Mobile Augmented Reality Archæological Reconstruction

Panagiotis RITSOS, David J. JOHNSTON, Christine CLARK, Adrian F. CLARK

VASE Laboratory, University of Essex, Colchester UK

ABSTRACT

This paper describes a mobile augmented reality system, based around a wearable computer, that can be used to provide *in situ* tours around archæological sites. Ergonomic considerations, intended to make the system usable by the general public, are described. The accuracy of the system is assessed, as this is critical for *in situ* reconstructions.

Keywords: augmented reality, wearable computing

1 INTRODUCTION

People are fascinated by archæology. They want to learn how their ancestors looked, what they did, how they lived — but the proportion of archæological sites where there are substantial remains above ground is small. Indeed, the curators of archæological sites are faced by a dilemma: they wish to attract interested visitors but need to do this without disturbing any remaining archæology. This usually rules out *in situ* physical reconstructions.

The traditional way of helping people visualize the appearance of archæological sites is through paintings and physical scale models. In recent years, these have increasingly given way to computer graphic "walk-throughs." Ultimately, however, both of these are unsatisfactory because it is difficult for a visitor to relate the appearance of the model (physical or virtual) to the archæological remains. A better solution is for the visitor to be able to visualize the appearance of the model as they walk around the site. This is conventionally achieved through labelled perspective views on plaques.

However, by exploiting the simultaneous reduction in size and increase in performance of computer technology, one can imagine a system that uses virtual reality technology to visualize the ancient buildings as the visitor explores the site. To achieve this kind of mobile augmented reality (though not specifically for archæological reconstructions), several groups have proven the potential of using a wearable computer with head-mounted display and position-sensing technology to superimpose the reconstruction onto the visitor's own view of the surroundings; see Section 2. This research follows the same general route but differs from preceeding work in two ways: it attempts to produce a prototype of a system that would be suitable for use by the general public; and a careful consideration is given to the sources of errors in position and orientation as accurate co-location of the computer model with the archæological remains is vital.

The remainder of this paper is organized as follows. The next section gives an overview of previous work on wearable tour guides. Section 3 then gives a brief introduction to the archæological site that has been formed the focus of the work, highlighting its important properties. Different possible designs for the overall system are considered in Section 4, and Section 5 goes on to describe the practical problems that were encountered during implementation and how they were overcome. Section 6 describes the system in use and considers the sources of positioning errors and their severity. Finally, Section 7 draws conclusions and outlines further work.

2 PREVIOUS AND RELATED WORK

One of the first mobile tour guide systems was Lancaster University's GUIDE [CHEVERST *et al.*, 2000]. GUIDE provides city visitors with location-related information using a tablet PC equipped with a radio-LAN (802.11b) interface. GUIDE does not attempt to produce 3D reconstructions of the surroundings but instead provides text and images related to the user's position. A number of base stations around the city provide each user with information relevant to their position. The authors argue that GPS, the *global positioning system* which they used in early versions of the system, does not provide any advantages in such an environment when compared to network-based location mechanisms. In particular, GPS requires at least four satellites to be in view in order to obtain a moderately accurate position fix, yet this is rarely possible in the "urban canyons" formed by tall buildings.

One of the earliest research efforts in the field of mobile augmented reality is Columbia University's *Touring Machine* [FEINER *et al.*, 1997], which has evolved into the *Situated Documentaries* [HÖLLERER AND PAVLIK, 1999]. This employs a GPS-equipped wearable computer to provide hypermedia presentations that are integrated with the actual outdoor locations to which they pertain. The prototype uses a tracked, see-through HMD and a hand-held pen computer to present 3D graphics, imagery and sound superimposed on the real world. A backpack wearable computer, equipped with a GPS receiver, provides the location information. The user roams within the university campus and is able to see information in the form of virtual tags, indicators and point-of-interest flags through the HMD.

A similar system, developed at Carnegie-Mellon University, is known as *Smart Site* [YANG *et al.*, 1999]. This utilizes a wearable computer and a laptop to provide an intelligent tourist assistant system. The system architecture caters for multi-modal input and output. Location information is obtained from a GPS unit, while a camera provides visual input capabilities. A microphone and headphones provide audio input and output. Speech and gesture inputs may be used. The authors argue that for a tourist system it is important to use a number of input modalities to accommodate different scenarios. A user can roam in a site, derive location information from the GPS unit but also request for further details by speech or by gesture.

One of the most significant systems in this area is Tinmith-Metro, an interactive augmented reality 3D constructive solid geometry modeller [PIEKARSKI AND THOMAS, 2001]. This system is based around a wearable (backpack) computer with machine vision capabilities utilizing a USB camera. User input is principally via a set of pinch gloves, allowing the user to execute commands using menus linked to finger gestures; indeed, the authors argue that desktop user interfaces should be avoided and replaced with speech recognition, camera input and hand-gesture tracking. Position is determined using (differential) GPS, while a digital compass measures head orientation. The system is capable of generating 3-D models of the external surfaces of buildings as the user roams, and can place pre-fabricated 3D objects within a scene. Objects are created by fitting infinite planes to surfaces and using their intersections to define buildings.

Archeoguide [ARCHEOGUIDE, 2002] is a EU-funded programme that aims to result in a virtual tourist guide. Its architecture, although similar in visual output to the one described here, has some fundamental differences. The system uses a centralized site information server, a mobile unit with GPS location, and a wireless LAN for communication. The site server includes an image database and a content creation mechanism. All rendering is performed using a VRML-based VR toolkit developed within the programme.

3 BACKGROUND

The invasions of Julius Cæsar of Britain in 55 and 54 BC were supposedly to avenge the murder of the king of the Trinovantes, who inhabited modern Essex and southern Suffolk, by his rival, the king of modern Hertfordshire. After Cæsar's departure, the two tribes were united into a single royal house which established its capital at *Camulodunum* ("fortress of the war god"), to the south-west of modern Colchester at a place now known as the Gosbecks Archæological Park. Camulodunum exported grain, meat, hides, hunting dogs and even oysters to the Roman world and reached its zenith under the reign of Cunobelin (Shakespeare's Cymbeline).

Shortly after Cunobelin's death, Claudius decided to bolster his reputation by conquering Britain. In 43 AD, his 40,000-strong army landed in Kent and fought its way northwards, stopping just short of Camulodunum. The Roman commander sent for Claudius who travelled from Rome to lead his troops into the city in person. In 49 AD, the first *colonia* (colony) of the new province of Britannia was established next to Camulodunum, on the site of modern Colchester. The Gosbecks area was also developed further with a 5,000-seat theatre, the largest in Britain, and a temple complex.

The colonists' treatment of the locals was apparently not good. In particular, when Prasutagus, king of the Iceni, died in 60 AD, his wife Boudicca ("Bodicea") was beaten by Roman officials and her daughters raped. The Iceni rose in fury and destroyed the colony, burning its temple (the site of the modern castle) where many colonists had taken refuge. Boudicca went on to ravage London and *Verulamium* (at the edge of modern St. Albans) before being apprehended and killed. The colony at Colchester was rebuilt, this time with the extensive walls that remain today.

Following the end of Roman rule circa 450 AD, future development was centred within the city walls, in modern Colchester. The Gosbecks site proved an excellent source of building materials for local inhabitants — so much so that all that effectively remains today are foundations, and these are below the ground level. Indeed, the only way that a visitor to Gosbecks is aware of the scale and layout of the former buildings is by white lines painted on the ground and accompanying signage.

Our work to date has concentrated on reconstructing the Roman temple complex at the Gosbecks site; we shall extend it to encompass the theatre in the future. Gosbecks has some characteristics that, as we shall see later, reduce some problems while introducing others. Firstly, it is situated in a reasonably flat area (modulo rabbit holes) at the edge of farmland near the brow of a low hill; the nearest buildings are over 500 m away. Hence, almost an entire hemisphere of sky is visible. There are no trees, pylons, or post-Roman remains on the site to interfere with our reconstructions. There are also no visible foundations that the 3-D reconstructions have to abut. On the other hand, much of the site has not yet been excavated, particularly the area surrounding the likely grave of Cunobelin; so English Heritage, the owner, is understandably protective of the undisturbed archæology and does not permit the ground to be disturbed, such as by erecting further signage or radio masts.

4 DESIGN CONSIDERATIONS

4.1 Accuracy Requirements

Our basic aim is that a person should perceive the 3D reconstruction as being "in the right place," corresponding to an accuracy requirement for position of the order of 1 cm, clearly not achievable with the kinds of systems currently being explored. A more achievable aim is that doorways in the model should be sufficiently positionally stable that the wearer can negotiate physical doorways without difficulty — say 0.2 m. Accuracy requirements for orientation are not, we believe, too stringent: 5° error seems to be acceptable to most users.

4.2 Determining position and orientation

In the laboratory, the most common way of estimating position and orientation is via magnetic trackers, exemplified by Ascension's *flock of birds*. However, their region of operation is a few metres at best, so they are unsuited to this application. The ideal solution for an archæological site would probably be to mount radio beacons on masts and use triangulation; but we cannot do that on the Gosbecks site for the reason expounded above.

If triangulation using fixed beacons is not permitted, the obvious fall-back is to use the well-known Global Positioning System (GPS) — *i.e.,* triangulation with moving, remote satellite beacons. Single-receiver GPS, even after the removal in May 2001 of the "selective availability" random error superimposed on signals, is unlikely to provide sufficient accuracy for this application; instead, *differential* GPS is required. This involves the use of two receivers, a stationary *reference point* sited at a known position, and a roaming one situated on the wearable computer. As long as the two receivers are within a few kilometres of each other, the GPS signals that reach them travel through virtually the same atmospheric section and therefore have the virtually the same errors and delays. The reference point uses its known position to derive what the signal travel time should be, measures the actual travel time and, by subtracting these two values, calculates the error of the received signal. This error is then transmitted to the roaming receiver for correction purposes. In the UK, differential signals are broadcast on FM radio or by coastal beacons. For our work, the preferred approach is to use a laptop computer with attached GPS receiver at a well-defined position on the Gosbecks site, and to use a radio-LAN to broadcast the error signal to the wearable systems. We currently use Aviator Pro wireless LAN cards with an effective range of about 150 m. Conveniently, each roaming computer can also broadcast its own position; in the future, we shall use this to superimpose toga-clad avatars on other visitors.

This basic differential method can be made more accurate by measuring the phase of the GPS carrier signal (as well as the signal itself) as this is necessarily at a much higher frequency than the GPS data. Some GPS receivers are designed to make available or exploit carrier phase information, but such devices are too expensive for this application. However, a group of people scattered around the world have decoded the internal protocols used by one particular GPS receiver, and this provides information that is accurate enough for our needs when used in conjunction with home-grown differential GPS software.

GPS, even differential GPS, yields only position information. A separate sensor is required in the user's HMD to determine orientation. The particular device used for this work, Virtual i/O's *I-glasses*, incorporates an electronic compass and tilt sensor; while not particularly accurate, we have found (rather to our surprise) that it works well enough outdoors and meets our orientation accuracy requirement.

4.3 Wearable computers

The most significant question regarding the capabilities of the wearable computer concerns where the rendering operation takes place. Most researchers have been concerned with only a single user and hence have performed all model visualization on the wearable itself. On the other hand, Archeoguide (see above) intends to render its models centrally, transmitting the resulting 2D imagery to the wearables. With a radio-LAN being necessary for differential GPS in any case, the latter approach could allow the wearables to be somewhat simpler. However, it is the authors' opinion that centralized rendering does not scale well to significant numbers of participants, so we have produced a distributed solution in which each wearable renders its own view.

The consequence of employing distributed rendering is that the wearable cannot use a low-powered processor such as a StrongARM: we are forced to use a high-performance processor such as an Intel Pentium, with the consequent power requirements and weight penalties.

5 IMPLEMENTATION ISSUES

5.1 Wearable computing issues

The wearable computer is based on the popular "Tin Lizzy" architecture and was initially belt-worn. The system currently uses a Digital Logic PC/104 motherboard equipped with a 266 MHz Intel Pentium processor with 64 Mbyte of memory, four serial ports, one USB, VGA and sound. (The motherboard can now be upgraded to a 700 MHz processor.) A PCMCIA adaptor is used to take a radio-LAN card. An IBM TravelStar hard disk is currently used, though we expect a production system to replace this with a compact flash card. An aluminium enclosure is used for protection, security, and to dissipate heat. The *I-glasses* HMD attaches to one serial port and the external GPS unit to another. The HMD's driver unit is powered from the wearable via a DC-to-DC converter. User input currently comes from a HandyKey Twiddler2 [HANDYKEY CORPORATION, 2002], attached via USB, on our prototype system. Power is supplied by two camcorder batteries, located in a separate unit for easy exchange. The operating system is Red Hat Linux 7.2. Our application code is written in C and uses the Mesa OpenGL-compatible graphics library.

Our initial system used a CMC Allstar 12-channel GPS board within the computer enclosure. Although this is reputed to provide carrier phase information, data from it were found to be unreliable; this was eventually traced to a firmware problem. Moreover, the card was found to be affected by electrical noise from the computer components. For this reason, we adopted an external receiver, a Garmin G12XL.

The major problem with the current configuration is power consumption. The two camcorder batteries provide a lifetime of only about forty-five minutes; we expect most users would spend about an hour on their tour, so we really need two hours' worth of power. However, increasing the number of cells adds significantly to the total weight and also increases recharge time.

The weight of the belt-carried system, including all cabling *etc.*, is 4.9 kg; we estimate that this figure can be reduced by about 20%. However, the system is bulky and becomes uncomfortable after about an hour's carriage. Consequently, we have investigated integrating the various components into a photographer's many-pocketed jacket. The main enclosure then fits into the large rear pocket and the GPS receiver and batteries into the side pockets. The GPS antenna can be attached to an epaulet. The only remaining cables are then for the Twiddler (stored in a front pocket when not in use) and the HMD. Trials suggest this is much more comfortable, especially when the wearer is not in the first flush of youth. Equally important, it is much easier for visitors to the site to put on and take off the system when installed in the jacket than when attached to the belt.

(a) Belt-mounted configuration in use

(b) Jacket-mounted configuration

Figure 1: Wearable tour guide system

5.2 The 3-D model of the temple complex

Roman architecture was based on the principles of the earlier, more elegant Greek building styles; however, their approach was more formulaic to enable faster construction and employ less skilled artisans. Fortunately, a guide to building design due to Vetruvius [Morgan, 1967] has survived from antiquity. With this guide, the authors have been able to write a program to generate 3-D models of Roman buildings from a few key measurements such as the number of columns and their diameters.

The temple complex comprises a square-shaped portico, with an outer wall, an outer circuit of Doric columns and in inner circuit of Ionic columns, all being covered by a tiled roof. The entrance faces roughly east. Within the portico is a ditch and, in the south-east of that, the main temple (Figure 2).

The model is a large one, and the wearable is currently equipped with only a generic VGA card — graphics cards with adequate 3-D acceleration are only now becoming available at a reasonable cost in the PC/104 form factor — so rendering performance on the 266 MHz processor is poor. The tour guide application is able to perform several simplifications to the model to improve performance. These include rendering the columns, which are nominally cylindrical, as sets of eight octahedral prisms; avoiding textures; and using simple lighting. In fact, the latter is essential in any case in order to ensure the model is easily distinguishable through the low-resolution HMD in daylight conditions.

Our initial implementation of a user-level application used Tcl to generate VRML models of the Gosbecks site and a VRML viewer for navigation, providing links, *etc.* However, we found that VRML viewers for Linux were either closed-source, slow, or unstable. Consequently, the application software has been written *ab initio* using OpenGL and GLUT.

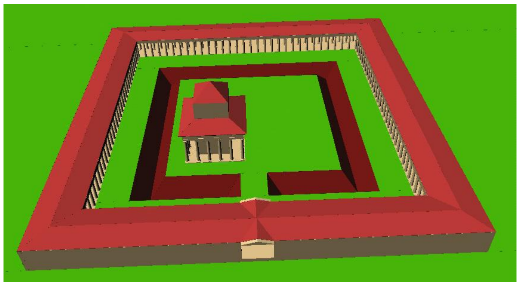

(a) View from above

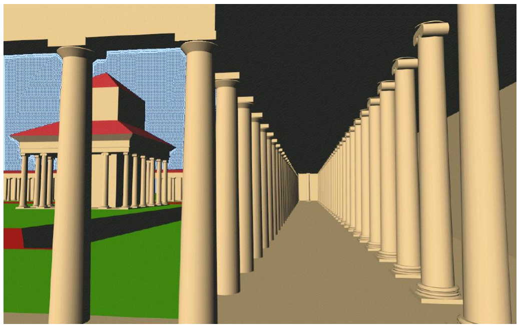

(b) Typical user's view

Figure 2: Views of the 3-D model of the Gosbecks temple complex

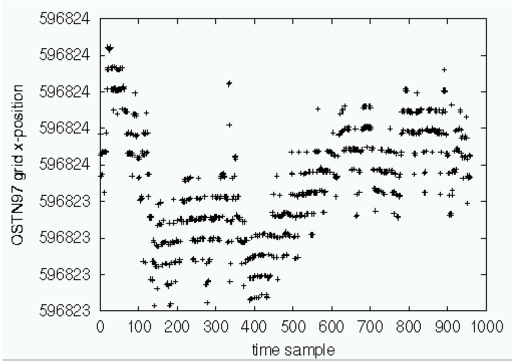

(a) X-position error

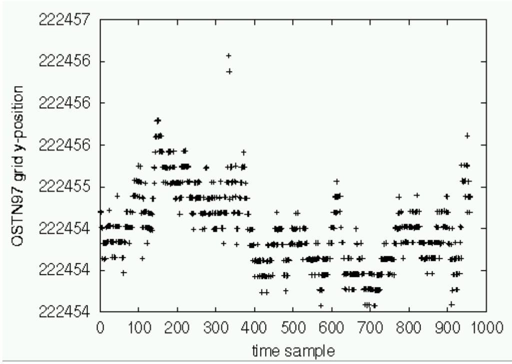

(b) Y-position error

Figure 3: The variation of GPS location with time

In order to make the model compatible with the real world as viewed with the *I-glasses*, the parameters governing the OpenGL camera model must be carefully set. This was done by producing a 3-D model of the inside of the authors' research laboratory, fixing the OpenGL camera location and orientation to be that of the HMD, and adjusting the parameters until the real and virtual worlds show the same perspective effects for each eye independently.

Stereo rendering of the scene is obviously desirable; but OpenGL does not support the stereo mode offered by the *I-glasses*, in which alternate video lines are directed to the left and right eyes. To achieve this, OpenGL is operated with a virtual stereo camera, rendering into an off-screen buffer, and the resulting two views interlaced. While this worked with the standard X server (XFree86 version 3), rendering performance was unexpectedly poor. Examination of the flow of execution through the X server showed that it was itself largely responsible for the performance penalty, so the X server code was modified to improve throughput. A speed-up of almost a factor of ten was obtained, allowing us to obtain stereo rendering of a reasonably complex model at about 15 frames/sec on the 266 MHz processor.

5.3 Interaction mechanisms

User interaction via a Twiddler is adequate for debugging but definitely not appropriate for a production system. We are experimenting with an interaction mechanism that employs only two buttons: to select something, for example a commentary on the nature of the column capitals, the user looks at a capital, observes a change in its colour, and presses the `select` button to invoke it. The other button then acts as a `stop` button if the user wants to interrupt the commentary. With the belt-mounted configuration, these may be mounted on the enclosure; and with the system integrated into a jacket, they will eventually be attached to the front of the jacket.

6 DISUSSION

As the mobile rendering performance we can currently achieve is less than a production system requires, we are investigating two mechanisms for improving it, both of which have been widely used elsewhere. Firstly, as the user is moving, the model can be rendered as a wire-frame rather than a shaded surface. This gives the user an idea of how the geometry is changing with a reasonable update rate as he or she moves. The second mechanism is to employ levels of detail processing: when the angular range subtended by the far side of the portico is small, its rows of columns are replaced by planes with a single image of the columns texture-mapped onto them — of course, the shading is then not strictly correct.

One of the problems with using GPS is that the location changes with time. Figure 3 shows the variation of location of a fixed receiver operating in a differential scheme over a period of time. The X and Y values shown on the graph are computed via the Ordnance Survey's OSTN97 transformation from the OSG-specified Mercator projection of latitude and longitude (to $1/10000$ min of arc) as delivered from the receiver working on the WGS84 datum. A typical differential GPS reading will be within a metre of the right place, but it can wander violently by this sort of amount in the short term. The authors speculate that this is at least partly due to cycle slips and the rising and setting of GPS satellites; whatever the reason, the magnitude of the variations can be problematic in practice. The effect can be ameliorated a little by maintaining a moving average of position (and orientation) values — but this also makes the system less responsive to the wearer's movements.

7 CONCLUSIONS

We are attempting to develop the concept of wearable augmented reality, proven in principle in several research prototypes, into a practical proposition. A "Tin Lizzy" system can be made (just) light enough to be worn on a belt but integration into a jacket appears to be a better solution. Extrapolating the performance of our 266 MHz system and performing experiments on laptop systems suggests that PC/104 cards now becoming available will be fast enough to achieve real-time rendering, though careful optimization of the model will probably still be necessary.

More worrying is the poor accuracy obtained in practice from GPS. Indeed, we believe that a better solutions for position estimation are required for this type of application. In the future, we intend to investigate other RF-based position estimation schemes. However, even the best current outdoors position and orientation sensors will not be accurate enough to superimpose 3-D reconstructions onto existing archæology. In an attempt to achieve that, we are starting to investigate the use of a camera mounted on the HMD, using computer vision processing of its imagery to identify existing buildings, and to infer the information through infra-red motion.

References

[ARCHEOGUIDE, 2002] Archeoguide, 2002. `http://archeoguide.intranet.gr/`.

[CHEVERST *et al.*, 2000] Keith Cheverst, Nigel Davies, Keith Mitchell, Adrian Friday, and Christos Efstratiou. Developing a context-aware electronic tourist guide: Some issues and experiences. In *Computer-Human Interaction 2000*. IEEE Computer Society, 2000.

[FEINER *et al.*, 1997] Steven Feiner, Blair MacIntyre, and Tobias Höllerer. A touring machine: Prototyping 3d mobile augmented reality systems for exploring the urban environment. In *The First International Symposium on Wearable Computers*. IEEE Computer Society, 1997.

[HANDYKEY CORPORATION, 2002] Handykey Corporation. Twiddler, 2002. `http://www.handykey.com/`.

[HÖLLERER AND PAVLIK, 1999] Tobias Höllerer and John Pavlik. Situated documentaries: Embedding multimedia presentations in the real world. In *The Third International Symposium on Wearable Computers*. IEEE Computer Society, 1999.

[MORGAN, 1967] Morris Hicky Morgan. *Vitruvius, the Ten Books on Architecture*. Dover publications, 1967.

[PIEKARSKI AND THOMAS, 2001] Wayne Piekarski and Bruce H. Thomas. Tinmith-metro: New outdoor techniques for creating city models with an augmented reality wearable computer. In *The Fifth International Symposium on Wearable Computers*. IEEE Computer Society, 2001.

[YANG *et al.*, 1999] Jie Yang, Weiyi Yang, and Matthias Debecke. Smart sight: A tourist assistant system. In *The Third International Symposium on Wearable Computers*. IEEE Computer Society, 1999.

ABOUT THE AUTHORS

The authors may be contacted at *VASE Laboratory, Department of Electronic Systems Engineering, University of Essex, Colchester, UK*. Email: `{pritso,djjohn,cc,alien} @essex.ac.uk`

Panagiotis Ritsos is a research student in the VASE Laboratory researching architectures for wearable computing. He designed and built the wearable computer hardware used in this work, and was responsible for integrating software components into the application.

David J. Johnston is a researcher in the VASE Laboratory who specialises in position estimation for augmented reality. He produced the GPS processing and head-tracker interface software used in this work.

Christine Clark is a researcher in the VASE Laboratory. She has developed software that generates 3D models in a number of formats that conform to the rules of Vitruvius. She was responsible for the Gosbecks model used in this work.

Adrian F. Clark heads the VASE Laboratory. His research interests encompass computer vision, virtual reality and wearable computers. He helped with the software integration phase of the work.