

GENERATING VIRTUAL ENVIRONMENTS TO ALLOW INCREASED ACCESS TO THE BUILT ENVIRONMENT

G. T. Foster, D E N Wenn, W. S. Harwin
University of Reading, UK
F. O'Hart
University of Dublin, IRELAND

ABSTRACT

The problems encountered by individuals with disabilities when accessing large public buildings is described and a solution based on the generation of virtual models of the built environment is proposed. These models are superimposed on a control network infrastructure, currently utilised in intelligent building applications such as lighting, heating and access control. The use of control network architectures facilitates the creation of distributed models that closely mirror both the physical and control properties of the environment. The model of the environment is kept local to the installation which allows the virtual representation of a large building to be decomposed into an interconnecting series of smaller models. This paper describes two methods of interacting with the virtual model, firstly a two dimensional aural representation that can be used as the basis of a portable navigational device. Secondly an augmented reality called DAMOCLES that overlays additional information on a user's normal field of view. The provision of virtual environments offers new possibilities in the man-machine interface so that intuitive access to network based services and control functions can be given to a user.

1. Introduction

The problems of ergonomic access and control are particularly relevant to people with special needs when faced with poor building design and information support. This is highlighted in studies of disabled or elderly people in the community. For instance in the 1991 RNIB survey by Bruce et al (1991) where a comparison is made between the level of independent mobility of young registered blind people in 1991 and the level found by Gray and Todd (1967). It was found that the percentage level of independent mobility had not increased over time in spite of the increase in available mobility training. The difficulties in accessing the built environment are a major factor for this lack of mobility and are common across a broad range of people with special needs.

A novel solution to this difficulty is addressed in the EU funded project ARIADNE "Access, Information and Navigation support in the *Labyrinth* of Large Buildings". ARIADNE is exploiting a new generation of networked intelligent buildings to provide an infrastructure that will enhance accessibility to the built environment. This technology is targeted at several levels of building user. Thus a casual user might use the building information structure for navigation directions. Spoken directions would be provided to a blind or visually impaired user, or to a person who is illiterate. A cognitively impaired individual might get the directions repeated at a regular interval with a variation of syntax on each occasion, and a wheelchair user might follow a succession of flashing signs that were guaranteed to avoid stairs and inaccessible regions of the building. The approach used by ARIADNE is based on new microwave smart card sensor technology that allows the building to sense a user is in a particular place in a building and respond in an appropriate way to that users requirements. The smart card readers are networked around the building and can trigger local environment changes in the building or transmit information directly to the user through hand held devices and sophisticated talking signs. The ARIADNE network offers a powerful information and control resource that can be used by user interfaces with different levels of sophistication (Foster, 1998).

Exploiting networked control technologies as promoted by products such as ARIADNE, it is possible to create a virtual environment that is a mirror of the real environment, a concept promoted by Gelernter (1992). This is possible because the network infrastructure allows the real time passing of information and control commands between the virtual model and the actual environment. This control and information infrastructure ensures that the Virtual Environment remains concurrent with the real world, producing a mirror of the actual environment as can be seen in Figure 1.

The use of a parallel Virtual Environment can offer the user an alternative method of interacting with the building. When interacting with the Real World the user has to physically deal with the environment and receives physical feedback from the building itself. This may be difficult or impossible if for example, the person is in a wheelchair or has a motor impairment. An accurate virtual model of the environment offers an alternative control route as information contained in the model can be presented in a variety of ways which are tailored to the requirements of the particular user. The network infrastructure provides a conduit for control and information to flow to and from the real environment. Providing this infrastructure not only benefits users with special needs but can also be used by regular building users, maintenance personnel, for statistical purposes or even to provide a structure to assist autonomous

mobile robots in the built environment (O'Hart and Foster, 1998).

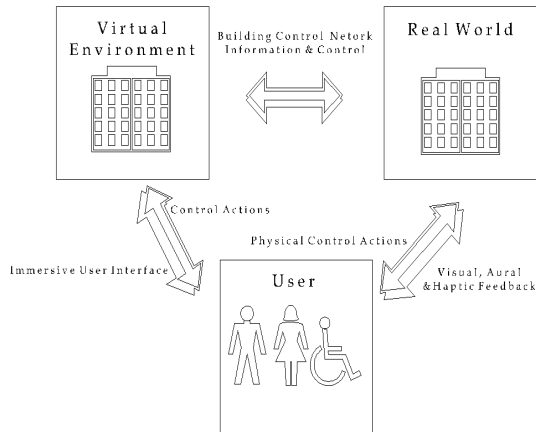


Figure 1: Interaction either with the Real World or a Virtual Environment.

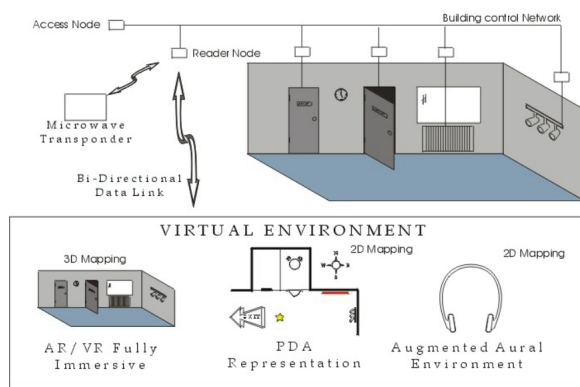


Figure 2: Obtaining a virtual model of the local environment.

The underlying concept is that the model of the environment can be provided by the building. The building is split down into manageable sections, often rooms and corridors, and the model of the local environment is held locally on the network nodes that are physically installed in the building (see Figure 2). A user with a suitable transponder is monitored by the building information system and can extract the local information via a microwave link managed by Reader Nodes distributed through the building. Associated with the Reader Node is an Access Node that holds the virtual representation in a Local Object Model (LOM). The complete building model is decomposed into a series of these Local Object Models allowing easy parallel processing of information with natural fault tolerance. Once a virtual model is available in

the LOM it can be presented in an appropriate form to the building user. Two methods are presented in this paper:

(1) *Two Dimensional Augmented Aural Environments:*

A novel navigational tool is described based on a speech synthesiser that translates the environments into useful spoken facets. The user holds a device that contains a compass and pre-recorded sampled speech segments. This forms the basis of a navigation device for visually impaired building users. The device is also appropriate for people with a cognitive impairment.

(2) *Three Dimensional Mapping using Augmented Reality Environments:*

The second method is a 3Dimensional augmented reality suite called DAMOCLES. The primary interface is a lightweight, low cost Head Mounted Display (HMD) based on the 'I-Glasses' headset from Virtual I/O Ltd. These have onboard tracking and allow the generation of both fully immersive Virtual Environments and semi immersive Augmented Environments that superimpose additional information upon a users normal field of view. The current generation of device is not portable although it could be mounted on an electric wheelchair. DAMOCLES is based on earlier work on an augmented reality that can be used to operate the control system of a smart house (Hammond et al 1996). DAMOCLES is an appropriate technology for a user with severe motor impairments including individuals using a wheelchair.




Figure 3: Model classes within the Local Object Model stored on a network node.

2. Generating Local Models

One of the major problems with generating virtual environments is that models quickly become large and unwieldy, this is a particular problem if the VR information is incorporated into a mobile unit. The Local Object Model

approach combats this difficulty and allows the VR models to be rendered within the memory constraints of the network infrastructure. (A more detailed account of Local Modelling techniques can be seen in Foster, 1998.) Reader and Access nodes are placed strategically at 'decision points' within the building. These decision points are the places where routes diverge forcing a user to make a navigational decision.

The Local Object Model can be seen in Figure 3, and is implemented on each Access Node in the ARIADNE network. Contained within the Local Object Model there are currently 6 model classes which form a distributed data resource across the network. These are:




Figure 4: The Adjacency Model.

- **Localisation Model.** The localisation model contains a physical description of the environment local to the node. Generally this is represented using VRML or a vector list of objects in Cartesian / Polar 2D co-ordinates.
- **Control Object Model.** The Control Object Model contains a list of devices that can be controlled through the network such as window openers, lights and heating.
- **Static Feature Model.** The Static Feature Model contains a list of objects that are not accessed through the network but are of interest to a user, for instance public telephones, seating areas, toilets etc.
- **Dynamic Feature Model.** The Dynamic Feature Model keeps a record of the tagged devices that are currently in the local area. This model can be used to perform a network search to locate a particular person or tagged object.
- **Adjacency Model.** The adjacency model represents the navigation links between neighbouring Localisation Models as can be seen in Figure 4. Connections between areas are given a weighting depending on distance and difficulty and the route must be achieved by passing through a connection feature which is typically a door or archway. A holistic view of the entire network produces an interconnecting lattice of adjacency models that provide a navigational map of the building an example being shown in Figure 5. This lattice can be used as a framework for network wide services such as the parallel searching agents as presented in Foster et al (1997).
- **Performance Model.** The performance model is related to network maintenance and keeps a model of

the performance of the control objects in case of sensor or actuator failure.




Figure 5: A navigational topology lattice generated from connected adjacency models.

3. Interpreting the Virtual Model

Users wandering around a building have widely differing needs dependant upon their experience, ability and special needs. Indeed many users require no assistance at all when dealing with the built environment. However when users are presented with a new environment, especially if they have special needs, then assistance needs to be given to permit them to navigate, interrogate and interact with that environment. There is a broad spectrum of stimuli that can be presented to a user but currently within ARIADNE these are typically limited to visual and aural information. Tests have been carried out with olfactory stimuli and there is the possibility of giving haptic information but the meaning of these is generally too far removed from most people's experience to allow them to be of any use without training.

In order to demonstrate the possibilities of a networked navigation aid within a building two case studies are presented here. The first example explains how navigation is supported within the network using a simple hand held device. Whilst the interaction does not primarily involve virtual reality, a virtual model is used to derive mapping information. The information presented by this system would be suitable for an ambulant user with perhaps, visual or cognitive disabilities.

The second case study (DAMOCLES) covers an augmented reality. The information presented is highly visual, but permits complex interaction with the networked devices without relying on any physical contact with the immediate surroundings. This opens the possibility of a user with motor impairments being able to open doors or dial local phones with eye gaze controls or voice commands.

3.1 A NAVIGATION AID USING A 2-DIMENSIONAL AURAL ENVIRONMENT

When a user is presented with a new environment there is often considerable confusion about the best way of accomplishing tasks within that building, such as locating a person, office or other feature. This problem is compounded when the building is complex and the user has special

access needs. The ARIADNE project demonstrates a network of control and information devices in a network around the building. These devices with their logical topology can be used to provide navigational clues to a disorientated user.

To interpret the virtual model from the building a prototype of a hand held navigational device has been developed. It consists of a small microprocessor, currently a Neuron 3150 device, which provides audible output through an ISD 33180 speech chip. This contains up to 180 seconds of analogue sampled audio, which can be split up into an arbitrary number of sample segments. The device also has a basic compass facility through the use of a Vector 2X magnetometer from Precision Navigation, which provides bearings with a resolution of 1° .

The virtual model is downloaded to the device via the microwave link. In this case a two dimensional environment can be compactly represented as a set of n features \mathbf{F} local to the decision point in the building. (1).

$$\mathbf{F} = \{f_0, f_1, f_2, \dots, f_n\} \quad (1)$$

$$f_n = \{i, \omega, \delta, \phi\} \quad (2)$$

Where a feature f_n contains the following information i an id number that is directly associated to a audible message on the hand held device, ω is the target size of the feature, δ is the distance and ϕ the bearing direction to the feature. (2).




Figure 6: Taking a bearing from a decision point to a target feature.

A typical usage scenario can be considered, the user moves around the building until a decision point is reached, at this juncture, being in range of a Reader Node, the hand held device receives the two dimensional local model from the building and alerts the user via a simple buzzer. The user can then place the device into a 'locate' mode via a simple button press. In this mode points the user scans the device around the room, when the bearing from the compass device matches the feature bearing given in the model, that particular feature is announced to the user who then is able to orient themselves with respect to the feature.

Having located the feature the user can then track towards the object by placing the device into a 'follow' mode (Figure 6). The bearing to the object is taken via the

compass device and the user can set off towards the object. To stay on course the user must stay within an error angle θ in order to reach the goal successfully.

If the user continues within that given error angle the destination should be found, however if the user deviates from the bearing simple audible messages are given to nudge the user back on track. A nominal template for this facility can also be seen in Figure 6, where the user is prompted to veer, make a sharp turn or reverse to get back on course. The effects of any errors in following the desired bearing are cumulative so successive errors can place the user in an unknown position until another decision point is reached. However assuming the user is able to keep more or less on track, it is possible to navigate through the building by travelling from decision point to decision point.

There are some drawbacks to this approach, the major problem is shared with many VR applications and that is tracking the user. The broad area covered by the microwave readers locates the user to an accuracy of roughly 10 metres which adds error to the initial starting point for the following mode. This problem can be addressed by using other communication media such as a directed IR link from the ceiling which can be used to improve the location accuracy to approximately a 2 metre radius before initiating the following mode.




Figure 7: A surface depicting the bearing error across a 10m x 10m room space.

Using bearings to navigate means that the system is sensitive to positional errors, considering a 10m \times 10m space with a user origin at (10,5) and a target at (0,5) the surface in Figure 7 shows the extent of bearing errors related to position. There is a zero error directly along the bearing between the origin and the location of the target feature. Imagining that the target bearing is to go directly north, when the user is far away from the destination feature the bearing error is tolerable. However this error quickly increases the closer the user is to the destination.

The other drawback with the device as it stands is that the compass is affected by magnetic fields and soft iron within the building, this can mean that magnetometer

readings are distorted by metalwork within the building structure, furniture and also magnetic fields from VDUs. The Vector 2X has a built-in hard-iron calibration algorithm. This compensates for magnetic fields generated by a host system such as the housing and electronics. Unfortunately it is extremely difficult to compensate for dynamic changes in the magnetic field strength. In practice the device is intended to be used in relatively open areas away from interference, these errors tend to drop off with distance due to the inverse square law. Another problem with the magnetometer is that it must be kept perfectly level for good results. In order measure the X and Y components of the magnetic field to accurately, the compass needs to be aligned parallel to the surface of the Earth. Typically this is achieved by using a gimbal mount or placing the compass in a bubble of oil allowing gravity to level the compass. Tilting the device gives increasing bearing error per degree of tilt and, due to the Earth's magnetic field, this is an effect that is more pronounced the nearer to the earth's magnetic pole the device is used.

3.2 AN IMMERSIVE 3D AUGMENTED REALITY ENVIRONMENT - DAMOCLES

Some building users require detailed information about the buildings infrastructure and services. Maintenance engineers require plans of electrical cabling, pipework, building operators require statistical information on usage and security guards need to be able to interact with access control systems. Exactly the same network and devices that provide this information can be used to supply advanced interaction methods for disabled users. In particular wheelchair bound users can benefit from being able to interact with the environment, examples would include opening windows, setting heating or controlling taps.

Another method of representing the virtual models contained on the network is through an immersive 3D environment where the user wears a Head Mounted Display (HMD) and stereoscopic information is produced. Augmented reality overlays image information on the normal field of view of the user much like a head up display. The augmented reality suite presented in this paper is DAMOCLES (so called because the original image used to test the HMD was a large sword which appeared to float above the head of the user).

Augmented reality falls between the two extremes of complete VR and full reality and depending on the system in question there exists a continuum of possible degrees of

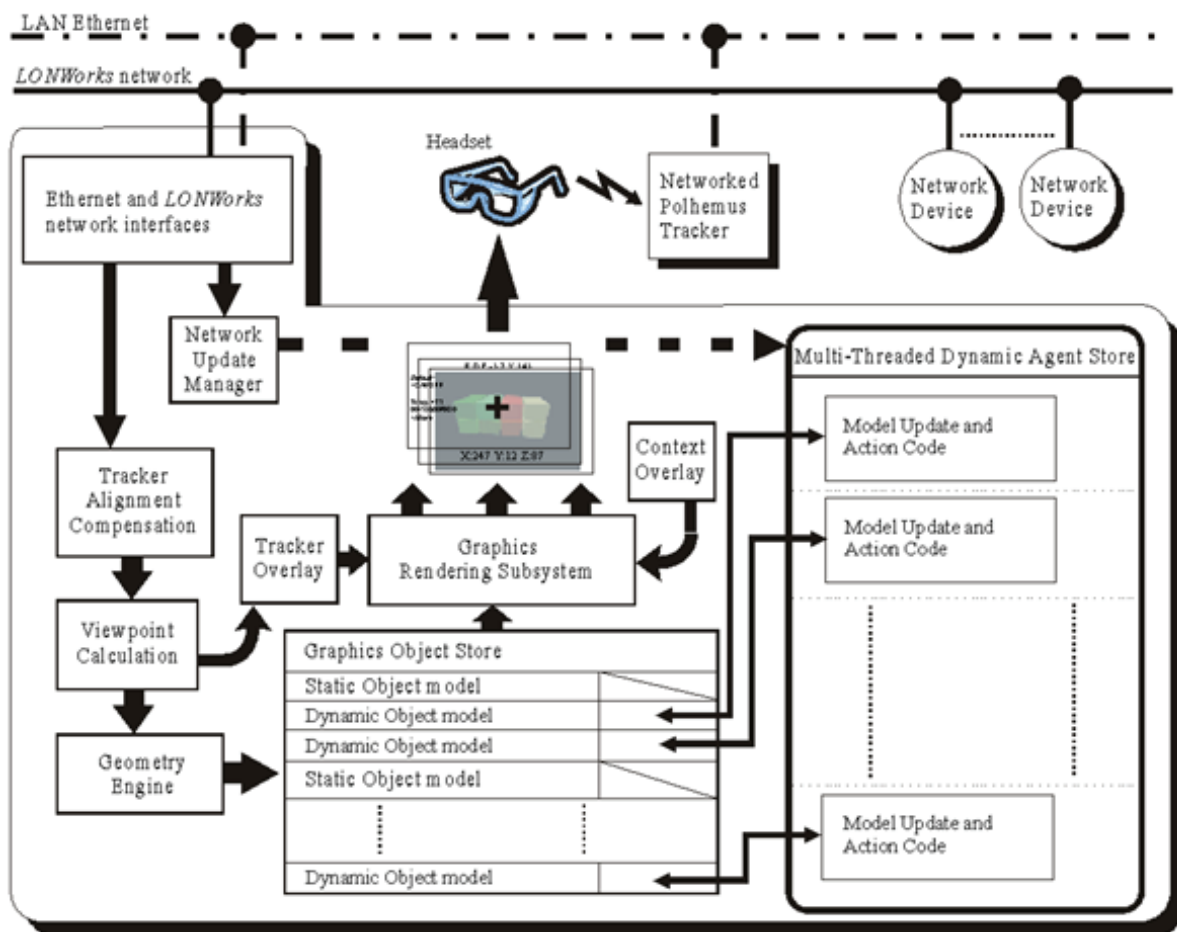


Figure 8: A block diagram of the DAMOCLES system.

augmentation that can be offered to the user as discussed by Milgram & Kishino (1994). This ranges from simple wireframe overlays to full colour enhanced images to present the abstracted display of measured values. A human user is certainly more comfortable with data that only slightly modifies his view of the environment, this approach eliminates the nausea effects often associated with fully immersive environments.

The aim of the DAMOCLES system is to permit a user to be physically present within a real world environment and to obtain enhanced interfaces to a variety of different objects within that world. This is achieved by the use of a stereo-graphic HMD with half silvered optics allowing the user to view overlaid computer generated imagery on the real surroundings. The ultimate goal is to allow a user to freely wander around a built environment in a similar way to the *Touring Machine* proposed by Feiner et al (1997). However unlike the unidirectional touring machine system DAMOCLES has the requirement for detailed information about devices in the network and the ability to interact with them.

A block diagram of the system is shown in Figure 8. At any point in the network a client computer may be connected and at the same location a magnetic tracking system is sited, the Polhemus tracker in this case. There are considerable problems associated with using a magnetic tracker in most locations, not least of which is its spatial calibration to ensure good static accuracy as discussed by Azuma (1995). Most accurate magnetic tracker systems rely on a generated field from a base transmitter. Since the field strength is quite small the useable range of these devices is only a few metres. In the case of the Polhemus system the user must be within 1-2m of the base transmitter (the same is true for other magnetic systems such as Ascension's 'Flock of Birds'). The *Touring Machine* of Feiner et al (1997) used differential GPS along with a small electromagnetic compass (similar to the Vector device) to provide its positional information. However GPS is restricted to outdoor locations if it is not to suffer from inaccuracy due to occlusion of the satellites, and only provides an accuracy of $\pm 30\text{cm}$; far too large for good feature alignment in an augmented reality. The UNC at Chapel Hill have developed a wide area tracker called

HiBall that uses optical means to calculate position but this relies on having key marks on the ceiling (Welch & Bishop, 1997). The reliance on a base station prevents wide area coverage by magnetic systems however they have been used here to prove the validity of the architecture.

Once connected to the network the remote client obtains information from the network about the devices in its location. Unlike the 2D mapping system DAMOCLES uses a more centralised client server architecture rather than a truly distributed approach so the initial transactions are with the central database server. It would be more efficient if all of the location information could be obtained without referring to the database server. However this would require that the users computer know where it was within the network, and since one of the features of the network is the location transparency of a device, it is simpler to refer to a main database for the initial datum at present. Once the virtual environment is generated it is then displayed upon the HMD with due regard to position and orientation.

Given localisation information from a server DAMOCLES proceeds to build a graphical model description. The system first takes graphical data from a device constructing an object within the virtual part of the AR. Attached to each device is a description of the dynamic behaviour of that object. This takes the form of code modules or agents that act upon the model description within the AR. If the construction of the object within the virtual world is complicated because of regular changes to the device, then instead the code module contains all the information required to construct a simulacrum ('or shadowy likeness') of the real device, these two scenarios are shown in Figure 9.

Both OpenGL and DirectX rendering libraries have been used, currently the latter has been selected due to performance considerations on the PCs and the option to easily implement high speed stereoscopic rendering. The choice of rendering system has an influence upon the type of graphical models that can be stored. If fast and efficient model transfer is to take place, size and complexity is dependant on the graphics library being used. Since the models are to be stored on the devices themselves then they must be as compact as possible whilst still conveying the required detail.

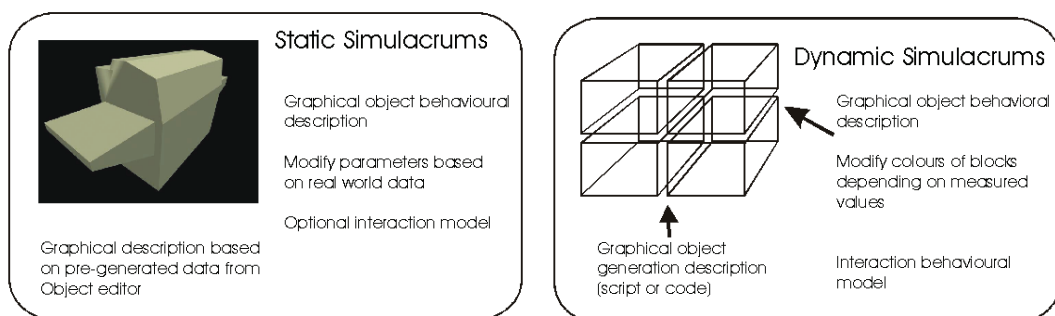


Figure 9: Graphical object model derivation

After the localised world has been generated from the Local Object Models in each of the nodes in the area interaction can then occur between the user and any devices that support external stimulation via the Local Object Model. Two forms of object update relevant for display purposes may occur. Firstly, changes in the physical world sensed and supplied by a device will entail changes in the graphical object within the virtual part of the augmented reality, for instance doors opening will require angular transformations of their simulacra. Secondly, multimodal data sets require abstracted information to be generated and displayed. Depending on the type of the measured parameter this can be quite complex. For example displaying a colour representation of black body temperature requires translation into a suitable range and then mapping to some colour. In both cases the modification to the graphical object is handled by a software agent running in a separate thread within the virtual world. The source of this code is either a system database or from the device itself. By breaking the overall system down into compact discrete objects the size of model update agents is small, and the virtual environment does not have to deal with a flood of information trying to update objects outside the users immediate vicinity.

When using an augmented environment any device in the network that has an actuator attached to it can be controlled, this introduces a number of issues. Safety must be considered, particularly if the users interactions could harm other people or cause the controlled object to go out of limits. Concurrency problems occur when the users requests conflict with the control algorithms already operating within the device. Concurrency is also an issue when multiple users may be interacting with the model and the actual device leading to issues of prioritising control, an example is one user requesting a window to open whilst another requests it closed. At the present time restrictions have not been imposed upon what a user of DAMOCLES can do to the controlled environment but there is provision within the existing agent control framework to add priority management and process limiting features.

DAMOCLES was designed to operate with a diverse range of networked devices. Consequently an interaction method that is consistent for all objects has been implemented by using a context sensitive linguistic paradigm. A method whereby each dynamic code agent implements a default verb specific to that device. For instance an electric window opener has a default reflective action that is to open or to close. Highlighting a selected window object within the users field of view selects that device, a button press then invokes the default verb for the object, and in this case the window opens. More complex verb sets are accommodated in the software by permitting a selection to linger on the object, after a delay a list of available actions is presented to the user via the HMD. A consistent approach has been presented to the user that is easily understood whilst allowing flexibility for more complex functions to be implemented.




Figure 10: An AR snapshot of heated zones in an oven.



Figure 11: Wearing the DAMOCLES HMD.

A snapshot of the augmented reality system can be seen in Figure 10, here the system is installed in the Distributed Systems Research Lab. Looking at the network controlled machinery in the lab displays the state of that equipment. Figure 10 showing an oven which has various heated zones with temperatures displayed in a suitable colour. In this case the third zone from the left is actually coloured red, indicating a high temperature being detected by the networked node controlling the device. Note some static registration errors on the wireframe overlay, a result of tracker calibration.

Figure 11 shows a picture of the DAMOCLES HMD in action, the Polhemus tracker is visible mounted above the head of the user. This picture has actually been enhanced with DAMOCLES highlighting of control features, in this case the wooden framed glass door in the background is actually highlighted as a controllable network object.

4. Conclusions

This paper has presented the use of networked building control systems as an infrastructure to support distributed virtual environments. Two methods of interacting with these distributed models have been discussed, the first being a prototype of a hand held navigational device which could be

used by a person with visual impairments as an orientation and navigation device. This device shows considerable potential although there are some basic problems associated with the use of a compass within a building.

The second system presented is the DAMOCLES augmented reality system that offers a novel method of controlling devices within the built environment. The system augments the normal field of view of a user by superimposing additional information derived from the network, and controllable objects can be selected and activated simply by the act of looking at them. Presently the system is a laboratory prototype and considerable development is required before it could be a truly portable device.

5. Acknowledgements

This work was part funded by the European Union under the TIDE (Telematics Initiative for Disabled and Elderly) initiative of framework IV.

REFERENCES

- R.T. Azuma (1995) "Predictive Tracking for Augmented Reality", PhD dissertation UNC Chapel Hill Dept. of Computer Science technical report TR95-007.
- I Bruce, A McKennell and E Walker (1991) *Blind and partially sighted adults in Britain: the RNIB survey*, Volume 1, Published by HMSO
- S. Feiner, B. MacIntyre & T. Hollerer, (1997) "A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment", *Proc. ISWC '97*, Cambridge, MA, pp. 74-81.
- G.T. Foster & J.P.N. Glover. (1997). Supporting Navigation Services Within Intelligent Buildings. *Proceedings of the First Mobinet Symposium*. Athens.
- G.T. Foster, D.E.N. Wenn & J.P.N. Glover. (1998) ARIADNE - Exploiting a new generation of Intelligent Buildings. In *Improving the Quality of Life For The European Citizen. Technology For Inclusive Design and Equality*. pp 388-391. Porrero & Ballabio (eds.) pub. IOS Press 1998.
- G.T. Foster (1998) "Local Modelling Techniques for the Management of Distributed Control Systems." PhD thesis. University of Reading.
- D. Gelernter (1992) "Mirror Worlds" pub. Oxford University Press. ISBN 0-19-507906.
- Gray & Todd. (1967) *Mobility and Reading Habits of the Blind*, Published by HMSO
- J.C. Hammond, P.M. Sharkey & G.T. Foster (1996) "Integrating Augmented Reality With Home Systems" *Proc. 1st European Conference on Disability, Virtual Reality and Associated Technologies*. ECDVRAT, Maidenhead, UK, 1996.
- P. Milgram & F. Kishino (1994) "A taxonomy of mixed reality visual displays", *IEICE (Institute of Electronics, Information and Communication Engineers) Transactions on Information and Systems*, Special issue on Networked Reality, December.
- F.J. O'Hart & G.T. Foster (1998) "Creating dynamic mobile robot environments from an intelligent building" *Proc. Third European Robotics and Intelligent Systems and Control Conference EURISCON*. Athens, June 1998.
- G Welch & G Bishop (1997) "SCAAT: Incremental Tracking with Incomplete Information," *SIGGRAPH 97 Conference Proceedings, Annual Conference Series*. ACM SIGGRAPH. Los Angeles, CA. August 1997

BIOGRAPHIES

Grant Foster received a BSc in Cybernetics and Control Engineering from the University of Reading in 1991. Since then he has been working on distributed microprocessor applications ranging from industrial control networks to domotics and intelligent buildings. He is currently the project manager for the EU TIDE ARIADNE project which is investigating the use of technologies in improving the accessibility of the built environment. His main research interests cover the application of distributed control systems from low level sensor and controller validation to high level information interfaces.

Contact information:

Grant Foster
Distributed Systems Research Group
Department of Cybernetics
University of Reading
Whiteknights, Reading, UK.
Email: grantfoster@cyber.rdg.ac.uk

Darren Wenn graduated from the University of London, Royal Holloway and Bedford New College in 1989. He is currently completing his thesis in Distributed Systems at the University of Reading. Research interests include integration of cheap distributed microprocessors in robot systems, interaction methods within large scale distributed systems and augmented reality. He is responsible for the design and construction of most of the networked nodes in use at the departments distributed system resource network and is responsible for the development and writing of all the Damocles software.

Contact information:

Darren Wenn
Distributed Systems Research Group
Department of Cybernetics
University of Reading
Whiteknights, Reading, UK.
Email: cubdenw@cyber.rdg.ac.uk

Fergal O'Hart received his BSc in Computer Science from University College Dublin in 1996. He has worked as

a research assistant for the Department of Cybernetics at the University of Reading within the EU-sponsored TMR Mobinet project. He is currently pursuing a PhD in mobile service robotics at Trinity College Dublin. His research interests are mobile robotics and intelligent building systems.

Contact information:

Fergal O'Hart
University of Dublin
Trinity College
Dublin, IRELAND
Email: Fergal.Ohart@cs.tcd.ie

William Harwin received a BA in engineering from Cambridge University in 1982 and an MSc in Bioengineering from Strathclyde University in the following year. Following a period as a research assistant at the Engineering Department of Cambridge University he began a PhD on computer recognition of the head gestures made by people with cerebral palsy, work he completed in 1991. Between 1990 and 1995 he directed rehabilitation robotics research at the Applied Science and Engineering Laboratories in Delaware, USA. He moved to the Department of Cybernetics at the University of Reading, England in 1996 where he directs work at the Human-Robot Interface Laboratory (tHRIL). Research interests focus on unusual human-machine interfaces include haptic interfaces, robot assistance of human movements, and design research to improve the availability and quality of products for niche assistive technology markets.

Contact information:

William Harwin
Distributed Systems Research Group
Department of Cybernetics
University of Reading
Whiteknights, Reading, UK.