Integrating augmented reality with home systems

J.C. Hammond, P. M. Sharkey, G.T. Foster

Department of Cybernetics, University of Reading
Whiteknights, Reading, RG6 6AY, UK

J.Hammond@HSOnline.co.uk, P.M.Sharkey@reading.ac.uk, cybgtf@cyber.rdg.ac.uk

ABSTRACT

Augmented Reality systems overlay computer generated information onto a user's natural senses. Where this additional information is visual, the information is overlaid on the user's natural visual field of view through a head mounted (or "head-up") display device. Integrated Home Systems provides a network that links every electrical device in the home which provides to a user both control and data transparency across the network. This paper discusses the integration of AR techniques with a real world infrastructure and demonstrates this through the integration of an visual AR prototype with an existing Home Systems technology test rig. The integration of these two technologies provides the basis of a specialised information/control device that allows the control and diagnostics of Home Systems devices from the basic action of looking at them.[1]

Keywords: augmented reality, home systems, control devices

1. INTRODUCTION

Virtual Reality has been the subject of intense development over recent years with much of the technology being driven by the military and entertainment's (computer games) industries. As a result of advances in headset technology, associated model and graphics handling techniques are now available that are opening up new application areas. Augmented Reality (AR) is one such area. Augmented reality has been previously investigated in the military and automobile industries, primarily through the development of the Head-Up Display (HUD), where visual information regarding the current dynamic state of the vehicle (e.g. and aeroplane, helicopter or car) is displayed to the pilot by projection of this data directly onto the windscreen of the vehicle, or (more expensively) onto the visor of a pilot's helmet. The principal advantage this type of display offers is the provision to the pilot of essential information (airspeed, altitude, attitude, etc.), whilst allowing the pilot to maintain visual coverage of the airspace through which he is flying. The European Commission research project *Prometheus* has taken this concept one stage further, by investigating how a computer generated scene may be overlaid unto the real scene to enhance the user's performance. Specifically, by instrumenting a car with cameras, other sensors and real-time image processing computers, the scene in front of a driver may be enhanced by overlaying outlines of the road and any obstacles directly onto the windscreen, even when such obstacles are obscured by fog or in night driving conditions.

This paper presents an adaptation of this concept with the specific application area of disability, and details the development of a prototype visual AR system. The central idea is to provide a person with (non-visual) disability additional visual information that may assist them on a day-to-day basis. Thus, a person with mobility difficulty may be able to look about an environment and, though the additional information provided by the AR display, ascertain what objects within their field of view may be influenced. By coupling such an information provider to some form of control one can easily envisage the user visually selecting a functional object and controlling the object by simply looking at it. Home Systems provides the ability to associate function with common objects in a typical home, such as doors, windows, lights, alarm and safety systems, and so on. We show: how a simple, inexpensive, head mounted display may be adapted to allow a user to view a typical room, where the 3D structure of the room is overlaid in wire-frame onto the actual scene; how this overlay can update in real-time as the user looks around; how functional items within the room can be identified and highlighted through the display as the user looks at them; and, finally, how the function may be accessed in a simple intuitive way.

[1] This project was developed as a final year project by the first author for the degree of BSc in Cybernetics and Control Engineering, from the Universit of Reading, UK

 ISBN 0 7047 1140 X

1.1. *Augmented Reality*

A visual AR system overlays three-dimensional computer images over a user's normal field of view through the use of a 3D head-up display. AR is different from VR because the model world that it creates is not an abstracted 'virtual' world but a close 'mirror' of the actual environment. The ideas of modelling closely correlated mirror worlds can be seen in Gelernter (1992). Virtual Reality systems are currently at a state where the computing power needed to generate a believable and useful environment is still not widely available. However, the graphical display requirements for AR systems are reduced since the amount of computer graphics can be relatively small compared to the field of view but still maintain useful information.

The close correlation with actuality makes the AR system particularly useful for real world Human Computer Interface (HCI). In order for any AR system to operate effectively it must have complete information about the environment in which it is operating. Further, to make the system a useful tool it must be able to alter the state of an object in actuality and reflect that change in the computer model. Hence a method is required to enable measurement, control and feedback between the computer model and the real world.

1.2. *Home Systems*

For the domestic environment a technological infrastructure capable of supporting an AR system exists in the Home Systems (HS) control network. HS provides a network that links every electrical device in the home, this provides both control and data transparency across the network. Integration of AR with HS opens up the possibility of modelling the domestic environment and providing a novel method for interaction within the home. A multitude of scenarios are possible, from opening a window simply by looking at it, to changing the heating control system from examining a radiator. The implications of such technologies for rehabilitation can be found in Cooper (1996). HS and AR are therefore mutually supportive technologies.

- Information infrastructure.
- Actual devices to control.
- Link to "actuality".
- Network model.

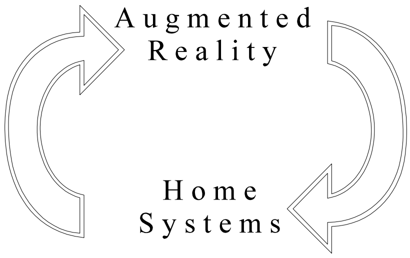

- Human Computer Interface.
- Extends HS location model.
- Novel alternate control method.

1.3. *Structure of Paper*

The paper is structured as follows. In Section 2 we review Home Systems and introduce a number of projects associated with this work. Section 3 discusses the hardware development of the Augmented Reality System and prototype head-mounted display, while Section 4 presents an integration of the Augmented Reality System within a Home Systems test rig set up within the HS Research Laboratory. This report concludes with a wider discussion on how such a prototype AR system may be further developed.

2. HOME SYSTEMS TECHNOLOGY

2.1. *Home Systems Overview*

The Home Systems (HS) specification, which defines networked control systems for the domestic environment, was developed by the ESPRIT Home Systems Consortium (Project 2431) and the integrated Interactive Home (Project 5448). The specification is upheld by the European Home Systems Association (EHSA). HS supports networked communications on a variety of media including twisted pair and power line. Nodes are based upon the common 8051 family of microprocessors and the HS system is an 8 bit addressed token passing network.

In HS terms, devices connected to the network fall into the following two categories:

- *Complex Device* (CiD): The CiD also has no system 'intelligence', however it does support token passing and is therefore able to broadcast data onto the network. An example CiD may be a door locking device that signals back onto the network whether or not it has successfully completed its task.
- *Feature Controller* (FC). A feature controller is a device that contains 'system intelligence', i.e. in the application layer of the device it is able to process and act upon information from network devices. The nature of the system information processing capabilities of a feature controller is entirely up to the designer. An

Proc. 1st Euro. Conf. Disability, Virtual Reality & Assoc. Tech., Maidenhead, UK, 1996
©1996 ECDVRAT and University of Reading, UK ; ISBN 0 7047 1140 X

example FC might be a security system feature controller. This controller is able to gather information from sensors and issue controlling actions. Also it would be able to communicate with other feature controllers where appropriate. For instance it might alert a communications feature controller to phone for help in an emergency situation.

2.2. *The HS-ADEPT Project*

The HS-ADEPT Project (Home Systems - Access of Disabled and Elderly People to this Technology), funded by the European Commission TIDE Initiative, is co-ordinated by the Department of Cybernetics, University of Reading, see Ferreira et al (1995). The aim of the project has been to develop emerging Home Systems technologies to provide the technological infrastructure to allow independent living for disabled and elderly people. The work at Reading has concentrated on the system integration and control issues relating to networked technologies. Project developments are currently being installed in the homes of 3 disabled people in the UK and a rehabilitation centre in Portugal, where they will be subject to independent real world trials.

2.3. *User Interface Feature Controller*

2.3.1. Structure and Model. The major problem when attempting to control any networked system is defining the physical location of nodes around a network. The problem manifests itself when trying to control a specific device across the network as each similar device actually looks the same. So choosing which particular light you want to turn on from a large list presents several control problems. This problem is made more acute with concepts such as 'plug and play' where devices can actively be moved around a home. The solution adopted by the HS-ADEPT project was the development of a User Interface Feature Controller (UIFC) that holds local information about control facilities available in a certain location.

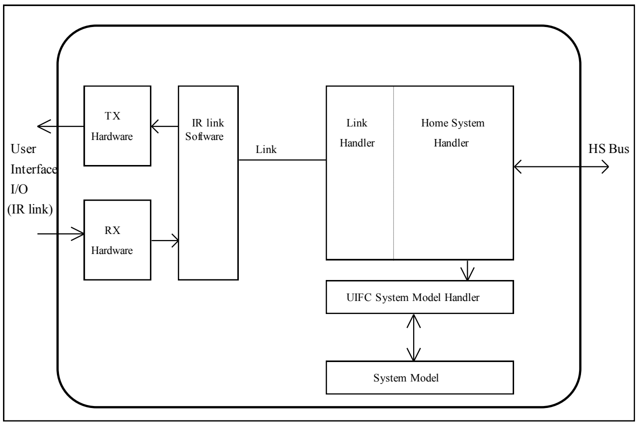

Figure 1. *Schematic of the User Interface Feature Controller*

Fig. 1 shows the schematic of the UIFC. The philosophy behind the UIFC is very simple: it works on the premise that devices are most likely to be operated in the immediate vicinity of the user. The user is tagged to a specific location through an Infra Red (IR) communication link, the relatively short range of IR communications effectively locating the user to the nearest transceiver. In practice, a UIFC is placed in each room of the house so control actions are defined on a room by room basis. The IR communications are handled by a separate microprocessor board and the validated codes are passed to the UIFC node via an RS232 link.

The UIFC maintains a local model of devices that are accessible from that point. The model is initiated by placing the UIFC in a 'Learning Mode' and enrolling the necessary devices. So, for a particular room, the UIFC may learn the local devices such as the lights, doors and windows. However, there is no constraint as to which devices are learnt in this way. So if global control actions are required, such as the security and safety system or the front door unit, they can simply be learnt as well. The System Model Handler then keeps a dynamic model of devices relevant to it's location.

In a typical control scenario the IR link receives a message, for instance to turn on the light. This message is parity checked and, if valid, is then passed on to the UIFC node. The model then attempts to find the local light and operate it. Feedback is passed back through the bi-directional IR link as to whether the control action was successful. The advantages of using the local FC concept are a direct consequence of the UIFC not having to hold a detailed representation of the entire Home System network to which it is attached:

- Software for the UIFC is less complex as the UIFC has less devices to address. This reduces the size of the dynamic model stored, the amount of ROM and RAM space required and the complexity of the software.

- Reduced software complexity produces an increase in execution speed.
- A reduction in the size of the dynamic model means that a control action requested by the user can be implemented faster as there is less time required to search the model for matching network information..
- For the UIFC to maintain an up-to-date dynamic model it must be informed of changes to devices on the HS network. If every UIFC is required to keep track of every device on the HS network this generates a high amount of bus traffic. With every UIFC having to monitor only a subset of the devices on the HS network there is a considerable reduction in bus traffic and, as a consequence, a reduction in the software activity within the UIFC.
- Local control will allow only the activation of devices which are local to the user. The user has visual feedback as confirmation of his action and the safety associated with that feedback.

2.3.2. IR Link. The IR link was designed to give a generic message structure that allows the construction of very cheap and simple User Interface devices. Effectively this means that control devices can be pre-programmed – the command to control a device is the same whichever room you are in – and the UIFC translates the requested action to the appropriate device. The IR link is bi-directional although unidirectional communications are adequate for normal operation – the return message largely ignored. However the return message does allow feedback to the user. For example, the Apple Newton based control devices developed at the University of Reading exhibit several useful features as a result of the bi-directional communications, which include:

- context specific menuing (for instance if the front door bell rings the door entry menu is automatically displayed);
- alarm display; and
- network diagnostics.

The full set of codes for objects can be found in the *Home Systems Specification Command Tables* in *Home Systems Specification Release 1.2*. (1995).

3. THE AUGMENTED REALITY SYSTEM

A visual augmented reality system comprises three major elements: computing, optics, and sensors. The computer stores a reasonably accurate 3D wire-frame model of the environment within which the user is traversing. Head tracking sensors can establish the location of the user within the environment and the direction in which a user is looking (this assumes that the user's eyes look along a 'Cyclopean' or gaze direction). The computer model estimates the field of view of the actual environment and presents this information, via Liquid Crystal Displays (LCDs) back to the user through dedicated optics. The optics (LCD screens, half-silvered mirrors and lenses) ensure that the wire-frame model coincides with the real world view. The hardware implementation of a prototype AR system is detailed below.

3.1. Overview

The AR system is shown schematically in Fig. 2, which details the information transfer between different devices within the system. The AR system is based around an application developed on a PowerMac 8100AV computer which controls the interaction between the head tracker, the graphics display devices and the IR link to the HS network. The primary link of the system is that between the PowerMac and the headset as well as the receiver for the head tracker. The direction of the arrows in Fig. 2 denotes the direction of communications in the current system – the link between the PowerMac and Home Systems is unidirectional, while the communications between the PowerMac and head tracker are bi-directional. For the integration to the Home Systems network (see Section 4) the PowerMac utilises an IR link for serial communications.

The display is produced using a pair of LCD television screens with 512 x 768 pixels. The screens are mounted on the headset in an assemblage of optics consisting of lenses and half-silvered mirrors, thus allowing the images generated by the screens to be overlaid onto the users visual field.

The position and orientation tracking of the headset is performed by a Flock of Birds™ (FOB) Tracker (from Ascension Technology Corporation). This has a range of c. 2.5m from its transmitter. The FOB works by transmitting

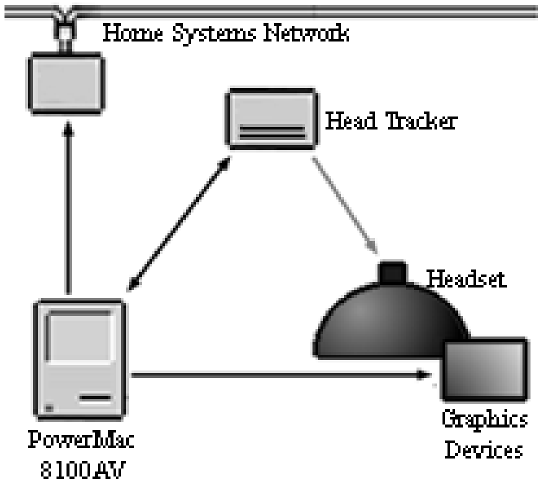

Figure 2. *System communication schematic*

a pulsed DC magnetic field which is received at the headset and all six degrees of freedom are determined by the measured field characteristics. The head tracker is capable of up to 144 measurements a second and has a translational accuracy of approximately 2mm with an angular accuracy of 30 arc-seconds. The sample rate of the FOB exceeds the required update rate for any display device by a factor of 8, and the resolution is approximately equal to that of the human eye when the distances for this application are considered. Due to the geometry of the AR headset prototype, one pixel represents approximately 50 arc-seconds of the user's field of view. As a result, there is little degradation in the quality of the display attributed to the accuracy of the head tracker.

The software on the PowerMac is an object orientated application that performs all calculations on the model of the environment. These allow the graphics to be updated given the position of the user in the environment. The image is generally a wire-frame representation of the environment so that the user is able to clearly see the outside world without obstruction. This also cuts down on the time it takes to render the graphics and therefore the latency of the system with degrading the information conveyed.

3.1. Optical Hardware

The optical hardware consists of half-silvered mirrors and lenses mounted in front of the user's eyes allowing each eye to simultaneously view both the outside world and the images on the LCD displays. The configuration of the assembly can be seen in Fig. 3. This shows one half of the optics on the right side of the headset. For each eye the optical components consist of a half-silvered mirror mounted at 45 degrees that allows the user to see through the headset and to simultaneously see the computer generated images. The images are focused, via a lens, to a suitable distance in front of the user. Without the lenses the user would not be able to focus on both the graphics and the environment at the same but would have to continually re-focus between the two. The main problems to be solved for the optical hardware were:

- *Graphics field of view:* the field of view should be as large as possible without degrading the quality of the graphics.
- *Brightness of the virtual images:* the brightness should be sufficient to allow the user to clearly see the virtual images under normal lighting or include the provision to alter the brightness and contrast to suit the conditions.
- *Focusing of the virtual images:* the images should be set at a comfortable distance in front of the user, probably between 3-5 metres.
- *Alignment of the virtual images:* the images should be aligned accurately enough to allow the user to easily associate the virtual images with their physical counterparts. They should also be accurate enough to prevent incorrect interaction with adjacent devices, for example, two windows physically close in an environment should have separate images with no overlap, thus preventing one window to be activated when the intention of the user was to activate the other. This covers physical alignment of the optics as well as factors affecting distortions in the images, for example aberrations caused by the lenses.

The half-silvered mirrors were chosen to have a 50:50 reflect-transmit ratio so that they give equal luminance to the virtual and real images. Any further adjustments are made by the brightness and contrast of the LCD displays. The mirrors are mounted 30mm from the users eyes covering the entire visual field not obstructed by the assemblage. This means that the real world is seen with uniform clarity .

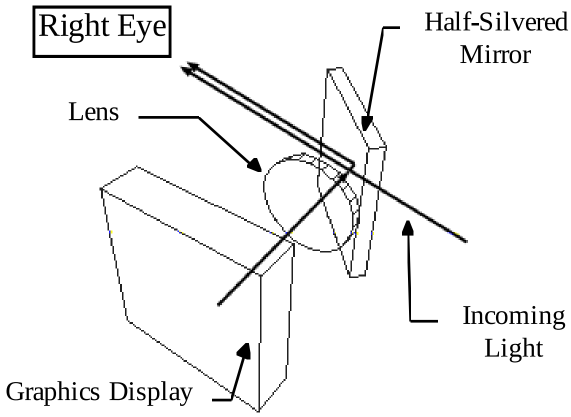

Figure 3. *Basic Optical Configuration (Right Hand Side)*

Each lens is mounted 40 mm from the mirrors and perpendicular to the user's line of sight. The lenses are plano-convex with diameters of 50mm and a focal lengths 125mm. Hence graphics screens placed at a distance equal to the focal length appear to be at infinity, allowing the user to see these combined with the real world at anything beyond approximately 3m. The size of the lenses mean the virtual image covers 36.32 degrees of the users visual field. The eyes normal field of view being approximately ±100° horizontally and ±60° vertically but most of the activity the eye concentrates on is focused on the fovea. This combined with the fact that most of the images displayed will probably be at distances of 3m or more means 36.32 degrees is sufficient for the application. Because each lens is less than 73 mm from the eyes no light is lost from the graphics displays and when viewing the images the effects of aberrations are minimal. These do exist but only at the extremes of the images away from the concentration of the fovea.

3.2. Display Software

The software written specifically for the application is based on the XYZ Fixed Angles transformation given by J. Vince (1995) to position the user in the virtual environment. This uses a single homogeneous equation to rotate and translate the environment before the graphics for each eye are mapped on to a two dimensional virtual plane (one for each eye) given the focal length of the lens. For speed these calculations manipulate integer values wherever possible and therefore do introduce slight errors in the transformation although in vital areas such as the measurement of angles long values are used. The errors noted in this transformation are negligibly small but it is the intention to eradicate these at a later date by scaling up all values to be significantly large integers, including those currently being represented by longs and only return the values to their correct size immediately before they are mapped to the display.

The model of the environment is stored as a series of shapes in a linked list. This list is constant once the model has been defined. A second identical list is used to store the transformed model after every movement of the user. The model is then clipped using a simple algorithm determining all those objects behind the user. More advanced clipping is not required as the graphics are only wire-frame and do not need such common 3D procedures such as back face removal. The simplicity of the operations cuts down on the latency of the system allowing the addition of more specific procedures for object selection.

The most application specific procedure detects whether the user is facing an object. Each object with functionality is assigned one or more facets, defined by a series of four 3D co-ordinates. If the vector projected from the user straight through the centre of the virtual screen intersects this space then the user can activate the object. The procedure works by projecting a vector in the model given the users position and orientation, three points are taken on the plane of the facet and a vector normal to this is calculated. After simple checking to ensure that the two normals are not parallel indicating that the users line of sight never intersects the plane, equation (1) shows how the point of intersection is ascertained.

There are a series of restrictions for the model that make the equations for the positioning of the point of intersection simpler and quicker. The four points of the detection facet must all be in the same plane and the facet must be a rectangle. This is to check whether the point of intersection is within the facet the vector between the point of intersection and each corner of the facet is found and the angles between these and the adjacent side of the facet is measured. If all four are less than 90 degrees then the point of intersection is deemed to be within the facet. These restrictions are not serious disadvantages since this only requires that all devices be represented by a flat rectangle. The distance of the user from the point of intersection is h given by equation (5). The procedure only activates the nearest of the facets indeterminate of which object it is on. In this way only one object can be selected and acted upon. Later additions could be to assign different actions to different facets on the same object, for example enabling a door be locked from one side only.

$$\mathbf{n}_x (u_x - x_p) + \mathbf{n}_y (u_y - y_p) + \mathbf{n}_z (u_z - z_p) = 0 \quad (1)$$

$$u_x = x + h \sin(\phi_{yaw}) \quad (2)$$

$$u_y = y + h \sin(-\phi_{pitch}) \quad (3)$$

$$u_z = z + h \cos(\phi_{yaw}) \quad (4)$$

$$h = \frac{\text{(N.Point-N.Position)}}{\text{(N.Unit)}^{-1}} \quad (5)$$

where

$\mathbf{n}_x$, $\mathbf{n}_y$ and $\mathbf{n}_z$ are vectors normal to the plane,

$\mathbf{u} = [u_x, u_y, u_z]^T$ is the point of intersection,

and $\mathbf{p} = [x_p, y_p, z_p]^T$ is a point on the plane.

These calculations can easily be modified to allow greater versatility by defining the facet by only three points. This constrains the facet to a single plane whereas four points does not. The principle for ascertaining the relative position of the point of intersection is similar but will not work on the 90 degree calculation rather some other specific to triangles. This unfortunately requires more storage area for the facets as well as a slower more complicated algorithm not required for this application.

3.3. Head Tracker and Calibration

The Flock of Birds position and orientation tracker runs directly out of the PowerMac modem port using the RS232 protocol. On issuing specific ASCII characters to the FOB it will return x, y, z, roll, pitch and yaw information of the receiver relative to the transmitter. The returned information is encoded in 12 bytes, 2 for each measurement which have to be decoded by a series of bit transformations as described in the reference documentation of the FOB. Unfortunately since the FOB only measures the position of the receiver relative to the transmitter an accurate position of the transmitter relative to the base co-ordinates of the model must be known.

The acuity of the human eye is relatively good at 40 seconds of an arc (or more simply the eye can distinguish between two points 1.5 mm apart 10m away). This means that the accuracy of the transmitter needs to be as exact as possible so that the alignment of virtual model and real environment is as close as possible. This is a problem known as static registration and is common in AR systems, mostly unseen in VR due to the lack of frame of reference for the user in a completely immersive display. Static and dynamic registration (the appearance of the virtual image to lag behind the real environment when the user is moving) are being researched by a number of groups including Azuma and Bishop (1994).

To place the head tracker accurately without the need to physically measure the position of the transmitter before the model is initially configured a feedback procedure was implemented. This procedure was designed to automatically calculate the orientation and position of the transmitter and in the process take into account the position of the receiver unit relative to the eyes of the user. The procedure also removes any errors in the positioning of the receiver on the top of the headset and should be specific to each user that configures it. The procedure works by prompting the user to align a cross they are able to see in their display to certain points in the environment of known position. The first point is a cross set in front of a mirror and is used to find the initial offset angles of the tracking equipment relative to the co-ordinate frame of the room. By aligning the cross in the users virtual display with the cross in front of the mirror and its reflection it is guaranteed that as long as the mirror is mounted flat on the wall the user will be on a line perpendicular to the wall and therefore the results provided by the tracker are offsets. After aligning with two more crosses with significant x, y and z displacements it is possible to calculate the position of the eyes of the user and therefore the relative position of the tracker transmitter. The model of the environment is configured from these calculations. This procedure provides a satisfactory static registration for the integration to HS and its control.

3.4. Model

The physical model used by the AR system is provided by the device information held in the Home Systems devices. The principle function of the AR system is the accurate display this model of the environment. The current system relies on an installation procedure for the definition of this model and is determined by physically measuring the environment and recording its characteristics in a pre-defined format enabling a dynamic interpretation by the application. Further developments will have such information available directly from the network.

The format for the environment model contains a variety of information, some unique to that model whilst other symbols represent objects contained within the application. The application is written in such a manner that allows any shape to be generated as long as it is definable by a series of straight lines. A shape may simply be for visual reference only but by the inclusion of specific characters the object is no longer assigned the attributes of a simple shape but takes on the functionality of a window, door, light or any of the objects in the applications hierarchy. By doing this a set of co-ordinates in three dimensional space might be allocated the functions of a window, for example with position and security (locked or unlocked) states. The only remaining set of information required is the 'Detection Frame' defined as the area of the object that when looked at, the user may act upon the object.

The information contained within the application for each object is currently constant although at a later stage this paper discusses the increased benefits of a completely dynamic model obtained from sources other than the user. The application stores data in each of the objects in the hierarchy on the possible states the object can be in as well as handles to menu descriptions that ultimately provide the user with their options; handles to audio sound-bites that provide an additional feedback to any changes in the state of an object and the necessary codes to interface at a basic level to the Home Systems network.

4. INTEGRATING AUGMENTED REALITY AND HOME SYSTEMS.

4.1. The Test System

The test system as shown in Fig. 4 has been implemented in the Distributed Systems Research laboratory with effectively a single room scenario. This employs one UIFC that holds the network control model for a variety of devices including a door actuator, window actuator, lights and a security system.

Presently the AR model has been synchronised by manually entering the devices present in the room. The various devices are highlighted by placing a wire frame representation over the object with basic on/off control actions implemented using a simple switch. Proportional commands such as dimmer intensity can be set through the keyboard. Continuing development will enable the physical model to be automatically loaded from the UIFC model into the AR model.

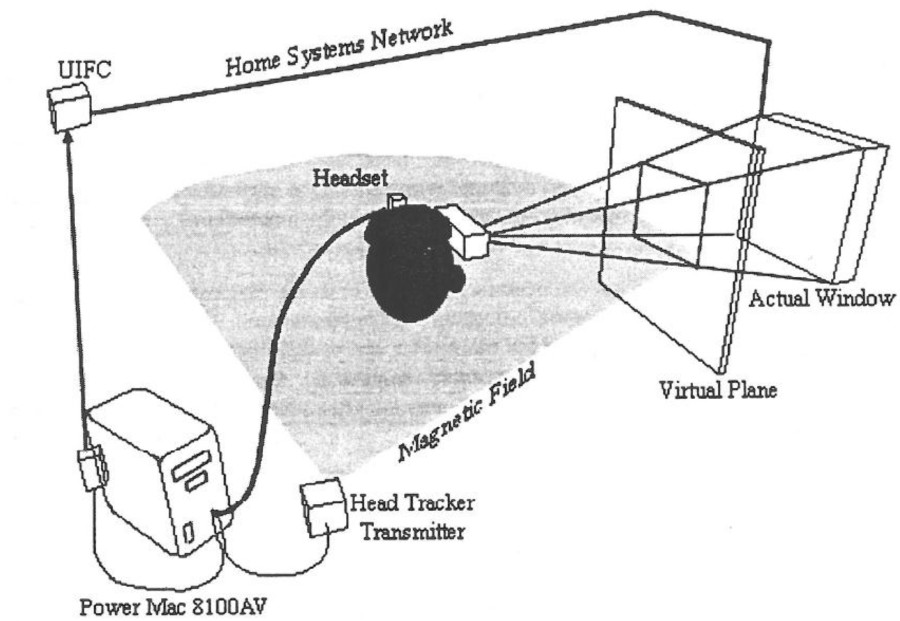

Figure 4. *The test installation*

4.1.1. Accuracy. The system at the moment is not as accurate as was originally intended due to the combination of errors in various separate areas. These are the errors introduced in modelling the environment, inaccuracies in the

Proc. 1st Euro. Conf. Disability, Virtual Reality & Assoc. Tech., Maidenhead, UK, 1996

calibration calculations as well as errors introduced in the transformation of the model. On their own these are negligibly small but combined they introduce registration errors of up to 10 - 12 inches. All these errors can be reduced with significant work but the system does perform the necessary tasks and the registration error can be avoided with intelligent modelling of the environment, for example, not modelling the edges of an object but using a symbol or hot spot overlaid onto the device to indicate functionality. There are also possible errors inherent in the design of the headset due to the proximity of the head tracker receiver to the metal frame of the optics assembly as well as the weight incorporated to balance the weight on the front of the headset. This could be easily be solved by changing the construction materials but the effect of the metal is currently unknown without removing the compounding errors from other parts of the system.

4.1.2. Comfort. The current headset is a prototype and is not in its smallest possible form. This means that the users head and neck are subject to a substantial weight. This makes wearing the headset intolerable for prolonged periods of time and means that the device would not be viable without serious ergonomic design to reduce weight and increase comfort.

4.1.3. Control. The system is designed to only pick the nearest of any overlapping devices and therefore to control any obscured devices the user must physically move so that the obstruction is removed. For the persons this system is intended this may not be a viable option but it would not be too difficult to re-engineer the software to allow, with specific commands, the user to toggle between different devices intersecting their line of sight. The system in its present state utilises the press of a mouse button to input user action and this means that the entire system is tied to the proximity of the PowerMac. With further redesign of this input device right up to speech recognition capabilities provided by the PowerMac this restraint will be lifted.

4.1.4. Display. It was intended that the wearer would be able to easily see both virtual and real images concurrently. In specific environments this is severely compromised due to ambient light almost drowning out the virtual images. Unfortunately the contrast of the current screens is fixed and only the brightness is variable. Shielding the graphics displays does increase the quality of the display and reduce this problem but the graphics still maintain a constant level and whereas looking to a dark corner of the room provides good conditions for observing the images viewing a bright window does not.

4.2. The User Perspective

At the time of writing there have been no extensive trials of the system covering a wide range of users. First impressions suggest that the system has potential applications for people with quadriplegia or other high level paralysis who are without visual impairment and are able to look around a room. However preliminary findings have raised a number of interesting points.

The optical display need not be high resolution. If a method of highlighting devices by ‘hot spots’ is employed a simplified headset could be employed.

The user feedback does not have to be optical. All the AR systems to date rely on an optical representation of the room. However this information could be provided in an audible form so a headset would not have the display devices just a small earpiece. This could be used to help a visually impaired person control devices or orient themselves in an unknown environment.

To initiate control actions most actuation devices are appropriate. Here any form of switch could be used from suck & puff switches to voice actuation. The only unsuitable switches are those that cause a gross head movement such as head switches.

Factors such as head tremor affect the system. Tremor is a major problem for the AR system, filtering of the tracking data along with broad tolerances on device locations goes some way to alleviate the problem.

4.3. Further Work

The power of integrating the technologies of Augmented Reality and Home Systems has been shown to be to the mutual benefit of information sharing, this is still at a limited level and highlights the need for the parallel development of the technologies. In particular the ability to download the model of the environment directly from the UIFC will be a significant advance, this direct feedback link makes true synchronisation of the AR model and the actual environment a reality.

The first stage would simply be to recreate the model from the devices on the network who would supply physical dimensions as well as detection frames and information determining the identity of each device. Increasing the dynamic nature of the objects in the AR system and removing the descriptions of a window, door etc. to the network there is the advantage of being able to slot new and previously undefined devices onto the system that carry their own dimensions, states, activation codes and handles without the need for re-engineering the hierarchy of the AR system.

Surpassing this the AR system has the potential of displaying almost any visual data the HS network can potentially store about itself. The discussion of this paper to this point has assumed the existence of a single model of the environment which could be described as the 'User level'. The network level error detection capabilities allow errors to be pin pointed to specific pieces of hardware. A model of the physical connections and apparatus on the network could be downloaded to the user under the heading to 'Engineer level'. With this model the user would be able to view the state of the network as well as any available statistics from individual components suitably superimposed onto the corresponding device. This would require that the AR system would hold links to various models of the network held within the UIFC that could be download as requested and updated with any change in state of the environment.

5. CONCLUSIONS

2.4. Extensions to specification to enable AR applications

While HS is able to offer the technological infrastructure and control capabilities that enables real world interaction with Augmented Reality systems, there is, as yet, no physical installation information (i.e. no local model) carried by a particular device. As a result, the AR model must be built up from physically measuring up a room. The next steps in the development process is to produce a number of devices that carry their physical dimensions and installation details. This level of detail may then be stored in the model so that when a person wearing the AR headset enters a room the local model downloads all the necessary physical data for the AR system to build its own environmental model.

As a proof of concept the results from this project have clearly shown the advantages of the integration of AR systems with existing control technologies. For commercial systems to be viable there are some technological advances to be made both in display, tracking and network technologies.

- *Users.* The AR system represents a different way on interacting with the environment, using the system as a control device has considerable potential for a broad range of users. The effectiveness of the dynamic AR model and display depends upon suitable actuation methods.

- *Display*. As the AR system only requires a relatively small amount of information to be overlaid on a users normal field of view there is the potential to use very small cheap LCD units to place 'hot spots' to represent devices. However it is clear that for practical purposes a smaller and cheaper head mounted display is required.

- *Home Systems.* Here HS is demonstrated as an ideal technological infrastructure for complimentary AR modelling. However for real Mirror World type developments physical locations must be provided in node data structures and protocol for the interaction of the HS model and AR model must be developed.

Acknowledgements. The authors would like to thank Wayne Rudge for work on the UIFC; Darren Wenn for hardware development of the IR link; Martyn Cooper and HS-ADEPT for the use of the Demonstration System; and John Wann for ue of the Head Tracker.

6. REFERENCES

R. Azuma & G. Bishop (1994) Improving Static and Dynamic Registration in an Optical See-Through HMD *Proceedings of SIGGRAPH '94* (Orlando, FL, 24-29 July 1994) *Computer Graphics,* Annual Conference Series.

M. Cooper & D. A. Keating. (1996) Implications of the emerging home systems technologies for rehabilitation. *Med. End. Phys*. vol. 18, pp176-180. April 1996.

J. M. Ferreira, J. L. Ramalho, M. Cooper, D. A. Keating. (1995) Smart Homes and the HS-ADPET Project: Applications and User Interfaces. *Proc. ECART3 Conference,* Lisbon, October 1995. pp341-343.

D. Gelernter (1992). *Mirror Worlds*, Oxford University Press, Oxford.

J. Vince (1995). *Virtual Reality Systems.* Addison-Wesley, London.

Home Systems Specification Release 1.2. (1995) Technical Report, Department of Cybernetics, University of Reading.