

Just-In-Time Multimedia Distribution in a Mobile Computing Environment

Michael J. O'Grady and Gregory M. P. O'Hare
University College, Dublin

Poor data rates often hinder the timely dissemination of multimedia content to users in a mobile computing environment. We describe an approach, based on dynamic and proactive precaching, to overcome these limitations. This approach, which we term intelligent precaching, is realized through the innovative deployment of intelligent agents on mobile devices.

Mobile computing is widely anticipated to become the next major computer usage model in coming years (see the sidebar, "Mobile Computing," for more details). Such a paradigm shift will have profound implications for practically all aspects of the computer industry. How hardware and software will evolve to address this new scenario, however, remains to be seen. Among the issues that researchers must address is how to effectively distribute information to users—information that may well have a substantial multimedia component.

The physical environment is by its very nature heterogeneous. Likewise, the electronic infrastructure available to mobile users can differ considerably even within small geographic regions. Catering to the complexity and differences inherent in such environments is a critical task facing those seeking to deploy applications and services to mobile users. One particular example that illustrates this heterogeneity is the varying quality and range of data communications services currently available to mobile users. Enabling the

timely delivery of multimedia information to mobile users under such constraints is the primary focus of our research, which we describe here.

Wireless data bottleneck

Users are accustomed to high-specification hardware and high-speed networks, so it can be something of a culture shock when they first experience mobile computing. The first obvious difference is the inferior computational power and facilities that state-of-the-art PDAs offer compared to conventional workstations (see the "Mobile Cellular Telecommunications" sidebar, p. 64). A second noticeable difference is the poor data rates supported by wireless networks, compared to those of their fixed-network cousins. These poor data rates are painfully obvious to users doing something as mundane as browsing the World Wide Web.

Consequently, remedying the data rate problem continues to preoccupy the telecommunications sector. Solutions based on satellite and microwave technologies have had limited success. However, solutions based on either WiFi (IEEE 802.11) or cellular telecommunications (any 3G solution ratified by the International Telecommunication Union [ITU]) offer the best hope for a solution in the coming years. Although frequently presented as competing technologies, WiFi and 3G complement each other in certain areas. WiFi offers higher data rates than 3G but is only deployed in select localities. In contrast, cellular telecommunications networks are deployed over vast regions of the world, although fully 3G-compliant networks are operational only in Japan for the most part. Most network operators, however, are in the process of migrating to 3G (albeit incrementally), so data rates are likewise increasing. Similarly, deployment of WiFi is accelerating, thus increasing its geographical availability. But 3G networks will never offer comparable data rates to WiFi (see Table 1). Likewise, WiFi will never cover as wide a geographic area as cellular networks. These circumstances give those aspiring to provide services to mobile users ample food for thought.

Intelligent precaching: A possible solution

As we researched the practicalities of realizing applications and services for mobile users, we realized that delivering dynamic information services would prove problematic against the backdrop of constraints under which such users work. We

Mobile Computing

Mobile computing has evolved rapidly over the past decade, aided by the explosive growth in wireless telecommunications. A popular vision of mobile computing involves a small handheld device of the PDA genre, connected to a fixed network via a wireless connection. Although this is currently the predominant vision, researchers have proposed several alternatives.

Location-aware computing

By incorporating a location-sensing component—for example, GPS—into a PDA and augmenting it with a wireless network, location-aware computing becomes a possibility. Users receive information that is relevant only to their location. A classic example is that of an electronic yellow pages that sorts its entries according to their distance from the user. Indeed, location-aware computing has grasped the imagination of many in the business community, and industry analysts expect the market for such services will grow substantially in the coming years. However, how best to realize location-aware computing remains the focus of much attention.¹

Ubiquitous computing

The ubiquitous computing concept was first articulated by the late Mark Weiser just over a decade ago.² This concept envisages an environment saturated with sensors and other computing devices that the user can interact with in a graceful and

seamless manner. This vision was radical and one that cannot as yet be entirely implemented. The term pervasive computing³ is frequently used interchangeably with *ubiquitous computing*. In each case, the core concept is simple: The user can access the necessary computer resources anywhere and at anytime.

Wearable computing

Wearable computing⁴ involves users actually wearing a computer. Philosophically, the concept is diametrically opposed to ubiquitous computing in that it envisages users having all the necessary computational resources with them.

References

1. C.A. Patterson, R.R. Muntz, and C.M. Pancake, "Challenges in Location-Aware Computing," *IEEE Pervasive Computing*, vol. 2, no. 2, April–June 2003, pp. 80–89.
2. M. Weiser, "The Computer of the 21st Century," *Scientific Am.*, vol. 265, no. 3, Sept. 1991.
3. M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Comm.*, vol. 8, no. 4, Aug. 2001, pp. 10–17.
4. S. Mann, "Wearable Computing as Means of Personal Empowerment," keynote address for the First International Conference on Wearable Computing (ICWC98), 1998. Available at <http://wearcomp.org/icwc/keynote.htm>.

Table 1. Data rates that users can expect using either cellular networks or WiFi.

Family	Technology	Generation	Maximum Data Rate	Expected Data Rate*	Expected Download Time**
Global System for Mobile (GSM) Communications	General Packet Radio Service (GPRS)	2.5G	115 kilobits per second (Kbps)	30 Kbps	4.4 minutes
	Enhanced Data Rates for GSM Evolution (EDGE)	2.5G	384 Kbps	60 Kbps	2.2 minutes
	Universal Mobile Telecommunications Service (UMTS)	3G	2 megabits per second (Mbps)	300 Kbps	26 seconds
Code Division Multiple Access (CDMA)	CDMA	2.5G	153 Kbps	60 Kbps	2.2 minutes
	CDMA	3G	2 Mbps	300 Kbps	26 seconds
Wireless Local Area Network (WLAN)	WiFi	—	11 Mbps	1 Mbps***	8 seconds

* Approximate data rate, because the obtained data rate depends on the prevailing network operating conditions.

** Assume 1 Mbyte file. Note that the average MP3 file would be 3 Mbytes.

*** Data rates supported by WLAN can vary substantially according to the distance from the access point.

focused on the tourism domain, investigating how best to disseminate information to tourists as they

explored a city—an endeavor of other research groups in various disciplines, as well. Although a

Mobile Cellular Telecommunications

Despite being one of the success stories of the past decade, cellular networks were severely constrained—and to a greater degree, remain so—in their ability to handle wireless data. In an effort to address this, the International Telecommunication Union launched its IMT2000 initiative,¹ which defined the specifications for so-called Third Generation (3G) networks. Among other requirements, the ITU ordained that 3G-compliant networks should support data rates of

- 144 kilobits per second (Kbps) for users moving quickly, such as in vehicles;
- 84 Kbps for pedestrians; and
- 2 megabits per second (Mbps) for users in a low-mobility or office environment.

If we consider the issue of multimedia dissemination in such an environment, the situation appears quite promising, at least initially. Closer examination reveals several problems. First, the data rates that a user can expect will vary according to the user's context. Obviously, the difference between 144 Kbps and 2 Mbps will have a significant impact on the end-user experience and must be considered when network operators plan services for mobile users. This issue is further complicated when we consider normal operating conditions and network operator policies. The spectrum allocated to operators is usually divided into channels, which are in turn shared by multiple subscribers. Some channels are reserved for voice calls, because voice traffic still makes a major contribution to an operator's revenue. All these factors can considerably reduce the data rates available to subscribers. Thus services that seemed attainable—

for example, videoconferencing and online games with significant streaming media content—may not be feasible except in limited circumstances. This situation has contributed immensely to the current turbulence in the cellular telecommunications sectors.

Most networks are currently in the process of being migrated to 3G via a series of interim technologies that are frequently termed 2.5G. Although data rates continue to improve, they may not be adequate in meeting users' expectations. In the longer term, the issue may well be addressed through 3.5G technologies, such as High Speed Downlink Packet Access (HSDPA),² or possibly the proposed 4G networks already on the drawing board.³ In the meantime, operators continue migrating their networks toward 3G, modestly improving the data rates available to customers as they do so. However, the effective delivery of services with a rich multimedia content is likely to remain, for several years, a significant obstacle to prospective application developers and services providers.

References

1. International Telecommunication Union – Radiocommunications (ITU-R), "International Mobile Telecommunications 2000 (IMT2000)," ITU-R Recommendation M.687-2, 1997. Available for purchase at <http://www.itu.int>.
2. European Telecommunications Standards Institute (ETSI), "Universal Mobile Telecommunications System (UMTS); UTRA High Speed Downlink Packet Access (HSDPA); Overall Description, Stage 2," ETSI TS 125 308, Version 6.0.0, release 6, 2003. Freely available at <http://www.etsi.org>.
3. T. Zahariadis and D. Kazakos, "(R)evolution toward 4G Mobile Communication Systems," *IEEE Wireless Comm.*, vol. 10, no. 4, Aug. 2003, pp. 6-7.

number of systems have been documented in the literature,^{1,2} CyberGuide³ and Guide⁴ are the two pioneering examples. Other notable efforts include Archeoguide,⁵ which is oriented toward historical reconstructions, and the hypermedia guide developed for the Costa Aquarium in Genoa, Italy.⁶

In our research, we envisaged the tourist being equipped with a PDA or similar device. Because navigation support would be essential, we assumed the use of GPS. The delivery of personalized multimedia-rich presentations on attractions that a tourist might encounter offered us significant opportunities to enhance the individual tourist experience. Delivering such presentations in an appropriate and timely manner would prove challenging, particularly when we considered the well-documented computational restraints of PDAs: for example, poor computa-

tional power and limited memory. One obvious solution was to host the entire multimedia repository on the PDA, using some combination of smart card or memory stick technologies, for instance. Although feasible, such an approach would severely curtail the possibilities to personalize content and include inherently dynamic content. In the latter case, such technologies could potentially lead to circumstances where content isn't revised or maintained systematically and diligently, resulting in the decay of content accuracy and, ultimately, a frustrated tourist.

The other alternative we considered was to host the content on a centralized server and access it via a wireless data network. But which wireless technology should we use? Although Guide successfully used WiFi, the WiFi technology would severely curtail the geographical area where the

tourist could operate. Crumpey,⁷ a project that provided location-aware services to tourists, had successfully used the 2.5G General Packet Radio Service (GPRS), which yields a data rate of about 30 kilobits per second (Kbps). However, this project wasn't overtly concerned with services that incorporated a significant multimedia component. Nevertheless, GPRS covers a large geographic area, and we therefore identified it as the most appropriate technology for our purposes.

Distributing information with a significant multimedia component over a GPRS network is problematic even under good network operating conditions.^{8,9} To counteract this drawback, we developed the following two-step strategy:

1. *Intelligently and dynamically precache information on the tourist's PDA.* Implementing such a strategy requires a model of the tourist's environment. Using this model in conjunction with the tourist's known position and cultural interests, it's possible to derive a ranked list of potential attractions that the tourist might visit. On identifying an attraction that the tourist is likely to visit, a presentation can be requested and downloaded in a just-in-time basis. The idea being that when the tourist finally encounters the attraction, a presentation is available on the PDA for instant consultation. As the tourist continues to explore, the cycle must be repeated, thus requiring the process itself to be dynamic and proactive. This approach reduces network latency, which gives the illusion of instantaneous information downloading and accessibility. However, immediate information access is sacrificed at the expense of the download of extraneous material that may be discarded later due to incorrect assumptions about the user's anticipated behavior.
2. *Minimize the amount of information sent to the tourist.* To achieve this, we dispatch only information relevant to the tourist's position. Second, all information sent is consistent with tourists' cultural interests, thus improving the probability that they'll be satisfied with the presented information while further minimizing the amount of information that must be dispatched.

Precaching (sometimes called prefetching) of data is a well-known, established technique in mobile computing for handling network disconnections.^{10,11} Traditionally, network servers might

have precached files on a user's laptop using criteria such as a "most recently used" basis. As most laptops can currently be augmented with a wireless modem, the issue of disconnection doesn't present the problem it once did. However, the recent surge of interest in location-aware services has again focused attention on the core issue of cache management. The invalidation and replacement strategies that servers should adopt when working with location-dependent data in mobile computing environments remain the subject of much research.^{12,13}

Although precaching has been used effectively to disguise network bandwidth limitations from users, it can also aid in obtaining maximum benefit from PDAs' limited memory, thus minimizing two inherent constraints of modern mobile computing systems. Indeed, these two constraints are most acute when users work with multimedia data, primarily due to the large size of the files involved. If a model of the tourist's immediate environment is available, along with a profile of the tourist's interests and preferences, these can provide a basis for the system to reason about what the cache contents should be. Furthermore, by making the precaching mechanism itself dynamic, the system can maintain the PDA's cache contents in such a way that it is always relevant to the tourist's position. In this way, information can be made available in an almost on demand fashion to mobile users.

System developers must consider several essential factors if they are to successfully implement such a precaching strategy. First, the entire process must be transparent to the tourist. This suggests a preference for an autonomous solution. Second, the cache contents must be maintained so that they're relevant both to the tourist's present position, and ideally, where the tourist will go next.

In practice, the system must react to the tourist's behavior as well as proactively seek to second-guess the tourist's likely future behavior and update the cache contents accordingly. In anticipating the tourist's possible future behavior, developers must reconcile a number of factors. These include the features of the tourist's immediate environment, the tourist's own interests and preferences, as well as the memory constraints of the PDA device. To achieve this, it's essential that developers provide a reasoning capability: Autonomy, reactivity, proactivity, and strong reasoning capabilities are synonymous with intelligent agents. In addition, agents' inherent social ability facilitates the intuitive

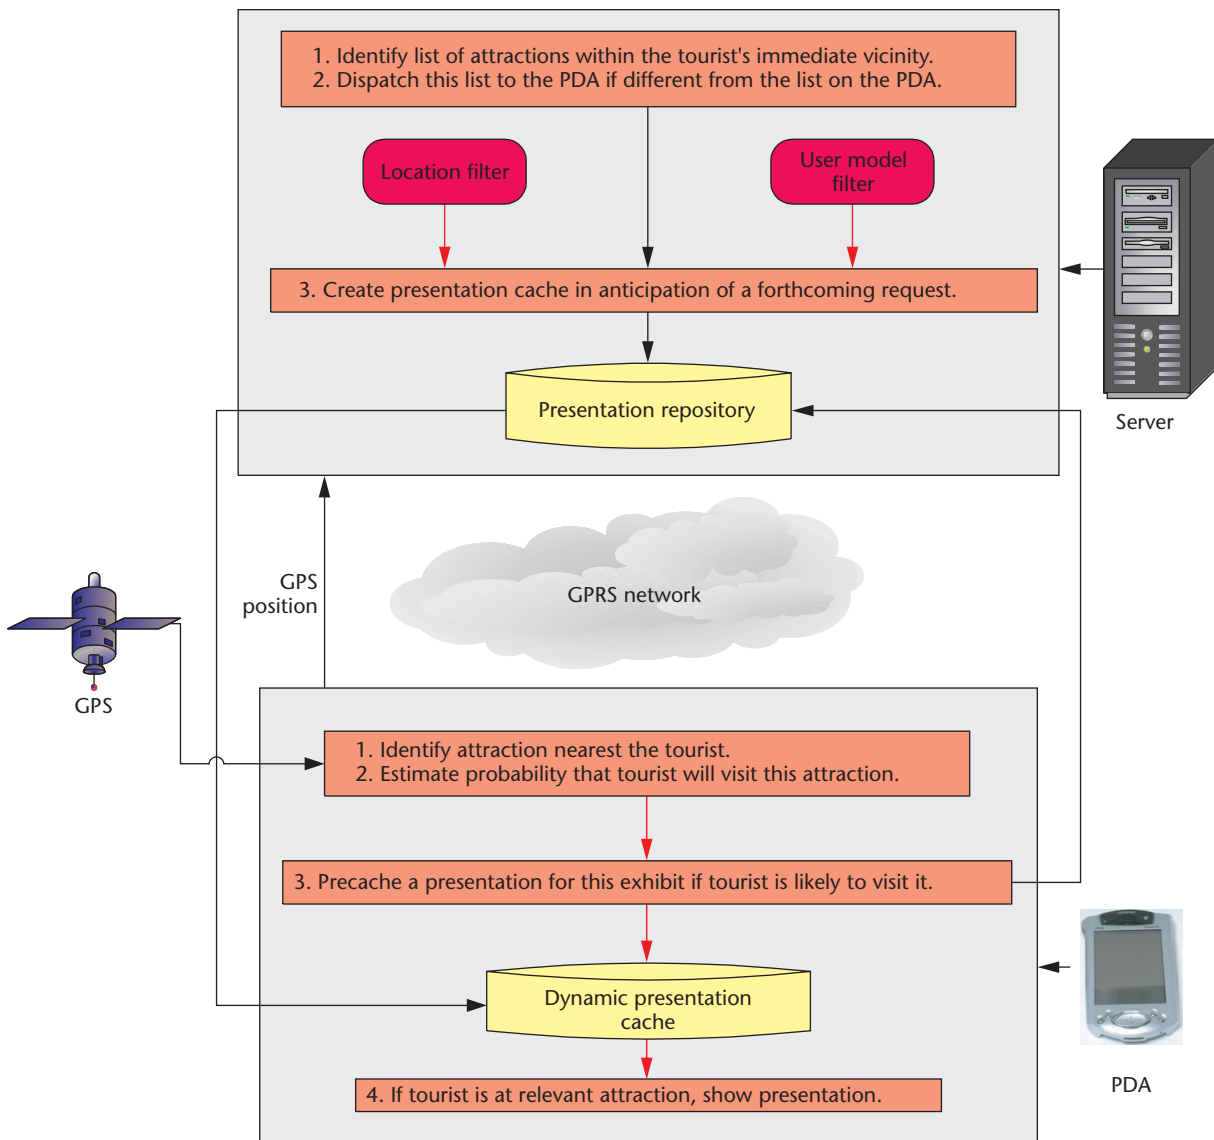


Figure 1. Essential stages of our intelligent precaching strategy.

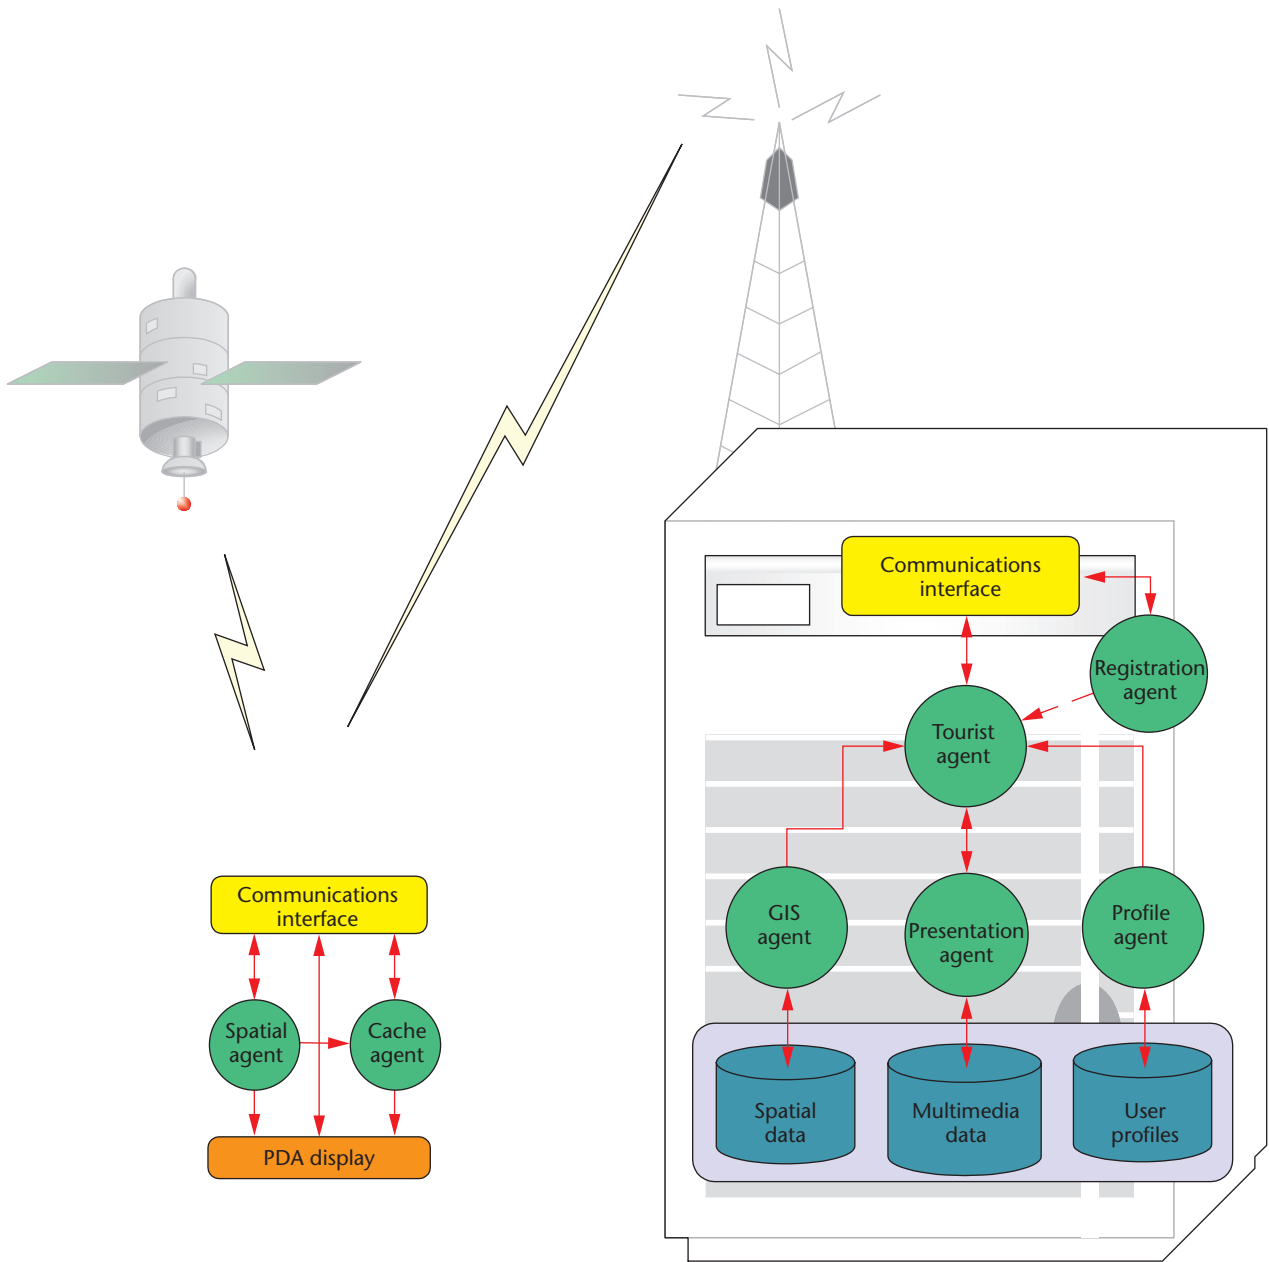
assignment of appropriate tasks to individual agents—an attractive facility from a system engineering perspective. In this way, the agents can then collaborate to deliver the necessary service to users. Therefore, intelligent agents form the basis of the architecture that we propose, next. Figure 1 shows some of the key steps in realizing intelligent precaching.

Architecture

To verify the validity of the intelligent precaching approach, we incorporated it into a system that we've developed in our research laboratory. This system, which we've dubbed Gulliver's Genie,¹⁴ is the result of ongoing research into various facets of mobile computing. Although Genie provides significant navigation support, it

is the Genie's strategy for intelligently delivering information to tourists that we focus on here. Tourists almost invariably seek out information of a cultural nature. The Genie's *raison d'être* is to deliver multimedia-enriched presentations to tourists on the various attractions that they encounter while exploring a city, an archaeological site, and so on. Such presentations consist of some combination of images, sound, and video. To maximize the tourist's interest, all presentations conform to the tourist's own personal interest model. For example, besides being delivered in the tourist's native language, presentations can be supplied on additional topics of interest that tourists choose, such as architecture, assuming that it's relevant to the attraction in question.

Gulliver's Genie conforms to a relatively



straightforward client–server architecture. The tourist’s PDA hosts the client component, and the PDA is connected to the server component via a GPRS connection. Intelligent agents (see “The Intelligent Agent” sidebar, next page) form the Genie’s fundamental building blocks. The Genie can be regarded as a multiagent system encompassing agents on the client and on the server, all collaborating to deliver the required services to tourists responsively, efficiently, and economically. Figure 2 outlines the architecture of Gulliver’s Genie, and its constituent agents are now described.

Spatial agent

The spatial agent monitors the tourist’s spatial context, that is, position and orientation. From these, the agent can make some deductions about behavior—for example, is the tourist moving or stationary? All of this information is shared with other interested agents, particularly the cache and tourist agents. To carry out its task successfully, the agent continuously monitors the GPS device and extracts position and orientation parameters. It also verifies that these parameters are consistent with previous values and that the quality of the sig-

Figure 2. Architecture of Gulliver’s Genie.

The Intelligent Agent

Intelligent agents offer an alternative paradigm for software development. In particular, complex and dynamic domains—which are difficult to model using traditional techniques—are potentially suitable for agents. An example of such a domain might well be the environment that the average mobile computer user operates within.

Agents come in many varieties, and researchers have proposed several taxonomies. One popular taxonomy classifies agents into weak and strong varieties.¹ Weak agents may be reactive, proactive, social, and autonomous. Strong agents may be rational, mobile, benevolent, and veracious (truthful) as well as inheriting some of the weak agents' characteristics. A classic example of strong agents has been realized through the belief-desire-intention (BDI) paradigm.² Briefly, beliefs represent a snapshot of the state of the agent's environment at a given time. *Desires* represent those objectives that the agent would like to fulfill or, in other words, its reason for existing. *Intentions* represent those desires that the agent is in a position to realize at a given time. Agents formulate intentions as a result of the agent's reconciling the existing model of its environment—represented through its beliefs—with its desires. The programmer will usually specify what beliefs or conditions are necessary before an agent can commit to realizing its individual desires.

Agents differ fundamentally from traditional software in that they're goal oriented and will only perform some task if it contributes to the realization of an agent's goal. In contrast, most

modern software systems are task oriented and comprise programs that perform tasks without any inkling as to what they're trying to achieve. Interestingly, the differences aren't clear-cut between objects—on which the object-oriented programming paradigm is defined—and agents. Some software practitioners consider agents as being nothing more than objects with some additional minor characteristics. Others regard both agents and objects as being intrinsically different, although sharing some characteristics. A detailed discussion of this viewpoint is beyond the scope of this article and may be found elsewhere.³ However, agents do bring a new way of thinking to the traditional software development process and offer an alternative mechanism for modeling and realizing solutions to problems in various domains.

References

1. M. Wooldridge and N. Jennings, "Intelligent Agents: Theory and Practice," *Knowledge Eng. Rev.*, vol. 10, no. 2, 1995, pp. 115-152.
2. M. Kinny, M. Georgeff, and A. Rao, "A Methodology and Modelling Technique for Systems of BDI Agents," *Proc. 7th European Workshop Modelling Autonomous Agents in a Multi-Agent World (MAAMAW 96)*, LNAI 1038, Springer Verlag, 1996, pp. 56-71.
3. J. Odell, "Objects and Agents Compared," *J. Object Technology*, vol. 1, no. 1, 2002, pp. 41-53.

nal used during the calculation was adequate before the agent briefs other agents about the tourist's status.

Cache agent

The cache agent implements the intelligent precaching strategy on the client. To do this, the cache agent depends on information from the tourist and spatial agents. From the tourist agent, it obtains a model of its immediate environment, which contains, among other things, a list of attractions within the tourist's immediate vicinity. From the spatial agent, the cache agent continuously receives position updates. Using both, it maintains a snapshot of the tourist's position with respect to nearby attractions. On deciding that the tourist is likely to visit a certain attraction, the cache agent requests a presentation on that attraction from the tourist agent and downloads it by the time the tourist reaches the attraction.

After downloading the presentation, the cache agent confirms that the user's position is still relevant and, if so, plays the presentation. If the

tourist hasn't yet reached the attraction, the agent waits until the tourist is in the attraction's vicinity before playing it. If the downloading process takes longer than planned, the tourist must wait. Alternatively, if the tourist has bypassed the attraction, the presentation won't be played; the cache agent will dispose of it when satisfied it's no longer required and that the tourist appears intent on visiting another attraction.

Tourist agent

All tourists are assigned their own individual agents on commencing a session with the Genie. In essence, this agent acts as the tourist's primary focal point for interfacing and accessing the Genie's services. In conjunction with the other agents, the tourist agent ensures that the environmental model is up-to-date and informs the cache agent of any changes. In conjunction with the presentation agent, it maintains a repository of prebuilt presentations in anticipation of a request from the cache agent. This repository is continuously updated to reflect any changes in the tourist's environmental model.

Presentation agent

The presentation agent builds presentations that are consistent with the tourist's individual profile. To do this, the tourist agent must provide it with details of both the tourist in question and the tourist's environmental model. It then proceeds to search the database for all relevant information on those attractions contained within the model. The Genie then filters these using the tourist's profile (obtained from the profile agent) and constructs the presentations.

Profile agent

The profile agent maintains the profiles of all users registered with the Genie and provides the relevant information to the presentation agent so that it can customize presentations accordingly. After a tourist has listened to a presentation, the cache agent returns an interaction record to the profile agent for further analysis and refinement of the tourist's profile. In this way, presentations are always consistent with the latest version of the tourist's profile.

GIS agent

The GIS agent advises on what a tourist's environmental model should consist of, given a particular geographical position. At present, it focuses on providing geographical information about attractions in the tourist's vicinity.

Registration agent

The registration agent administers the creation and termination of tourist agents as tourists join and leave the Genie community. It provides the initial contact point for tourists wishing to use the Genie's services.

Database

Underlying Gulliver's Genie is a multifaceted database, in which data is stored as belonging to one of the following sections:

- *Geospatial*: Information concerning the position of all tourist attractions known to Gulliver's Genie is stored in this section. In particular, the storage and management of electronic maps is crucial.
- *Multimedia*: All Genie presentations consist of some combination of sound, images, text, or video. These media elements, including relevant metadata, are stored in the multimedia section.
- *Profile*: All prospective users of Gulliver's Genie must first register with the registration agent. Details required include language, nationality, age group, and gender. Users are invited to indicate their cultural interests, for example, art, architecture, and popular culture. This initial snapshot, together with their personal details, is stored as a unique user profile.

The geospatial and the multimedia data comprise what we call a region's information space. Before the tourist can use the Genie, we must first define an information space for the geographical area in question. Based on this, the GIS agent can construct models of the tourist's immediate environment for closer observation.

Implementation details

We selected an iPAQ H3850 to host the Genie. Because it can be augmented with a dual-slot expansion sleeve, we can readily incorporate GPS and GPRS PCMCIA cards. The PDA connects to a server over GPRS using the standard Internet protocol HTTP. The runtime environment on the client and server is based on Java with the Jeode Java virtual machine (JVM) being used on the iPAQ. We implemented all the agents with Agent Factory (<http://www.agentfactory.com>).¹⁵ Agent Factory, which features a lightweight implementation of its reasoning engine for PDAs, provides a complete environment for designing and realizing BDI agents (that is, agents incorporating *beliefs*, *desires*, and *intentions*). We used IBM's DB2 to host the database.

Deploying multimedia on PDAs has proved somewhat problematic aside from the wireless communications issue. Displaying images is straightforward because PDAs support most of the important formats. The situation becomes more complicated when audio and video are involved. The Java Media Framework (JMF) provides developers extensive and sophisticated features for working with multimedia in a multiplatform environment but is limited on PDAs. While we have successfully deployed the JMF on the iPAQ, it supports only a fraction of the formats that it does on a normal workstation environment. Furthermore, although the JMF can successfully render a number of formats, in practice some are too slow on a PDA to be practical. Eventually, we settled on the reliable but somewhat low-quality au format for rendering audio. Unfortunately, rendering a video at an acceptable rate was impossible, so we aren't including video in any presentations yet.



Figure 3. Screen snapshots of the Genie, showing presentations in (a) and (b).

What the tourist experiences

When the tourist encounters an attraction, Genie automatically displays a presentation. All presentations follow a similar template and consist of an image, image title, and a number of hyperlinks (see Figure 3a).

Each link constitutes a concise information snippet about some aspect of the respective tourist attraction and consists of

- a name to indicate the link's subject: for example, architecture. By convention, the first link is usually used to identify the attraction, thus allowing the tourists time to orient themselves.
- a sound recording on the link's subject.
- an optional image to reflect the link's subject, should it merit its own image. If not, the Genie automatically displays the default image of the relevant exhibit.

Tourists activate all links using a stylus, the default interaction mechanism for interacting with a PDA, although we have programmed some shortcuts using the iPAQ's navigation pad.

An interface structure like the one we've just described would be adequate for most tourists, but the Genie must also attempt to cater to those tourists who might have more experience in certain topics and thus have higher expectations. For example, a general description of a building's architecture might be adequate for most people.

However, for those with experience or additional expertise in the field, a general description might be inadequate, and an opportunity could be lost if we don't customize the presentation accordingly. To cater to such tourists, the Genie offers users the possibility of accessing further information through what we call follow-ups. If a follow-up is available for a link, a plus or expansion icon is placed beside it; otherwise, the default negative symbol is used. When the user selects this, the link expands into a sublist of highly specialized sublinks that the tourist can access in the normal way. Figure 3b shows a simple example.

How agents deliver intelligent precaching

As tourists explore an area, the Genie's spatial agent continuously monitors their movements and informs the tourist agent of the latest position readings. Based on these, the GIS agent advises on what the contents of the cache agent's environmental model should be, and the tourist agent informs the cache agent accordingly. Simultaneously, the tourist agent arranges with the presentation agent to build and store presentations consistent with the attractions in the environmental model, in anticipation of a future request from the cache agent.

Meanwhile, on the PDA, the cache agent, acting on position updates from the spatial agent, attempts to anticipate where the tourist is likely to visit. To do this, it continuously monitors the tourist's position relative to the nearby attractions (obtained from the environmental model). When it deduces that the tourist is converging on a certain attraction, the cache agent requests the corresponding presentation from the tourist agent and stores it in its cache on the PDA. When the tourist encounters the attraction, the Genie automatically displays the presentation. Should the tourist have a late change of heart and move elsewhere, the presentation will remain in the cache until another is requested, whereupon it is discarded. This entire process takes place transparently to the user and is continually repeated as long as the tourist uses the Genie services.

Intelligent agents are particularly suited to implementing the precaching mechanism. In continuously monitoring the tourist's behavior, agents can react to events as they happen, and reason about the implications of any action before proactively seeking to anticipate and address possible future needs and requests. From a modeling and implementation perspective, the

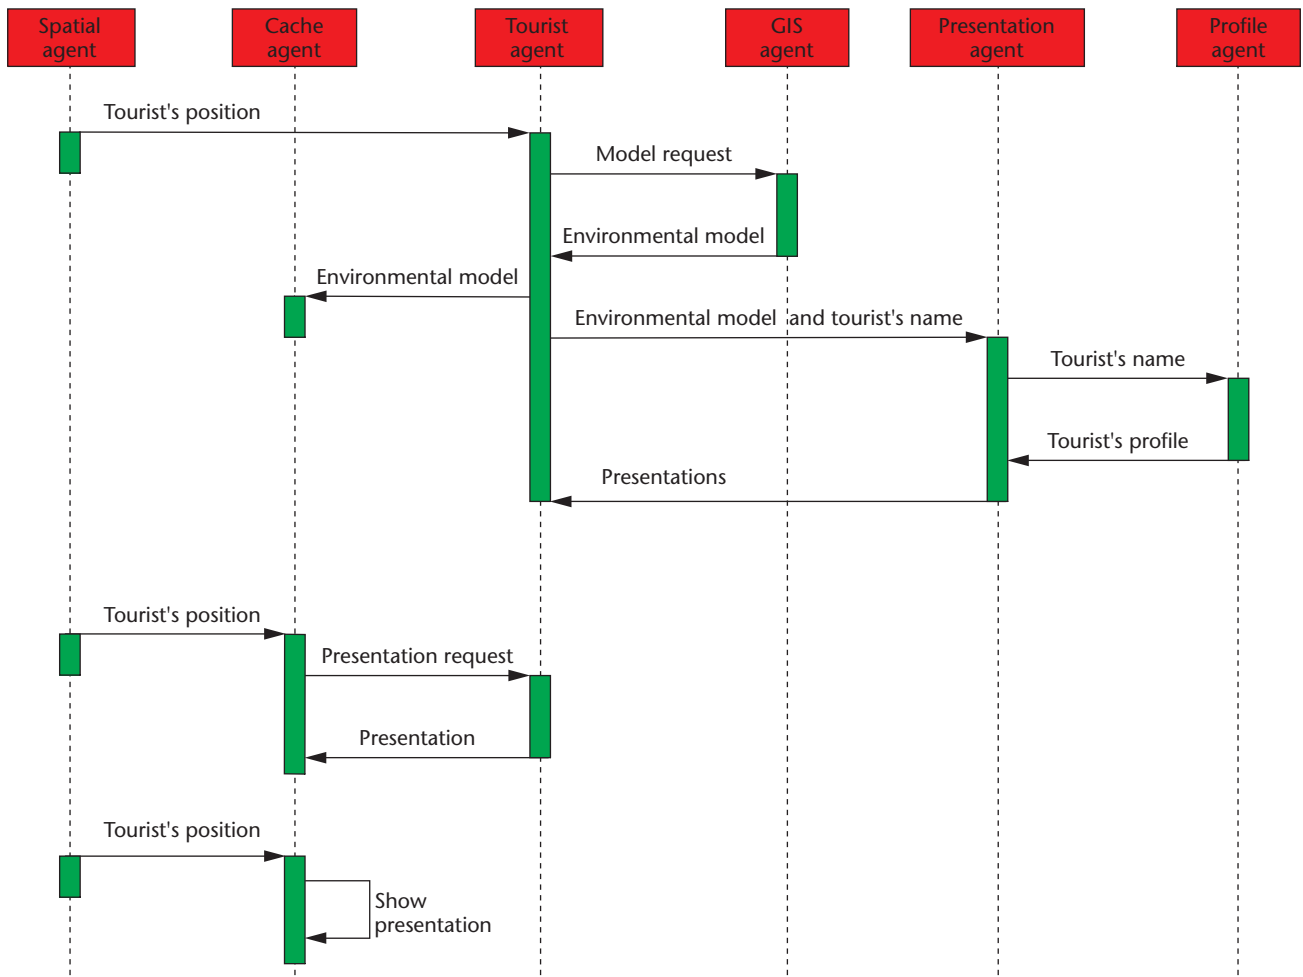


Figure 4. Unified Modeling Language (UML) sequence diagram illustrating the essential interactions between agents as they collaborate to deliver a precaching solution.

agents' social ability enabled us to adopt an intuitive and modular approach to system design. Figure 4 shows an example of the pertinent interactions.

Evaluation feedback

To verify the approach taken with the Genie, we conducted user evaluations and trials. Before doing this, we defined an information space by identifying the most interesting buildings and art works on the campus of University College Dublin. Fortunately, the campus is spread out, comprises a number of historical buildings, and is well endowed with artworks of various kinds. Therefore, we were able to include attractions that would appeal to a wide spectrum of users (see Figure 5).

A sample group of 40 users took part in the evaluation. This group—56 percent male, 44 per-



Figure 5. A tourist is consulting the Genie about a sculpture that he has encountered.

cent female—tended toward a younger age group with 65 percent being younger than 30 years of age. Most had some computing experience, as expected, but only 44 percent had some prior experience with PDAs. From an occupational perspective, the group was mixed, comprising students, visiting students from other universities, campus employees, and some day visitors to the university. All evaluations took place during daylight hours and under varied, although not wet, weather conditions. Each evaluation took about 45 minutes after which we asked the users to complete a questionnaire.

The questionnaire asked them to rate various aspects of the Genie on a Lickert scale of 1 to 7 (7 meaning being completely satisfied) as well as make some recommendations for improvements. Overall, 85 percent of participants were quite satisfied with the system, scoring it within the 5–6 range; 56 percent were very satisfied with the navigational aspects of the Genie, scoring it within the 6–7 range. The level of satisfaction regarding the presentations was excellent with 84 percent of participants scoring it within the 5–7 range. Finally, more than 70 percent were satisfied with the system functionality, scoring it within the 5–7 range.

We used open feedback to tease out what users considered to be the Genie's strengths and weakness as well as to identify how the system could be improved. Users regarded navigation support and the Genie's intuitive interface as its principal strong points. A third characteristic that users rated highly was the Genie's potential to deliver information at the right place and time.

Users perceived Genie's critical limitation to be its responsiveness when using the stylus. More interestingly, users noted that the timing of the presentations was occasionally out of phase because they were displayed after the user had passed a particular attraction. This, however, was an implementation decision we'd made because we decided that even if the user had moved beyond the attraction, the presentation would still be shown. Although this "feature" was observed only by 10 percent of our test group, it does suggest a need to review and refine the pre-caching process. In particular, the agents' ability to make up for the inherent inaccuracy of GPS and the variable data rates delivered by GPRS needs to be strengthened further, possibly through the incorporation of 3G and Satellite-Based Augmentation System (SBAS) technologies.

That said, GPS and GPRS performed reliably

during the evaluations, contrary to our initial expectations. Due to the nature of the campus, GPS coverage wasn't an issue, and we didn't experience a discernible loss of signal. In the case of GPRS, only 4 percent of users experienced signal loss, but because almost instant reconnection was available, it wasn't perceived as a major problem.

A third issue that arose was the perceived control of the Genie. Several users had reservations about the Genie's autonomous nature and expressed a desire to actively decide for themselves when the Genie should download data. It may well be that such users don't quite understand the implications of this. However, it seems reasonable that they should be given the choice, particularly as they may have to pay for such a service at some point, so we'll include a facility to let them specify this feature of the Genie's behavior in the next version.

Observations

Overall, we were encouraged by the results of the evaluations, despite the fact that several areas have been identified for improvements. Therefore, we intend to continue augmenting and refining the agents' behavior, thus improving the Genie's robustness and reliability. We're satisfied that intelligent precaching offers a viable approach to multimedia dissemination; nevertheless, some improvement is necessary before it could form the basis of a commercial application. Although reliability isn't an issue, users would expect the Genie to have completed the precaching of the appropriate content at the appropriate time, which is difficult given the unpredictable nature of the data rates supported by wireless networks and the inherent position errors in GPS.

Fortunately, two on-going developments will address these problems. First, the somewhat protracted migration to 3G networks will increase the available data rates, thus reducing download time. Second, the development of SBASs, for example, the Wide Area Augmentation System (WAAS; <http://gps.faa.gov/Programs/WAAS/waas.htm>) and the European Ground Navigation Overlay Service (EGNOS; <http://www.esa.int/export/esaNA/index.html>), will reduce positional errors. Such technologies are currently being integrated into GPS receivers and will lead to position estimates in the 5-meter range; a significant improvement in position accuracy over the current 20-meter norm.

In considering when such an approach should

be contemplated, two situations immediately spring to mind:

- The mobile user has a device of limited storage capacity.
- The multimedia content of a device has a dynamic component, such as new exhibits in a museum.

Finally, we must consider the nature of the application domain. The tourism domain is ideal, because attractions are usually liberally dispersed and there's sufficient time to deduce the tourist's destination before the Genie dynamically assembles the presentation and downloads it to the device. Alternatively, some domains may be particularly suitable to a hybrid solution. Static content could be stored on the device using a smart card, and the dynamic content could be precached on the device as the occasion demands.

Future developments

Given the experience and expectations people currently have for technology, the Genie's lack of support for video could be viewed as a significant deficiency. However, several promising solutions are already on the horizon. The most prominent of these, MPEG-4,¹⁶ is a comprehensive open standard aimed at delivering audio and video content to heterogeneous devices across networks of varying bandwidth. If we regard the Genie as a potential 3G service, as distinct from an application in its own right, the use of a standard mechanism for defining and delivering presentations would prove highly desirable. The Synchronized Multimedia Integration Language (SMIL) 2.0¹⁷ offers an intriguing possibility here, particularly when we consider that SMIL forms an intrinsic component of the 3GPP transparent end-to-end packet-switching streaming service specification.¹⁸ Should either or both of these possibilities come to fruition, the possibility would exist for identifying heuristics that agents could exploit to determine the circumstances when either a streaming or precaching solution would be preferable.

Conclusion

Our work with Gulliver's Genie will continue in a number of areas. Initially, we plan to incorporate more agents to provide additional services. The navigation component of the user interface requires additional enhancements to cater for users of differing abilities, and we intend

to expand the Genie's support for further multimedia data types. For example, we plan to augment the 2D map with 3D constructs. We also envision incorporating panoramic images and the use of Virtual Reality Modeling Language (VRML) or Java3D for historical reconstructions. Another of our plans is the use of scrolling text commentaries to accommodate individuals who might have hearing difficulties.

Intelligent precaching provides an alternative mechanism by which multimedia data may be made accessible to mobile computer users. In particular, it can help minimize and disguise two of the prominent limitations facing prospective suppliers of services to mobile users: poor data rates supported by the current generation of wireless data networks and the limited computational resources of popular mobile devices. Given that mobile computing is likely to be the prevailing computer usage paradigm in the coming years, the necessity to deliver a complex range of multimedia data in a timely manner will become increasingly urgent if mobile computing in its various guises is to fulfill its potential. **MM**

Acknowledgments

Michael O'Grady gratefully acknowledges the support of the Irish Research Council for Science, Engineering and Technology (IRCSET) through the Embark Initiative. Gregory O'Hare gratefully acknowledges the support of Science Foundation Ireland (SFI) under grant no. 03/IN.3/1361.

References

1. T. Pfeifer, T. Magedanz, and S. Hübener, "Mobile Guide—Location-Aware Applications from the Lab to the Market," *Proc. 5th Int'l Workshop Interactive Distributed Multimedia and Telecommunications Services (IDMS 98)*, LNCS 1483, Springer Verlag, 1998, pp. 15-28.
2. M.J. O'Grady and G.M.P. O'Hare, "Accessing Cultural Tourist Information via a Context Sensitive Tourist Guide," *Information Technology and Tourism J.*, vol. 5, no. 1, 2002, pp. 35-47.
3. G. Abowd et al., "Cyberguide: A Mobile Context-Aware Tour Guide," *ACM Wireless Networks*, vol. 3, no. 5, ACM Press, 1997, pp. 421-433.
4. N. Davies et al., "Using and Determining Location in a Context-Sensitive Tour Guide," *Computer*, vol. 34, no. 8, Aug. 2001, pp. 35-41.
5. V. Vlahakis et al., "Archeoguide: An Augmented Reality Guide for Archaeological Sites," *IEEE Computer Graphics and Applications*, vol. 22, no. 5, Sept./Oct. 2002, pp. 52-60.

6. F. Bellotti et al., "User Testing a Hypermedia Tour Guide," *IEEE Pervasive Computing*, vol. 1, no. 2, April-June 2002, pp. 33-41.
7. S. Poslad et al., "CRUMPET: Creation of User-Friendly Mobile Services Personalised For Tourism," *Proc. 2nd Int'l Conf. 3G Mobile Communication Technologies*, IEE, 2001. pp. 28-32.
8. C. Hoymann and P. Stuckmann, "On the Feasibility of Video Streaming Applications over GPRS/EGPRS," *Proc. IEEE Global Telecommunications Conf. (Globecom 02)*, IEEE Press, 2002, pp. 2490-2494.
9. R. Prasad and M. Ruggieri, "Future Trends in the Wireless Multimedia Environment," *Proc. Int'l Conf. Communication Technology (ICCT 03)*, IEEE Press, 2003, pp. 11-20.
10. D. Barbar and T. Imielinski, "Sleepers and Workaholics: Caching Strategies in Mobile Environments," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 1994, pp. 1-12.
11. Z. Jiang and L. Kleinrock, "Web Prefetching in a Mobile Environment," *IEEE Personal Comm.*, vol. 5, no. 5, Oct. 1998, pp. 25-34.
12. J. Xu, X. Tang, and D. Lee, "Performance Analysis of Location-Dependent Cache Invalidation Schemes for Mobile Environments," *IEEE Trans. Knowledge and Data Eng.*, vol. 15, no. 2, Mar./Apr. 2003, pp. 474-488.
13. S.-Y. Wu and K.-T. Wu, "Dynamic Data Management for Location Based Services in Mobile Environments," *Proc. 7th Int'l Database Eng. and Applications Symp. (IDEAS 03)*, IEEE Press, 2003, pp. 180-191.
14. G. O'Hare and M. O'Grady, "Gulliver's Genie: A Multi-Agent System for Ubiquitous and Intelligent Content Delivery," *Computer Comm.*, vol. 26, no. 11, Elsevier Press, 2003, pp. 1175-1185.
15. R.W. Collier et al., "Beyond Prototyping in the Factory of Agents," *Proc. 3rd Int'l Central & Eastern European Conf. Multi-Agent Systems (CEEMAS 03)*, LNCS 2691, Springer Verlag, 2003, pp. 383-393.
16. F. Periera and T. Ebrahimi, *The MPEG-4 Book*, Prentice Hall, 2002.
17. World Wide Web Consortium (W3C), "Synchronized Multimedia Integration Language (SMIL 2.0)," W3C Recommendation; <http://www.w3.org/TR/2001/REC-smil20-20010807/>.
18. 3GPP TS 26.246 V1.0.0, "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Transparent end-to-end packet switched streaming service (PSS); 3GPP SMIL Language Profile (Release 6), 3GPP, 2003.



Michael O'Grady holds a post-doctoral research fellowship at University College Dublin. His research interests include agent technologies, multimedia, and mobile computing. O'Grady

received a BSc and a postgraduate diploma in mathematical modeling, both from the National University of Ireland. Following this, he spent a number of years in industry where he worked in diverse areas including expert systems and telecommunications. He is a member of the ACM and the IEEE.



Gregory O'Hare is head of the Department of Computer Science at University College Dublin and director of the department's Practice and Research in Intelligent Systems & Media

(PRISM) Laboratory. His research focuses on multi-agent systems and on mobile and ubiquitous computing. O'Hare has a B.Sc. in computer science and an M.Sc. in information technology from the University of Ulster. He is a Fellow of the British Computer Society and a Science Foundation Ireland (SFI) Investigator.

Readers may contact Michael O'Grady at michael.j.ograde@ucd.ie.

For further information on this or any other computing topic, please visit our Digital Library at <http://www.computer.org/publications/dlib>.