



Gulliver's Genie: Agency, Mobility, Adaptivity

M. J. O'Grady*, G. M. P. O'Hare

PRISM (Practice and Research in Intelligent Systems & Multimedia) Laboratory, Department of Computer Science, University College Dublin, Belfield, Dublin 4, Ireland.

Abstract

The challenge presented by ambient intelligence and pervasive computing is an environment within which computing artifacts become seamlessly merged into our surroundings, whereby interaction with such artifacts becomes intuitive and unobtrusive, where the devices become sensitive to the presence of people and are imbued with the ability to anticipate and service the individual needs of each and every user. In this paper, we present Gulliver's Genie, an archetypical ubiquitous computing application. Gulliver's Genie embraces three central constructs; those of agency, mobility and adaptivity. The architecture adopts a Multi-Agent Systems (MAS) metaphor whereby agents manage and maintain a context within which mobile users exist and, based upon this context, seeks to adapt and personalize content based upon perceived individual user needs. Agents are mobile and may migrate to or from environmental artifacts reflecting their computational constraints.

Keywords: Intelligent agents, pervasive computing, context-awareness, BDI agents, mobile augmented reality

1. Introduction

Developing applications that operate reliably and robustly is a difficult task at the best of times. This difficulty becomes more acute when the application must be deployed in an environment that may be considered hostile, unpredictable and highly dynamic. The everyday outdoor setting in which many ambient and pervasive computing applications are expected to operate is one such example environment. The challenge presented by ambient intelligence and pervasive computing is to deliver an environment within which computing artifacts become seamlessly merged, whereby interaction with such artifacts becomes intuitive and unobtrusive, where the devices become sensitive to the presence of people within their environment [1] and imbued with the ability to anticipate and service the individual needs of each and every user.

In this paper, we present Gulliver's Genie, an archetypical ubiquitous computing application. Gulliver's Genie embraces three central themes in its construction; those of agency, mobility and adaptivity. The architecture adopts a Multi-Agent Systems (MAS) metaphor whereby agents manage and maintain a context within which mobile users exist and based upon this context seeks to adapt and personalize content based upon perceived individual user needs. Agents are mobile and may migrate to or from environmental artifacts reflecting their computational constraints. We describe the agent architecture and animate it via some examples. We provide brief comment on comprehensive user evaluations conducted. Based upon these, we outline some ongoing and future work in the area of Augmented Reality (AR) for mobile service delivery.

*Corresponding author.

E-mail addresses: michael.j.ograd@ucd.ie (M.J. O'Grady), Gregory.OHare@ucd.ie (G.M.P. O'Hare).

2. Intelligent Agents

Intelligent agents have been the focus of much research and debate within the academic community over the last decade. At the very least, agents bring an extra dimension to software development approaches though it could be reasonably argued that they offer an alternative and radical development paradigm. While the term *agent* has been in use since the 1970s at least, it has had the misfortune to be used quite casually and an agreed rigorous definition has been conspicuously lacking. It is therefore instructive to review the implications of ascribing the term *agenthood* to a piece of software.

A number of incarnations of agents exist. At the low end of the scale are so called reactive agents that function in a simple stimulus-response manner and minimise the need for a model of their environment. Such agents are generally termed *weak agents* [2]. At the high end of the scale are sophisticated intentional agents. Such agents maintain complex models of their environment and a list of certain objectives that they seek to fulfil as well as detailed plans on how to achieve their objectives. Such agents are frequently termed *strong agents* and support rational reasoning in a collaborative context, usually realised as a Multi-Agent System (MAS). A comprehensive introduction to MAS precepts may be found elsewhere [3].

A popular manifestation of strong agents has been realised through the Belief-Desire-Intention (BDI) paradigm [4] (Fig. 1), an intuitive and computationally tractable delivery of the fundamental intentional agent stance. Beliefs represent a method by which the agent can model its environment. Desires represent those goals that the agent wishes to achieve, that is, its *raison d'être*. This construct challenges the traditional view of software, which is task-oriented and comprises of a program performing tasks without knowing why or what it is trying to achieve. In contrast, agents are goal-oriented and will only execute a task if it aids in the achievement of its goals. The notion of intention was formulated to encapsulate the process of specifying those goals that can be realistically achieved through the execution of the appropriate task(s). Once this has been achieved, the agent proceeds to review its belief set and the entire cycle commences anew. In the following sections, the term agent is assumed to imply a strong intentional agent of the BDI genre.

Though agents have been generally the preserve of the academic community up to now, a number of toolkits are available including FIPA-OS [5], LEAP [6], Zeus [7], JACK [8] and AgentBuilder [9]. The latter two being commercial products. More recently, the specification of a methodology for developing agent-oriented software has been receiving considerable attention. This area, generally termed Agent Oriented Software Engineering (AOSE) [10], is fundamental to the incorporation of agent technologies into conventional software engineering practice. A significant step towards achieving this is currently being

undertaken through the Agent Unified Modelling Language (AUML) initiative [11], a proposal to incorporate agent-related constructs into the ubiquitous UML.

So when should agents be considered as a basis for a software solution? Some view agents as a possible remedy to the complexity inherent in modern computing systems. At an application level, any system that can be logically represented as consisting of a number of dynamic interacting components is perceived as being potentially suitable. In particular, domains that are extremely dynamic and complex are suitable and it is here that the benefits of

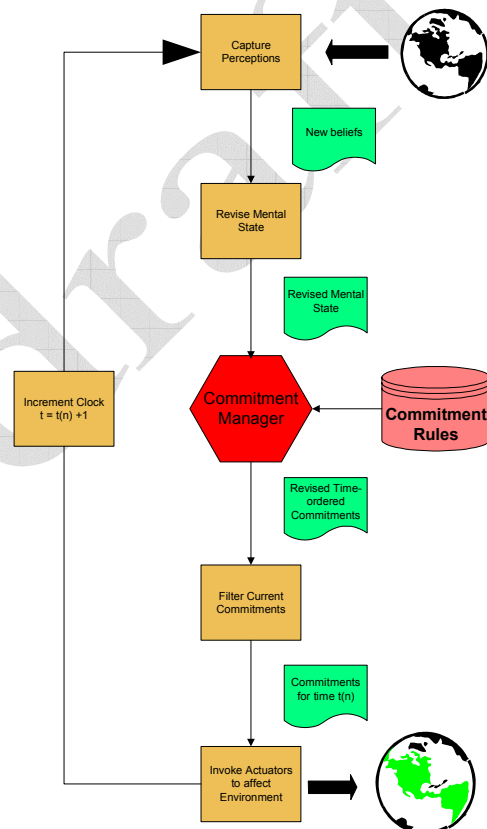


Fig. 1. BDI Perceive-Deliberate-Act Cycle.

agents in mobile computing applications may be harvested fruitfully. Traditionally, mobile computing would not have been considered suitable for the deployment of agents, primarily due to the computational limitations of Personal Digital Assistants (PDAs) and other handheld devices. However, recent developments both in hardware and software are rendering such limitations obsolete.

3. Gulliver's Genie

At its simplest, Gulliver's Genie [12] may be regarded as a mobile context-aware tourist guide. In many respects,

it is a classic example of a mobile application or service that many people would anticipate being available in the coming years. In this, it is not unique, as the literature is sprinkled with prototypical systems for various application domains. In the case of the tourist domain, there are a number of well documented systems including GUIDE [13] and CRUMPET [14]. As well as the mobile computing domain, researchers in wearable computing and Augmented Reality (AR) have also found the tourism domain a fertile one for experimentation. Well known examples include MARS [15] and ARREAL [16].

3.1. Some Design Considerations

When designing the Genie, a number of sometimes conflicting factors needed to be reconciled. In the first instance, the primary objectives for the Genie need to be elucidated clearly. In the case of a tourist, what would the Genie do that would contribute to a more satisfactory experience? From a software engineering perspective, what principles should be adopted when designing and implementing the Genie? And how would these affect and improve the resultant system?

Though tourists are by nature a heterogeneous group, they do have a number of needs in common. One of these is an interest or curiosity, exhibited to greater or lesser degree, in the history and culture of the area they are visiting. Therefore, the primary service offered by the Genie would be the delivery of multimedia presentations on the various attractions encountered during exploration. Other potential services, for example, standard location-aware services, could be incorporated at a future date.

By judiciously harvesting a number of techniques from differing disciplines, the tourist's experience could be significantly enhanced. In particular, significant opportunity existed for tailoring the content to the needs of individual tourists. Essentially, the Genie should be an *adaptive* application. Any information presented should be adapted to the tourist's immediate situation or context: at its simplest, position and orientation. However, precepts from user modeling could be gainfully employed for further refinement. For example, is the tourist an adult or a teenager? Do they have particular interests, for example art, architecture? Is there information that would be of particular interest to a particular nationality? The nature of the tourist's interaction need also be considered. Is a traditional explicit interaction modality preferable? Or can any implicit interaction [17] be captured and interpreted?

It is now instructive to review some of the practicalities of availing of those services offered by the Genie. It is assumed that the tourist has access to a device of the PDA genre that is capable of displaying various multimedia formats including sound, color imagery and, ideally, video. Access to the tourist's position or spatial context is also essential. Currently, using a satellite navigation service is

the only option, and, in particular, GPS is the only viable, commonly-available technology. In the future of course, technologies based on Galileo [18] as well as techniques based on the topologies of wireless cellular telecommunications networks [19] may well offer realistic alternatives.

Gulliver's Genie, like all such applications, necessitates a rich multimedia content set. The second issue concerns policies for the creation, access and maintenance of this multimedia repository. Three approaches may be considered:

1. All multimedia content may be hosted on the PDA – assuming that adequate memory is available, possibly through the use of a Compact Flash (CF) memory card. Though simplifying the design and implementation, this would severely curtail the ability of the Genie to adapt any content presented to the tourist. In particular the potential for using sophisticated user modeling techniques would be non-existent. A further issue that may well affect the end-user experience would be the risk of the data becoming redundant or out-of-date. The accuracy of all presented data is essential. If, perhaps, a museum should change its opening hours, it would prove well-nigh impossible to inform all Genie users of this alteration.
2. All multimedia may be hosted on a central server. This resolves the issues listed previously thus ensuring that the Genie may exhibit a more adaptive nature. However, it introduces further complications in that the tourist must be equipped with a wireless modem in order to access the central multimedia repository.
3. A hybrid approach may be considered in which static multimedia elements are hosted on the PDA and dynamic elements stored on a central server. This is an attractive approach but the issue of repository maintenance still lurks in the background thus compromising the potential of the Genie to adapt and expand as the number of users grows.

After considering each of these options, it was decided that the Genie would adopt the second option therefore demanding that the tourist be equipped with a wireless modem. Ensuring the relevancy of the delivered content, a factor essential to the tourist's satisfaction, would now be guaranteed. Selecting a wireless technology was relatively straightforward. An approach based on WLAN, such as that used by GUIDE, necessitated the deployment of a dedicated infrastructure thereby increasing complexity and restricting deployment. In the CRUMPET project, the 2.5G General Packet Radio Service (GPRS) was utilized successfully. GPRS has also been deployed in large swaths of the world. This, it was felt, maximized the deployment potential of the Genie and, as such, was selected as the technology of choice. One critical limitation of GPRS was its poor data rates averaging about 30 kb/s; a problem that would accentuate the difficulty of the timely delivery of multimedia data.

In reality, the data rates supported by GPRS is a function of the number of subscribers actively using the network at a given instance, as well as being subject to the policies implemented by the local network operator. This means that predicting the available data rate is practically impossible. In our experience, it averages at about 30 kb/s.

proposed here, termed intelligent precaching [21], assumes the availability of a model of the environment in which the user operates. Given a certain amount of knowledge about the user, for example their likely activities, in this case sightseeing, it is possible to anticipate with a reasonable degree of certainty what their information requirements will be in the short to medium term. Therefore content can be downloaded to their devices in a just-in-time basis thus disguising network limitations from the user. In addition, by considering the tourist's context and user model, significant scope exists for restricting the amount of information sent over the network without compromising the essential quality of that information. In the following section, an architecture is proposed for meeting the original objectives set out for the Genie while simultaneously allowing for the restrictions inherent in operating in an outdoor environment.

3.2. *Why Agents?*

Mobile computing is still its infancy. Though it is likely to become the dominant computing paradigm in the coming years, researchers are not near reaching a consensus about various issues critical to its success. Not least among these is the issue of how best to engineer applications for mobile users. In light of this, an alternative and radical approach was chosen for the Genie, namely that of one based on intelligent agents.

As has been seen, intelligent agents are perceived as being particularly suitable to situations that do not lend themselves readily to a conventional solution. An example of that is one that is extremely dynamic, possibly in a spatial or temporal sense, or just inherently complex. In the case of the Genie, the tourist is almost continuously mobile; the environment in which they are situated changes in a corresponding manner. In addition, rich multimedia content must be disseminated to such users over an unpredictable network and displayed on a device of limited computational capacity. All of which suggests that agents may well prove a suitable paradigm around which the Genie could be designed and implemented.

3.3. *Agent Abstraction & Identification*

Identifying and assigning tasks to individual agents is the single most important task facing the designer of an agent-based system. Though a number of methodologies have been proposed for the design process, for example, GAIA [22] and Prometheus [23], nevertheless, this task

To address this problem, a solution based on precaching was conceived. Traditionally, precaching has been used to overcome both network disconnection and latency, and, given the ongoing interest in mobile computing, has been the subject of much research activity [20]. The solution

remains more an art than a science. An initial temptation is to identify all the principle components, in much the same way as in Object Oriented Programming (OOP), and assign agent attributes to them. Conceptually, and from a visual perspective, this approach may appear attractive. However, assigning trivial tasks to agents is best avoided, particularly in the case of BDI agents where the reasoning capability is bought at considerable computational cost, especially on a PDA. A rule that was employed during the design of the Genie was that, in addition to considering the spatial and temporal dynamics associated with a task, a case must demonstrably exist for using some, though not all, of the standard agent characteristics, for example, autonomy, mobility, proactivity and so on. In addition, the potential of a task for incorporating additional intelligence in the future was allowed for thus increasing the longevity of the design. Traditionally, the Genie would have been implemented using a standard client-server architecture. Superficially, it may still be regarded in this light. However, in practice, it is a Multi-Agent System (MAS) encompassing a suite of agents hosted both on the tourist's device and supporting server.

One of the most pressing concerns facing the agent research community and prospective agent application developers is what methodology to follow. As has been stated, a number have been proposed in the literature. However, one initiative that is receiving considerable attention is that of Agent UML (AUML). One of the primary objectives of this initiative is to develop a suite of diagrams that captures the essence of agenthood. These diagrams would be consistent with existing UML notation. One of the obvious benefits of such an approach is that it allows agent proponents to position agents as enhancements of existing de facto technologies and standards. In this way, it is hoped that resistance to the agent paradigm would be minimized in industry thus increasing the deployment of agent-based solutions.

AUML formed the basis of the design of Gulliver's Genie. Though it is not a de jure standard as yet, a number of proposals have appeared in the literature and the design has been based on one of these [24]. One example of this is the AUML protocol diagram, similar to the standard UML sequence diagram, which contains some additional notation for handling agent constructs.

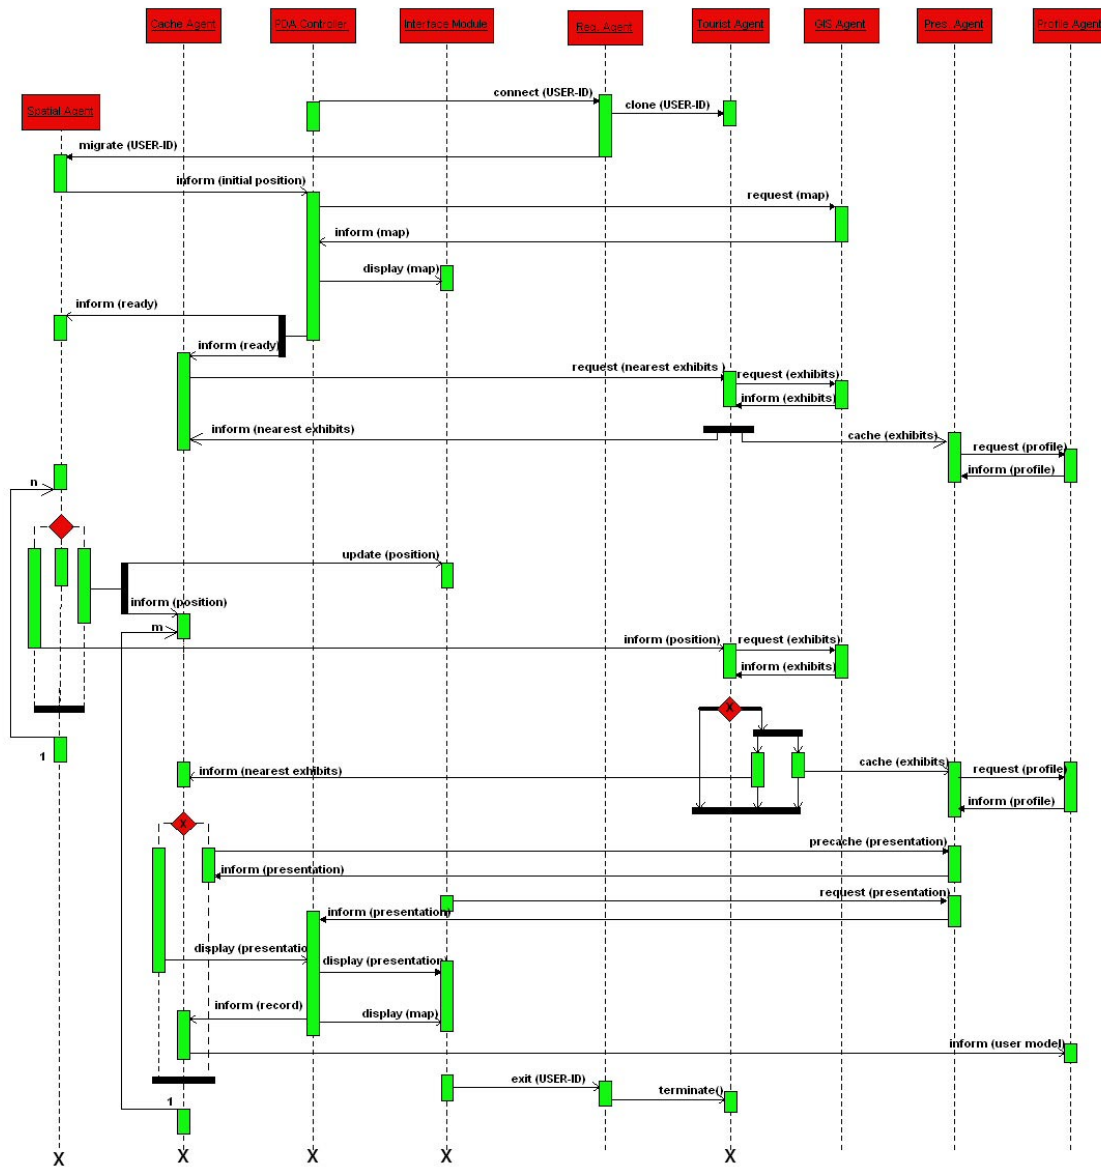
4. Agent Implementation

In the following sections, the tasks associated with each of the agents comprising the Genie MAS are described. An AUML Protocol diagram detailing some of the principal interactions between these agents (and some non-agent components) may be seen in Fig. 2. Furthermore, an outline

of the constituent agents of the Genie system may be seen in Fig. 3. These agents collectively deliver the Genie functionality, through opportunistic collaborations.

Fig. 2. An AUML protocol diagram illustrating some of the principal between the agents in the Genie MAS.

Final draft



4.1. Spatial Agent.

Deducing the tourist's spatial context, that is their position and orientation, is the primary task of the Spatial Agent. To do this, it monitors the GPS device autonomously and reacts to any perceived changes in the user's position or orientation. It must first verify the integrity of the signal used for calculating the current position after which it verifies that the reading is consistent with previous readings. Once satisfied that the reading is valid, the agent must then decide on the necessity and desirability of updating the tourist's navigational display accordingly. Depending on the scale of the map, it may not be

appropriate to update the display. Finally, it provides spatial data to any other interested agents, usually the Cache Agent and the Tourist Agent.

At first sight, one might assume that the Spatial Agent would be essentially reactive. This is not the case. The GPS receiver is continuously broadcasting data with a position reading available every 4 seconds. Keeping continuous track of this datafeed is needlessly expensive

Fig. 3. An illustration of the Genie MAS.

computationally as the tourist is not likely to have made significant movement in such a time interval. Rather the



sensor is polled at appropriate intervals, on average every 10 seconds. Thus the Spatial Agent is essentially proactive.

The Spatial Agent is mobile and migrates to the tourist's device on commencement of the service. In this sense it is representative of the dynamic nature of our agent community. A general principle adopted within the design

of Gulliver's Genie is to ensure a thin client. In certain circumstances it may prove prudent to migrate agents to or from the client in order to ensure load balancing and or agent self-preservation. Depleting PDA battery power for example may result in simple self-management *autonomic behaviour* [25] that ensures agent migration to an artifact of increased power resources and thus enhanced agent longevity.

Though programmed to interpret the NMEA protocol that is standard amongst GPS receivers, it may well be that a future device will use cellular network techniques for determining position. Should this become a reality, the Spatial Agent must be augmented with an additional capacity for handling this. By using a mobile agent, the core application on the PDA need not be updated: merely the agent itself. In this way, the future operation of this particular aspect of the Genie is protected.

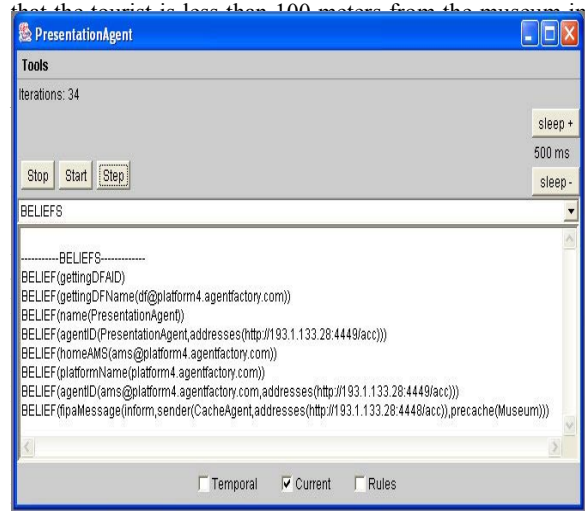
4.2. Cache Agent.

Intelligent precaching of multimedia data is a strategy adopted by the Genie to handle the poor data rates supported by cellular networks. An additional benefit accruing from this approach is that the limited memory resources available on a PDA can be managed more effectively. Fundamental to the operation of the Cache Agent is the availability of a model of the tourist's environment. This model could contain, for example, details of restaurants, bus terminals or shops within their vicinity. At present, it is limited to those tourist attractions within their immediate surroundings. In practice, an

environmental model of a city is assembled and stored in a database. As it is large, it cannot be maintained on the PDA, therefore partial models are used. The GIS Agent assembles these models dynamically in response to requests from the Tourist Agent. Note that these requests are triggered by position updates from the Spatial Agent. The Tourist Agent, aware of the state of the Cache Agent, may dispatch a new partial model, or not, as the occasion demands.

In essence, the Cache Agent is continuously reconciling the position of the tourist with the position of the various attractions. By continuously doing this, it can make an informed guess at where the user is likely to go and what attraction they are likely to encounter with a reasonably high degree of certainty. It then arranges to download the relevant multimedia files from the server. When the user finally encounters the attraction, the presentation ought to be available and can start automatically. Should the agent have deduced incorrectly and the tourist not acted as anticipated, the contents of the cache can simply be discarded. As collaboration between agents is essential to this process of intelligent precaching, it is instructive to review the process in some detail.

Consider the case of a tourist exploring a particular city. The GIS Agent determines that the museum, church and art gallery are the nearest exhibits. The Cache Agent is briefed of this situation, courtesy of the Tourist Agent. As the tourist explores, the Cache Agent keeps track of their position relative to the nearest exhibits. At some stage, the tourist appears to be converging on the museum. The Cache Agent, sensing this, decides to precache a presentation now that the tourist is less than 100 meters from the museum in



in anticipation of a request. On subsequently being informed that the Cache Agent wishes to precache a presentation for the museum (Fig. 5), it dispatches the appropriate presentation to the Cache Agent.

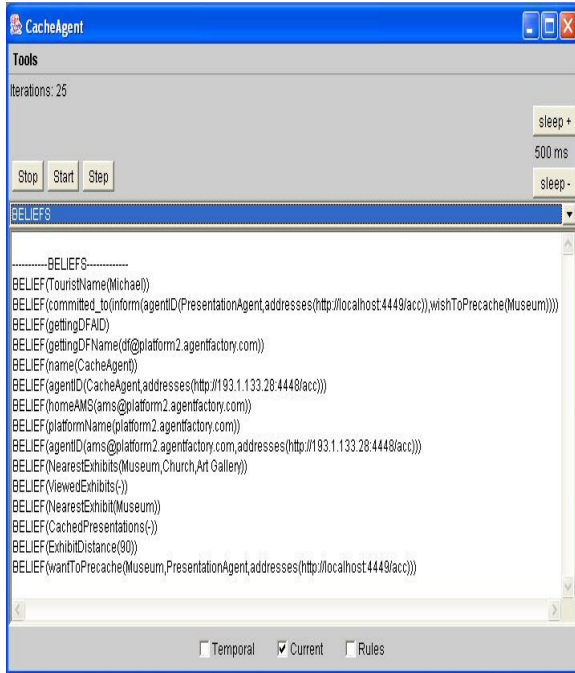


Fig. 5. The Presentation Agent is notified that the Cache Agent wants to precache a presentation.

Having received the presentation, the Cache Agent updates the cache and its belief set accordingly (Fig. 6). Knowing that it has a presentation for the museum in its cache, it proceeds to monitors the tourist's movements as before.

As the tourist approaches the museum, they will eventually pass an activation proximity threshold that triggers the following commitment rule:

$$\begin{aligned} &BELIEF(NearestExhibit(?exhibit)) \& \\ &BELIEF(ExhibitDistance(?delta)) \& \\ &BELIEF(CachedPresentations(?exhibit)) \Rightarrow \\ &COMMIT(Self, Now, BELIEF(true), \\ &DisplayPres(?exhibit)); \end{aligned}$$

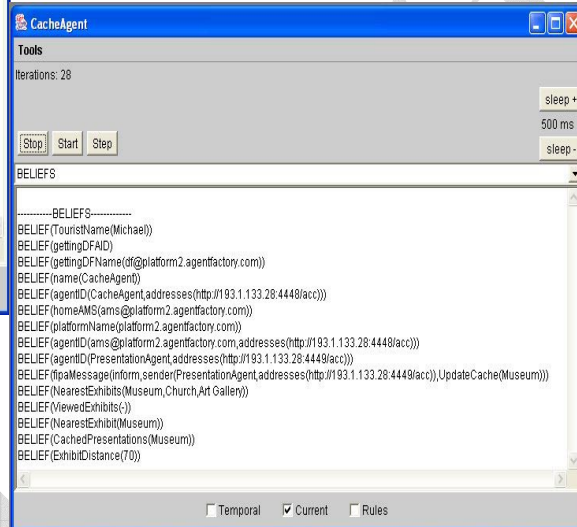
Fig. 6. The Cache Agent is asked by the Presentation Agent to update its cache with a presentation for the museum.

Fig. 7. The Cache Agent mental state that triggers a presentation.

The mental state that triggers this rule is illustrated in Fig. 7, and, as can be seen from Fig. 8, this commitment is now triggered.

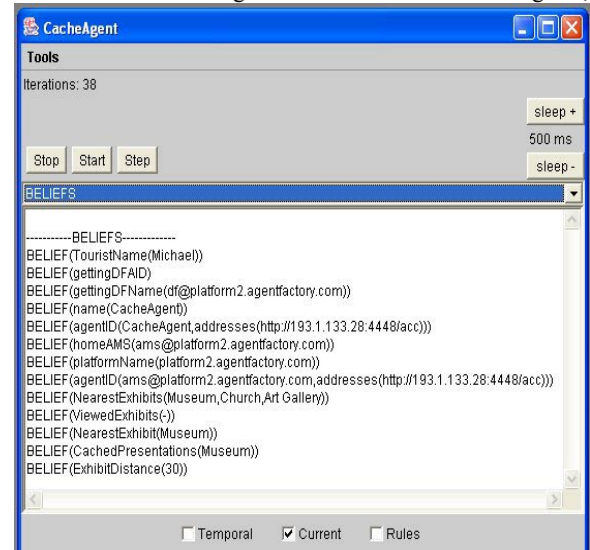
After a presentation has been shown, a record of the tourist's interaction is returned to the Profile Agent for further analysis.

Fig. 8. The Cache Agent is now committed to showing a presentation for the museum.



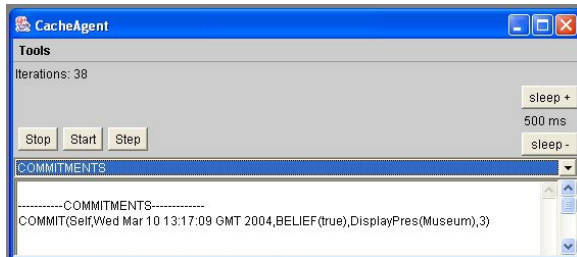
4.3. Tourist Agent

All users are assigned their own individual agents,



called Tourist Agents, on commencing a session with the Genie. This is achieved via an agent cloning capability contained within Agent Factory. Each user registration

results in a cloning of the generic Tourist Agent and the subsequent specialisation of that agent to monitor the specific user in question. The primary benefit obtained from this approach is that it ensures that the Genie can easily cater for a large number of users thus ensuring its scalability. Trials conducted have evaluated system



performance with some 1000 system clones. However these are mere simulations and actual evaluations with 1000 simultaneous users have not been possible for all the obvious reasons, including cost and the manageability of the experiment. The Tourist Agent continuously monitors the tourist's behaviour, courtesy of the Spatial Agent, and seeks to anticipate future services requests. To provide services, it collaborates with the GIS Agent to ensure that the model of the environment currently being monitored by the Cache Agent is up-to-date. It arranges with the Presentation Agent to maintain a number of presentations that correspond to each exhibit that the Cache Agent is monitoring. In this way, no time is lost when a request for a presentation is received and a presentation can be dispatched immediately.

4.4. GIS Agent

Given the importance that an accurate model of the tourist's environment is to the Genie, the GIS Agent has the responsibility of assembling the partial models that the Cache Agent uses for close monitoring of the tourist's behaviour. In addition, it advises Tourist Agents on whether updates to the environmental model currently being monitored by the Cache Agent on the PDA are desirable.

4.5. Profile Agent

User models are inherently dynamic constructs and the Genie continuously refines individual user models in the light of ongoing tourist behaviour. This is a critical task as it directly affects the end-user experience. In addition, as the user model is a primary filter for all content distributed to tourists, it facilitates content focus thus minimising irrelevant dispatch of information to the tourist; an important consideration in the light of GPRS limitations.

4.6. Presentation Agent

Assembling presentations is a collaborative process undertaken in conjunction with the Profile Agent. The Presentation Agent maintains a list of presentations for each exhibit in the vicinity of each user. This list is continuously updated in light of refinements to user models and, in particular, user movement, as the environmental model monitored by the Cache Agent changes. However a rich corpus of factors is considered in the assembly of the personalized presentations. These include such factors as user preferences and interests as contained within their user profile, line-of-sight upon approaching an attraction, tourist device limitations (screen real estate, memory, processor speed monochrome/colour) and network characteristics, for example data rates achieved using GPRS, are factored into the agent's calculations. Obviously, a large multimedia presentation of a few megabytes cannot be downloaded quickly. Therefore the agent prioritizes the individual presentation entities and removes some of the less important, depending on presentation size. The use of M-PEG 4 [26] as a technology for segmenting media content has been considered. It may prove beneficial to deliver and display text, audio and images upon the mobile device as a holding exercise while more demanding media rich (video and 3D graphic imagery) content is being transferred prior to display.

4.7. Registration Agent

The Registration Agent facilitates the initial registration process. When a tourist wishes to use the Genie service, it first arranges for the migration of the Spatial Agent. It must then assign an agent to the individual tourist. To do this, it clones a Tourist Agent and initialises it with some essential parameters including the tourist's identification and assigned port. For the remainder of the session, that Tourist Agent is essentially the tourist's interface to the services provided by the Genie. At the end of a session, it terminates the cloned Tourist Agent thus making more resources available for other agents in the Genie community.

4.8. Non-Agent Components

Having identified and justified the existence of the various agents, the outstanding components can now be described.

4.8.1. Interface Module

This component manages the user interface on the PDA. It shows presentations as appropriate and manages all explicit interaction that the tourist may have with the Genie.

4.8.2. PDA Controller

Somewhat of a misnomer, the PDA Controller is responsible for initializing the Genie, monitoring the Interface Module and organizing the orderly termination of the Genie service. In reality, it generally acts under the direction of both the Spatial Agent and the Cache Agent.

4.8.3. Database

Fundamental to the operation of the Genie, the database stores various types of data including:

- **User Profiles:** Standard data concerning the tourist is stored including name, age, native language and so on. In addition, an Interest Model is maintained detailing the tourist's cultural interests.
- **Multimedia data:** All tourist attractions within a geographic area covered by the Genie must have a number of associated media files with them. These include images, sound snippets and video files. In addition, these files are tagged with various attributes such that the Presentation Agent can reason about their inclusion in a presentation for a tourist with a particular model.
- **Environmental Model:** A model of the environment containing data pertaining to the various tourist attractions is essential to the Genie. Other information that may be stored includes any information necessary to the delivery of standard-location-aware services, for example, hotels.
- **Map Repository:** Navigation support is an essential prerequisite of a mobile tourist information system and access to a range of maps is of critical importance.
- **Presentation Cache:** This temporary cache contains pre-assembled presentations that a Cache Agent may need at any time. The cache is continually updated as the environmental model used by the Cache Agent changes.

4.8.4. Management Toolkit

A sophisticated toolkit is essential for database management. Systems like the Genie are inextricably dependent upon the relevance, completeness, accuracy and timeliness of the content contained within the multimedia database. It is also important to the portability of the Genie as it enables the rapid deployment of the Genie in new areas. A number of tools are provided which facilitate the rapid population of the multimedia database. These provide a simple and intuitive interface thus enabling content providers to upload content and associate it with a given location by merely using a point and click map-based interface.

Specifically a facility for geocoding maps is essential. Tourist attractions must be located on the map and

principal approaches identified, thus facilitating the view or line-of-sight that the tourist will have. In this way the environmental model is specified. Associating multimedia content with the various attractions is essential. In particular, it must be possible for the designer to assign various tags to individual entities thus allowing the presentation to reason about them when dynamically assembling the presentation. Finally, a tool is available for registering tourists with the Genie and specifying the various personal and cultural-interest parameters necessary for presentation personalization. Detailed consideration of this toolset will be reported elsewhere.

5. Realisation of Gulliver's Genie

Gulliver's Genie has been implemented on a HP IPAQ, a device of the ilk that tourists may be expected to possess in the future. This unit is expandable via a dual-slot expansion sleeve. This can contain two slots for PCMCIA or CF cards. A Holux CF GPS card (GM-270) is used for position determination. The GlobeTrotter GPRS wireless Internet card from Option International enables data communication with the Genie server. DB2 is used for hosting the Genie database. The standard WWW communications protocol HTTP is run over the GPRS connection. A servlet engine (Apache Tomcat) provides the primary interface with the Genie server. An illustration of the deployment of the various components comprising the Genie may be seen in Fig. 9. All components have been developed in Java. In the case of the agents however, the Agent Factory toolkit [27] was used for design and implementation. Agent Factory is a specialized toolkit for the fabrication of BDI agents. It supports the design, implementation and deployment phases of the standard software engineering lifecycle. Facilities for enabling agent mobility and cloning are supported. The reasoning engine itself is compliant with PersonalJava thus ensuring its operation on lightweight devices such as PDAs.

5.1. User Evaluations

In an effort to verify our design and obtain some independent feedback, user evaluations were conducted on the Genie (Fig. 10). A detailed description of the methodology followed and the results is beyond the scope of this paper but may be found elsewhere [28].

However, it is appropriate to reflect briefly on the more important findings. Initial reaction was quite favorable.

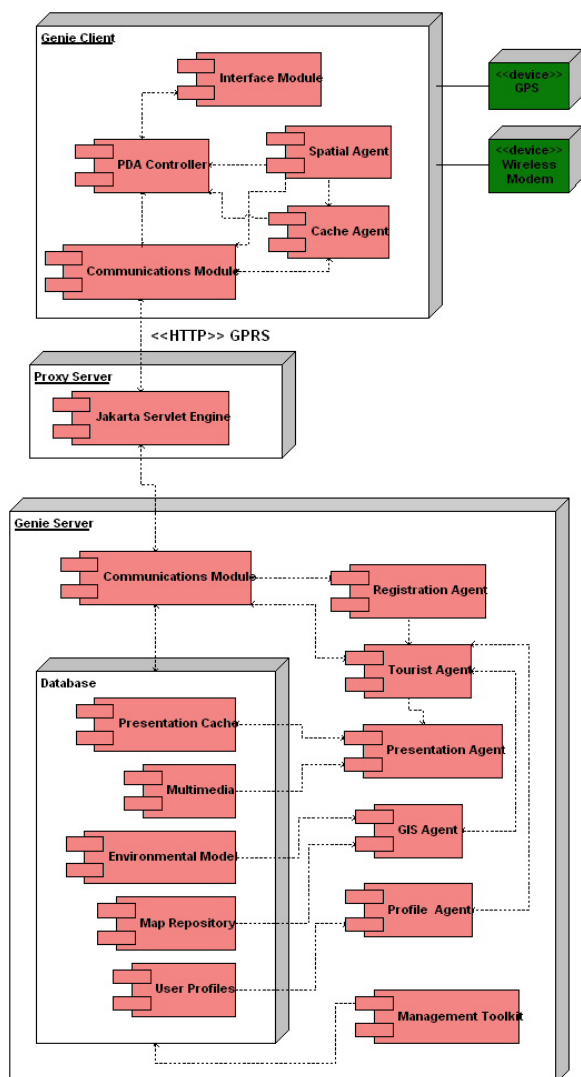


Fig. 9. A Component & Deployment Diagram of Gulliver's Genie.

Over 85% of participants (40 people took part in the evaluation), expressed overall satisfaction with the Genie. Over 84% considered the cultural presentations adequate. When asked how the Genie could be improved, over 62% said its general responsiveness could be improved, albeit in various ways, for example responding to user interaction as well as updating the display faster. Interestingly, a number of participants expressed unease with the autonomous nature of the Genie and felt that they had been moved out of the control loop, as it were.

Finally, a number of suggestions were made for improving the Genie. These included offering a greater range of services as well as supporting more sophisticated user interfaces. In particular, the use of 3D graphical constructs was recommended.



Fig. 10. A visitor consulting the Genie

6. Future Challenges

The genesis for Gulliver's Genie was that of the HIPS project [29] [30]. HIPS sought to achieve the simultaneous navigation and exploration of both the physical space that the user moved through together with the associated digital information space associated with this area. The user trials conducted with the Genie demonstrated that this simultaneous navigation was difficult with users finding it hard to fuse seamlessly the information content and the physical world.

In order to address this, recent work with the Genie has investigated the integration of these two spaces through the use of Augmented Reality (AR). The deployment of AR within the mobile computing sphere is in of itself not novel. Several systems exist which offer valuable beginnings and insights in this regard. Notable amongst these are systems like ARCHEOGUIDE [31] which offers augmented overlays of archeological sites for tourists, GEIST [32] which provides historical overlays, and Touring Machine

[33] which facilitates the exploration of urban scenes using a rich set of GPS and other sensors.



Fig. 11. The MicroOptical micro-HeadUp Display.

The augmented reality extension to Gulliver's Genie employs the MicroOptical micro-HeadUp Display (HUD). This is integrated with the Genie via a VGA adaptor. Fig. 11 depicts the wearable nature of the display while Fig. 12 illustrates a Genie presentation contained upon the micro display. Clearly the use of this micro display placed in front of the line of sight of one eye addresses the issue of the *disconnect* experienced by users of the PDA based Genie display.

However the re-engineering of the interface to avail of and, suitably harness the possibilities afforded by such technologies, is non-trivial. Gulliver's Genie constitutes the first agent-based mobile AR information delivery system offering context-sensitive and personalized content. While other systems like Touring Machine, MARS, GEIST support context-sensitivity and ARCHEOGUIDE supports personalized tours, the Genie architecture offers a dynamic, evolving and inherently intelligent system.

In moving to accommodate AR, several key architectural components needed to be revised. The Presentation Agent needed to be aware when content was to be displayed on a micro-display. The de facto presentation mode of follow-up links would be replaced by augmented labels overlaid upon the objects in view. While this is not difficult to achieve at a high level of granularity it would demand complex localization technologies to both locate the user and HUD sensors to record the angle of view. In addition the multimedia database needs to associate these labels with particular locations and thus multimedia content. The Profile Agent needs to record and monitor a richer and differing set of user interaction modalities including head movements and hand held interaction device. While this work is on-going the agent-based Genie architecture has elegantly supported the AR extension. Current work investigates the mechanisms for delivering scene annotation akin to *AR Graffiti* where

users can annotate the scene leaving comments that may be based upon user profiles may be deemed beneficial and consequently displayed.



Fig. 12. What the tourist sees via the HeadUp Display.

7. Conclusion

This paper has presented Gulliver's Genie, a ubiquitous computing application based around the three key enablers of agency, mobility and adaptivity. The Genie offers a mobile context-sensitive tourist service, which manages a rich context and utilises this to empower the personalization of content reflective of the perceived needs of the individual user.

Responsiveness, adaptivity, ease-of-use and robustness are all expectations of the mobile user. These are delivered through the deployment of a community of BDI agents that perceive the environment and maintain a corresponding mental state. This mental model drives the deliberative behavior of the system and enables personalization adaptivity and intelligent precaching. The system agents are agile in nature and may migrate from environmental artifact

to artifact. In the main, such migration acts merely as an enabler for system load balancing and thus enhancing system responsiveness. However one could envisage migrations between a rich set of situated environmental devices rather than the client-server-client migrations considered within this paper.

Recent work fuelled by feedback from user trials has begun to harness Augmented Reality as a conduit for addressing the discontinuity experienced by users between the physical and information space. The Genie thus delivers a mobile augmented reality experience to the user again mediated by the same set of agents albeit they need to be tuned to accommodate the differing interaction modalities.

Acknowledgements

Michael O'Grady gratefully acknowledges the support of the Irish Research Council for Science, Engineering & Technology (IRCSET) through the Embark Initiative postdoctoral fellowship programme. Gregory O'Hare gratefully acknowledges the support of Science Foundation Ireland under Grant No. 03/IN.3/1361.

References

- [1] Aarts, E. Ambient Intelligence: A Multimedia Perspective, *IEEE Multimedia*, 11 (1), Jan-Mar 2004, pp. 12-19.
- [2] Wooldridge, M., Jennings, N. Intelligent Agents: Theory and Practice, *Knowledge Engineering Review*, 10 (2), 1995, pp.115-152.
- [3] Wooldridge, M. An Introduction to MultiAgent Systems, John Wiley and Sons Ltd, February 2002.
- [4] Kinny, M., Georgeff, M., Rao, A., A Methodology and Modelling Technique for Systems of BDI Agents, In: Proc. 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96), Springer Verlag Lecture Notes in Artificial Intelligence, vol. 1038, 1996, pp. 56-71.
- [5] Foundation for Intelligent Physical Agents, <http://www.fipa.org>
- [6] Bergenti, F., Poggi, A. LEAP: A FIPA Platform for Handheld and Mobile Devices, In: Proceedings of the 8th International Workshop on Agent Theories, Architectures and Languages (ATAL-2001), Seattle, WA, USA, August 2001.
- [7] Nwana, H. S., Ndumu, D. T., Lee, L. C., Collis, J. C., ZEUS: A Tool-Kit for Building Distributed Multi-Agent Systems, *Applied Artificial Intelligence Journal*, 13 (1), 1999, pp. 129-186.
- [8] JACK, <http://www.agent-software.com>
- [9] AgentBuilder, <http://www.agentbuilder.com>
- [10] Wooldridge, M., Weiß, G., Ciancarini, P. Agent-Oriented Software Engineering II, Second International Workshop, AOSE 2001, Montreal, Canada, May 29, 2001.
- [11] Agent UML, <http://www.auml.org>
- [12] O'Hare, G.M.P., O'Grady, M.J., Gulliver's Genie: A Multi-Agent System for Ubiquitous and Intelligent Content Delivery, *Computer Communications*, 26 (11), 2003, pp. 1177-1187.
- [13] Cheverst, K., Mitchell, K. Davies, N., The Role of Adaptive Hypermedia in a Context-Aware Tourist Guide, *Communications of the ACM*. 45 (5), 2002, pp. 47-51.
- [14] Poslad, S., Laamanen H., Malaka R., Nick A., Zipf, A., CRUMPET: Creation of User-friendly Mobile Services Personalised for Tourism, In: Second International Conference on 3G Mobile Communication Technologies, London, UK, March 2001.
- [15] Hollerer, T., Feiner, S., Terauchi, T., Rashid, G., Hallaway, D., Exploring MARS: Developing Indoor and Outdoor User Interfaces to a Mobile Augmented Reality System, *Computers & Graphics*, 23 (6), 1999, pp. 779-785.
- [16] Baus, J., Krueger, A., Wahlster, W., A Resource-adaptive Mobile Navigation System", In: Proceedings of the International Conference on Intelligent User Interfaces (IUI'02), San Francisco, California, January 2002.
- [17] Schmidt, A. Implicit Human Computer interaction through Context, *Personal Technologies*, 4 (2&3), June 2000.
- [18] Galileo homepage - <http://www.esa.int/esaNA/galileo.html>
- [19] Zhao, Y. Mobile Phone Location Determination and Its Impact on Intelligent Transportation Systems, *IEEE Transactions on Intelligent Transportation Systems*, 1 (1), 2000.
- [20] Jiang, Z., Kleinrock, L., Web Prefetching in a Mobile Environment, *IEEE Personal Communications*, 5 (5), October 1998, pp. 25-34.
- [21] O'Grady, M. J., O'Hare, G. M. P., 2004. Just-in-Time Multimedia Distribution in a Mobile Computing Environment, *IEEE Multimedia* (In press), 2004.
- [22] Wooldridge, M., Jennings, N. R., Kinny, D., The Gaia Methodology for Agent-Oriented Analysis and Design, *Journal of Autonomous Agents and Multi-Agent Systems*, 3 (3), pp. 285- 312, 2000.
- [23] Padgham, L., Winikoff, M., Prometheus: A Methodology for Developing Intelligent Agents, In: Proceedings of the 3rd International Workshop on Agent-Oriented Software Engineering (AOSE-2002), Bologna, Italy, July 2002.

- [24] Bauer, B. UML Class Diagrams Revisited in the Context of Agent-Based Systems, In: Proceedings of Agent-Oriented Software Engineering, Montreal, Canada, 2001.
- [25] Kephart, J.O., Chess, D.M. The Vision of Autonomic Computing, IEEE Computer, 36(1), January 2003, pp. 41-50.
- [26] Periera, F., Ebrahimi, T. The MPEG-4 Book, Prentice Hall, ISBN 0-13-061621-4, 2002.
- [27] Collier, R.W., O'Hare, G.M.P., Lowen, T.D. Rooney, C.F.B., Beyond Prototyping in the Factory of Agents, In: Proceedings of the 3rd International Central & Eastern European Conference on Multi-Agent Systems (CEEMAS 2003), Prague, Czech Republic, 16th-18th June, 2003.
- [28] O'Grady, M. J., O'Hare, G.M.P., Sas, C. Mobile Agents for Mobile Agents: A User Evaluation, Gulliver's Genie, submitted to Interacting with Computers, 2004.
- [29] Benelli, G., Bianchi, A., Marti, P., Not, E., Sennati, D. HIPS Hyper-Interaction within Physical Space, IEEE International Conference on Multimedia Computing and Systems (ICMCS99). Florence, Italy, 1999.
- [30] O'Grady, M.J. O'Hare, G.M.P. Accessing Cultural Tourist Information via a Context-Sensitive Tourist Guide, Information Technology & Tourism, 5 (1), 2002.
- [31] Vlahakis, V., Ioannidis, N., Karigiannis, J. ARCHEOGUIDE: Challenges and Solutions of a Personalised Augmented Reality Guide for Archeological Sites, Computer Graphics in Art, History and Archeology, Special Issue of IEEE Computer Graphics and Applications magazine, 22 (5), Sept-Oct. 2002, pp. 52-60.
- [32] Kretschmer, U., Coors, V., Spierling, U., Grasbon, D., Schneider, K., Rojas, I., Malaka, R., Meeting the Spirit of History, In: Proceedings of the Conference on Virtual reality, Archeology and Cultural Heritage, Glyfada, Greece, November, 2001.
- [33] Feiner, S., MacIntyre, B., Hollerer, T., Webster, T. A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment, In: Proc. ISWC '97 (Int. Symp. On Wearable Computing), Cambridge, MA, USA, 13th-14th October, 1997, p74-81.