might be realised, it is necessary to reflect on some of the research activities that have focused on the moblogging area.

4.2 Research in Moblogging

Moblogging offers a different prospective on how blogs may be created and retrieved. Clearly, this raises a number of issues that has attracted the attention of some researchers. BRAINS (Leu et al, 2007) examines issues relating to the ubiquitous access to personal blogs on the WWW. Some of the technological issues involved in the development of moblogs are illustrated in SmartBlog (Beale, 2006; Beale, 2005). GTWeb (Spinellis, 2003) is concerned with maintenance of trip diaries, and focuses on the issue of presentation of geotemporal data. Also in the travel domain, Axup and Viller (2005) consider the use of moblogging for exchanging experiences and travel tips by backpackers in Australia. The critical issue of information overload is examined in GeoNotes (Espinoza et al, 2001), as is the use of content filtering techniques. A number of researchers have focused on the essential issue of middleware provision, for example, FrameDrops (Gross et al, 2005) and MobShare (Sarvas, 2004).

The increased interest and use of moblogging has not escaped the notice of some of the major telecommunications manufacturers and other commercial companies. ShoZu

(ShoZu, 2007) facilitates the sharing of camera phone videos and images, and their uploading to popular WWW sites. Nokia have developed Lifeblog (Nokia, 2007), an application that maintains a timeline for photos and videos taken with a phone. Researchers at Ericsson has developed weShare (Bjorling et al, 2006), an IMS based application that facilitates sharing and viewing of multimedia files, amongst others.

5. DESIGN CONSIDERATIONS FOR A MOBLOGGING APPLICATION

A number of prototypes have been described in the literature that all seek to improve the tourist experience. Most facilitate instantaneous access to content, though the methodologies used and modalities employed differ. Two of these incorporate intelligent agents. The first CRUMPET (Poslad et al, 2007), uses agents as wrappers to enable the delivery legacy services to tourists. The second, Gulliver's Genie, encompasses agents on mobile clients and network services, and uses GPRS for connectivity. It is described extensively in the literature (O'Grady and O'Hare, 2004; O'Hare & O'Grady, 2003) and has been the subject of user evaluations (O'Grady et al, 2005), the results of which are being used for a major revision that is currently ongoing. In essence, content adaptivity is a critical feature of the Genie. Multimedia content is disseminated to tourists in a just-in-time basis. All content is filtered according to the tourist's current spatial context, and according to their personal profile (age and so on) as well as cultural interests (art, literature and so on). However, the objective of providing a mechanism for enabling tourists to record, annotate and share their experiences has not been achieved as yet. However, coincidental yet timely developments in blogging and mobile phones have now made this objective achievable. Blogging offers a vehicle by which the service can be encapsulated and modelled. The current state-of-the-art in mobile devices offers the technology to make the service viable, though it is probable that it will be the next generation of the devices that will support all the necessary features. While existing blogging infrastructures offer interesting possibilities, tourists have unique needs and expectations, and significant scope exists for adapting and personalising the services to meet these expectations.

5.1 Enabling Blogging through agents

To introduce a blogging capability to the Genie, it is necessary to model the required functionality and implement it via the addition of appropriate agents.

On the mobile device, an agent, the Annotation Agent, is imbued with the necessary capability for handling the blog recording and retrieval process. In the case of adding a blog, it coordinates with two pre-existing agents, namely the Spatial Agent and the GUI Agent. User interactions are captured by the GUI agent but interpreted by the Annotation Agent. The Spatial Agent provides the current spatial context which is used for tagging, in the case of blog construction (Figure 3). In the case of blog retrieval, the spatial context is used for the defining the physical search space.

The server component of the Genie Multi-agent system resides on a network node. As agent architectures are inherently scalable, it was a case of encapsulating the blogging functionality into a separate agent – the Blog Agent. In the case of blog construction, the Tourist Agent acts as a routing agent, passing the uploaded media to the Blog Agent for parsing and storage in the multimedia database. For blog retrieval, the Blog Agent must collaborate with the GIS Agent to identify blogs nearest the tourist's current position. The Profile Agent is then consulted regarding the profile of the tourist in question such that blog list can be personalised for the tourist. As a simple example - if an audio blog is available, it is necessary to ensure that tourists can understand the language of the blog, and that their device can render the appropriate format. A list of available blogs is then dispatched to the

tourist (Figure 4).

Figure 3. Interactions between the agents on the mobile device during the blog creation process.

All agents on the server are implemented using Agent Factory (O'Hare, 1996), and on the mobile device, Agent Factory Micro Edition (Muldoon et al, 2006) has been deployed. The current implication focuses on simple geocoded text blogs (Figure 4), but implementing full multimedia enriched blogging is currently ongoing. This has not been practical up until now, as it only now at the time of writing that devices capable of supporting multimedia enriched blogging are appearing on the market. For example, camera phones have been on the market for some time but the programming interfaces necessary to control them are only now being made available to third-party developers.

In its present incarnation, most of the intelligence in the system resides with the agents on the server. However, by deploying agents on the client, the opportunity exists for augmenting the intelligence there as the situation warrants. Furthermore, the use of agents on the client offers a particularly attractive and intuitive option for modelling the various components of the system and the interactions between them.

Figure 4. Interactions between the agents on the server to identity and filter blog entries for the user to browse and select.

Figure 5. Recording and retrieving a blog entry

5 Discussion & Future Work

Ongoing developments in mobile hardware and software have enabled the deployment of intelligent applications and services at the extremities of mobile telecommunications networks, namely on the subscriber's or end-user's device. How such intelligence is harnessed is at the discretion of the service provider. As has been described, one useful application concerns the determination of the prevailing context, or strictly speaking, selective aspects of the prevailing contextual situation. In simple cases, this endeavour can be completed on the user's device. However, more sophisticated context streams may need to be

interpreted on a fixed server node. The distributed and collaborative aspects of the agent paradigm offer one elegant solution for modelling and implementing the constituent processes necessary to delivering an adaptive and personalised solution.

Moving multimedia data over the network has a cost element; hence the eagerness of network operators to include blogging software on their devices. This raises particular challenges. Users will not be eager to use services that cost more than a nominal amount. While tourists may be more flexible, assuming they feel the services are useful, monetary costs will still affect how frequently they use the service. As tourists mostly operate in urban areas, there is an opportunity to avail of the increasing prevalence of WLAN hotspots, most of which are free. As consumers demand devices that will also work with WLAN, as well as with traditional voice and data networks, an opportunity exists to intelligently switch between technologies as the occasion demands. Should a tourist wander into a hotspot, their blogs can be transparently unloaded for free. Likewise, by maintaining an accurate model of the environment and of the tourist's spatial context, data can be intelligent precached on their devices, in anticipation of short-term future access.

Finally, given the importance of the tourist spatial context in the scenario outlined, ongoing developments in global positioning systems offer an interesting opportunity for enhancing the accuracy of the determined position. Differential GPS (DGPS) techniques, traditionally the preserve of surveying companies as well as being prohibitively expensive, can now be accessed for free via the internet. SISNet (Chen et al, 2003), an initiative of the European Space Agency (ESA), is already in operation and disseminates DGPS corrections to interested parties. In practice, this means calculating positions to within 5 metres, as distinct from the normal 20 meters with conventional GPS. Utilising this technology may offer a significant opportunity for unobtrusively bringing the tourist's attention to blogs and other items of interest that may be in their immediate vicinity.

6 Conclusion

Deploying intelligent techniques on all the major components of a mobile network is technically feasible. However, recent developments in mobile device technologies have enhanced the options available to service providers. Hence, intelligent technologies can be deployed at the extreme periphery of the network, and in the most important place, the subscriber's handset. Given the computational limitations of the average mobile device, and the inherent complexity and variety of environment in which the average mobile subscriber operates, intelligent techniques on mobile devices offer software engineers a new and promising approach for consideration during design and implementation. This will have significant implications for HCI and usability aspects of the service at hand, and may be manifested in various ways, Intelligent User Interfaces being one obvious example.

In this paper, key concepts pertaining to the intelligent network as understood in the traditional telecommunications sense has been examined. Recent developments concerning the deployment of AI techniques on mobile devices thus enabling the deployment of intelligent services on such devices has been described. In this way, it can be seen how critical elements that were once the preserve of disparate disciplines are now merging. As an illustration of the mutual benefit that accrues to each discipline, and ultimately to the end user, the techniques of context-awareness and personalisation were described and illustrated through a discussion on content adaptivity. Finally, to demonstrate the viability of deploying AI techniques on mobile devices, a design and implementation of a moblogging service was described.

In conclusion: the dawn of a new technological era where intelligent technologies become embedded in everyday objects is imminent. As ambient intelligent environments increase in number and sophistication (and initially, specialisation), a significant opportunity will be afforded to customise and personalise services for their inhabitants. Though the

outstanding challenges must not be underestimated, it may be reasonably hoped that the harnessing of intelligent techniques may help launch a new wave of adaptive and personalised products and services.

References

- Aarts, E, Marzano, S. (2003), *The New Everyday: Views on Ambient Intelligence*, Rotterdam:010 Publishers.
- Abowd, G. D., Dey, A. K., Brown, P. J., Davies, N., Smith, M., Steggles, P. (1999) Towards a Better Understanding of Context and Context-Awareness in Proceedings of the 1st international Symposium on Handheld and Ubiquitous Computing (Karlsruhe, Germany, 1999); Berlin: Springer-Verlag. Lecture Notes in Computer Science (LNCS) 1707:304-307.
- Abowd, G.D., Mynatt, E.D., Rodden, T. (2002) The Human Experience, *IEEE Pervasive Computing*, 1(1):48-57.
- Adorni, G., Bergenti, F., Poggi, A. & Rimassa, G. (2001) Enabling FIPA Agents on Small Devices. In: Klusch, M., Zambonelli, F. (eds): *Cooperative Information Agents V*. Berlin:Springer-Verlag. Lecture Notes in Computer Science (LNCS) 2182: 248-257.
- Axup, J., Viller, S. (2005) Augmenting travel gossip: design for mobile communities, in Proceedings of the 19th Conference of the Computer-Human interaction Special interest Group (Chisig) of Australia on Computer-Human interaction: Citizens online: Considerations For Today and the Future (Canberra, Australia, 2005); ACM International Conference Proceeding Series, 122:1-4.
- Barfield, W. and Caudell, T. (2001) *Fundamentals of Wearable Computers and Augmented Reality*, Lawrence Erlbaum Associates: New Jersey.
- Beale, R. (2005) Supporting Social Interaction with Smart Phones, *IEEE Pervasive Computing* 4 (2):35-41.
- Beale, R. (2006) Mobile Blogging: Experiences of Technologically Inspired Design, In *CHI '06 Extended Abstracts on Human Factors in Computing Systems* (Montréal, Québec, Canada, 2006); New York: ACM Press. 225-230.
- Bellifemine, F., Caire, G., Greenwood, D.: (2007) *Developing Multi-Agent Systems with JADE*. New Jersey:John Wiley & Sons.
- Bjorling, M., Carlsten, J., Kessler, P., Kruse, E., Stille, M. (2006) Sharing everyday experiences, *Ericsson Review*, 1:10-19.
- Brown, P.J., Bovey, J.D., Chen X. (1997) Context-Aware Applications: From the Laboratory to the Marketplace. *IEEE Personal Communications*, 4(5):58-64.
- Chen, R., Toran-Marti F., Ventura-Traveset, J. Access to the EGNOS signal in space over mobile-IP, *GPS Solutions* 7(1):16-22.
- Dastani, M., van Riemsdijk, B., Dignum, F., Meyer, J.J. (2003) A Programming Language for Cognitive Agents: Goal Directed 3APL. Proceedings of the First Workshop on Programming Multiagent Systems: Languages, frameworks, techniques, and tools, (Melbourne, 2003); Berlin: Springer Verlag. Lecture Notes in Computer Science 3067:111-130.
- Espinoza, F., Persson, P., Sandin, A., Nystrom, H., Cacciatore, E., Bylund, M. (2001) GeoNotes: Social and Navigational Aspects of Location-Based Information Systems, in Proceedings of the 3rd international conference on Ubiquitous Computing, (Atlanta, Georgia, USA, 2001); Berlin: Springer-Verlag. Lecture Notes in Computer Science (LNCS) 2201:2-17.
- Foundation for Intelligent Physical Agents (FIPA), <http://www.fipa.org>. (accessed 27th June 2007).
- Gemmell, J., Bell, G., and Lueder, R. (2006) MyLifeBits: a personal database for everything. *Communications of the ACM* 49(1):89-95.
- Gross, T., Kleppe, M. (2005) FrameDrops: A Mobile VideoBlog for Workgroups and Virtual Communities, in Proceedings of the International ACM SIGGROUP conference on Supporting Group Work, (Sanibel Island, Florida, USA, 2005); New York:ACM Press. 128-131.
- Huang, Z., Zeng, D., Chen, H. (2004) A unified recommendation framework based on Probabilistic Relational Models [WWW document] http://papers.ssrn.com/sol3/papers.cfm?abstract_id=9065 13 (accessed 27th June 2007). International Telecommunications Union (1997) Q.1200 - General Series Intelligent Network Recommendation Structure, ITU: Geneva.
- JACK - The Agent Oriented Software Group, <http://www.agent-software.com>. (accessed 27th June 2007).
- Leu, J-S, Chi, Y-P., Shih, W-K. (2007) Design and implementation of Blog Rendering and Accessing INstantly system (BRAINS), *Journal of Network and Computer Applications*, 30(1):296-307.
- Montaner, M., Lopez, B., Rosa, J.L. (2003) A taxonomy of recommender agents on the internet. *Artificial Intelligence Review*, 19:285-330.
- Muldoon, C., O Hare, G.M.P., Collier, R.W., O Grady, M.J. (2006) Agent Factory Micro Edition: A Framework for Ambient Applications, in Proceedings of the International Conference on Computational Science (ICCS2006), (Reading, UK, 2006); Berlin: Springer-Verlag. Lecture Notes in Computer Science (LNCS) 3993 :727-734.

- Nokia Lifeblog – <http://www.nokia.com/lifeblog>. [WWW document] (accessed 27th June 2007). O'Grady, M. J., O'Hare, G. M. P. (2004) Just-in-Time Multimedia Distribution in a Mobile Computing Environment, *IEEE Multimedia*, 1 1(4):62-74.
- O'Grady, M.J., O'Hare, G.M.P., Sas, C. (2005) Mobile agents for mobile tourists: a user evaluation of Gulliver's Genie, *Interacting with Computers* 1 7(4):342-366.
- O'Hare, G.M.P. (1996) Agent Factory: An Environment for the Fabrication of Distributed Artificial Systems, In: *Foundations of Distributed Artificial Intelligence* (Eds. O'Hare & Jennings), Wiley Publishers.
- O'Hare, G. M. P., O'Grady, M. J. (2003) Gulliver's Genie: A Multi-Agent System for Ubiquitous and Intelligent Content Delivery, *Computer Communications*, 26 (11): 1177-1187.
- Phelan, D., Strahan R., Collier, R., Muldoon, C., O'Hare, G.M.P. (2004) Sos: Accomodation on the fly with ACCESS, in *Proceedings of the 13th International Florida Artificial Intelligence Research Symposium Conference (FLAIRS 2004)*, (Miami Beach, Florida, 2004); California: AAAI Press. 2-7
- Poslad, S., Laamanen H., Malaka R., Nick A., Zipf, A. (2001) Crumppet: Creation of user-friendly mobile services personalised for tourism, in *Proceeding of the Second IEE International Conference on 3G Mobile Communication Technologies* (London, UK, 2001); London:IEEE. 28-32.
- Rao, A. and Georgeff, M. (1991) Modelling Rational Agents within a BDI Architecture, in *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning* (San Mateo, CA. 1991.); Morgan Kaufmann Publishers. 473-484.
- Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J. (1994) GroupLens: an open architecture for collaborative filtering of netnews, in *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, North Carolina, United States, 1994); New York:ACM Press. 175-186.
- Sadeh N., Chan T., Linh,V., Kwon O. Takizawa K. (2003) Creating an Open Agent Environment for Context-aware M-Commerce, [WWW document] <http://www.cs.cmu.edu/~sadeh/Publications/Small%20Selection/Agent%20Cities%20Book%20Chapter%20final.pdf> (accessed 4th April 2007).
- Sarvas, R., Viikari, M., Pesonen, J., Nevanlinna, H. (2004) MobShare: controlled and immediate sharing of mobile images, in *Proceedings of the 12th annual ACM international conference on Multimedia* (New York, 2004); New York:ACM Press. 724-731.
- Satyanarayanan, M. (2001) Pervasive Computing: Vision and Challenges, *IEEE Personal Communications*, 8(4): 10-17.
- Schmidt A., Aidoo K.A., Takaluoma A., Tuomela U., Van Laerhoven K, Van de Velde W. (1999) Advanced interaction in context, in *Proceedings of First International Symposium on Handheld and Ubiquitous Computing (HUC'99)* (Karlsruhe, Germany (1999); Berlin: Springer-Verlag. *Lecture Notes in Computer Science (LNCS)* 1707:89-101.
- ShoZu. <http://www.shozu.com>. [WWW document] (accessed 27th June 2007).
- Smyth, B., McCarthy, K., Reilly, J. (2005) Mobile Portal Personalisation: Tools and Techniques, in B. Mobasher and S. S. Anand (eds.) *Intelligent Techniques in Web Personalisation*, Berlin: Springer-Verlag. *Lecture Notes in Computer Science (LNCS)* 3169: 255-271.
- Spinellis, D. D. (2003) Position-Annotated Photographs: A Geotemporal Web, *Pervasive Computing*, 2(2):72-79.
- Strahan R., O'Hare, G.M.P., Phelan, D., Muldoon, C., Collier, R., (2005) ACCESS: An Agent based Architecture for the Rapid Prototyping of Location Aware Services, in *Proceedings of the 5th International Conference on Computational Science (ICCS 2005)*, (Atlanta, USA, 2005); Berlin: Springer-Verlag. *Lecture Notes in Computer Science (LNCS)* 3516:695-702.
- Tarkoma, S., Laukkanen, M. (2002) Supporting Software Agents on Small Devices, in *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2002)* (Bologna, Italy, 2002); New York: ACM Press. 565-566.
- Vasilakos, A. and Pedrycz, W. (2006) *Ambient Intelligence, Wireless Networking, Ubiquitous Computing*, Artech House: Norwood.
- 3APL-M: Platform for Lightweight Deliberative Agents [WWW document] <http://www.cs.uu.nl/3apl-m/> (accessed 27th June 2007).