A Survey of Virtual Reality Literature

Anthony Steed
Department of Computer Science
Queen Mary and Westfield College
Mile End Road
London
E1 4NS
email steed@dcs.qmw.ac.uk

October 22, 1993

Abstract

This paper reviews the current technologies that have been called *Virtual Reality*. Much interest has arisen in these technologies because of the powerful sensation of being elsewhere that they create. Many aspects of the design of a virtual reality system are considered; the interface devices, system management software and actual virtual world structure. Also described are a few specific applications that have been developed and possible avenues for future research.

Keywords Virtual reality, virtual environments, interactive 3D graphics, head mounted displays, presence.

Contents

Chapter 1

Introduction

There have been many attempts at defining what Virtual Reality as a technology and even a social phenomenon consists of. Generally it is agreed that a system claiming to be a *virtual reality system* should provide the participant with a sense of *presence*, that is it should provide an illusion that the participant is physically present in a place other than that where their real body is. Many types of system give such a belief to a certain extent, indeed films are often classified as good or bad depending on whether or not the viewers felt 'there' and had observable reactions to the on screen action. On this criterion we should include such media as video, computer games, photography, compact discs and so on in our definition of virtual reality systems. Warren Robinett [Rob92] proposes such a taxonomy that covers the whole breadth of technologically mediated experience.

Obviously we can't get away with such a broad definition and hope to cover all the relevant technology. Taking a narrow definition, a virtual reality system is one similar to Ivan Sutherland's 'Ultimate Display' [Sut65], a system that can present information to all the user's senses at a resolution equal to or greater than that he or she can discern so there is no way to tell that the artificial world is not real. Such a system does not yet exist, but its enabling technologies are under development.

The history of 'immersive media' goes back maybe four decades to work on teleoperator systems for use in the nuclear fuels industry. Following the definition of Thomas Sheridan [She92a] *a teleoperator is a machine that operates on its environment and is controlled by a human at a distance.* The use of a typical teleoperator may involve the operator strapping a rig on to his arm that senses the joint movements and through which a robot arm is controlled. Such a basic system could be extended by having the operator's vision substituted for a view provided by a camera in the robot's environment and providing other feedback such as sound and force. In its original conception though it was just a device for manipulating nuclear fuel safely at a distance, but is an early system that tried to fool one of the human senses into believing artificially created sensations were real.

One of the first systems whose purpose was to provide the user with a sense of presence was created by Morton Heilig with his 1962 'Sensorama Simulator' [Rhe91, pages 49-53]. This was designed as an arcade machine that would provide the sensation of riding a motorcycle through Brooklyn. It provided an experience where the participant saw a 3D movie of a motorcycle ride complete with engine vibrations, wind and smell effects.

The major difference between Morton Heilig's system and current systems is in the fact that the Sensorama machine was non-interactive whereas today the emphasis is on systems that provide the user with control of the experience and the ability to alter it at will.

Myron Krueger various 'Artificial Reality' systems do provide the high degree of interaction that is missing in Morton Heilig's Sensorama [Kru90, Kru91]. Krueger's fundamental concept is that of an environment that responds to the participant's movements. In his system VIDEOPLACE the participant enters a darkened room and is confronted with a wall-sized display whose image is their own silhouette. A hidden computer can add graphic objects to the display and the participant can interact with these objects in a multitude of ways. Many scenarios have been created, one of the most amusing of which involves a CRITTER, a small graphic creature which with the right coaxing will jump on to an outstretched hand, dangle from a finger or perform a jig on the participants head. Whilst this scenario is for entertainment there are serious applications for this system, one of which is remote conferencing as the images of several environments can be combined so that the participants in their individual environments can all interact in the same artificial space.

VIDEOPLACE also differs from most current systems in the fact that the user sees a third person view of himself within the computer generated environment. A typical system today presents a first person view which moves realistically as the user moves. That is to say that if the user moves forwards, the object in front of him gets closer and if he turns his head right(left) the scene moves left(right).

Virtual Reality as has been popularised by the media over the past few years has come to mean a computer system such as described in [FMHR86, Fis90, SIG89, SIG90, Bri93] that has several key components:

- A head mounted display that has two screens, one for each eye, through which a computer presents a stereoscopic 3 dimensional image.
- A tracking system that can determine the position and orientation of the user's head.
- An input device that the user holds in or wears on his hand, that is also tracked, which operates as a locating and picking device.
- Some form of audio output.
- A database that holds information about the virtual world, its component objects and possible interactions.

Typically the virtual world will consist of a scene which the user can move through and alter by operating the device he is holding.

These three systems fit very well in to the Autonomy, Interaction and Presence (AIP) cube proposed by David Zeltzer as a means of classifying virtual reality systems [Zel92]. A virtual reality system needs to be interactive, that is the user is able to change events and not just be a spectator as in Sensorama. An ideal virtual reality system would have to provide a high degree of presence so that the user felt that they were in the world created by the computer. It would have to be autonomous, that is the virtual world would not be a passive structure, but should have the ability to react to a wide variety of stimuli. For example our world might have gravity so that any objects that are above the ground will fall automatically.

In the rest of this paper I shall discuss technologies that attempt to satisfy the above three criteria. Firstly I shall discuss the sense of 'presence' factors that might enhance it and why it is desirable. Secondly I shall describe the hardware components of virtual reality systems and the various current limitations. Next I shall describe various issues concerned with the design of the virtual world itself and the objects that will inhabit it. Then I shall consider some issues concerning navigation and object control in virtual environments. In the next section I shall present some current applications and specific problems in their designs. Finally I discuss research directions in virtual reality and possible future applications.

Chapter 2

Issues of Presence

The method of providing interfaces to computer systems has changed remarkably over the past 50 years. John Walker [Wal90] identifies five stages that the development of user interfaces has gone through to date.

1. Plugboards, dedicated set-up
2. Punched card batch
3. Teletype timesharing
4. Menu systems
5. Graphical controls, windows

Each stage is a step towards providing a system which is more intuitive to use and which doesn't require specialised knowledge of the actual workings of the computer. Modern computers mostly provide a 'desktop' view, where various icons on the screen represent files that can be dragged onto other icons representing applications with the end effect of having that application perform its function on the file. Frederick Brooks [Bro77] in comparing the process of designing a graphical user interface to making a new tool, provides evidence that the use of computer graphics can improve user efficiency.

However it is not obvious that the desktop metaphor extends to all applications. For instance alongside a architectural design package we may want a system that allows us to 'walk through' the building we have created, that is be able to move around inside a computer generated model of the building and to visually experience the space [Bro86]. This is just one of virtual reality's driving problems and the advantages of a system with a sense of presence are obvious. For architectural walkthrough we require: the ability to move through the model, the ability to effect changes in the structure of the model and the ability to judge the aesthetic qualities and physical limitations of the model. The first two we can provide with a single screen projection, as used in Virtus Walkthrough [Vir90] and Superscape [Dim93] systems.

However if someone wishes to judge the aesthetic qualities and physical limitations of a building then it is necessary for them to be able to move around it in a natural fashion and judge objects against their perceived location and size, which may happen properly only if they present. Teleoperator systems also benefit from a degree of presence noted by Held and Durlach [HD92] *the best general purpose system known to us (as engineers) is us (as operators).* This means that if we want to do a job remotely then it is more desirable for the operator to be fooled into thinking that the robot arm is his own arm than to think of the arm as another tool. This is because the operator knows how to perform tasks with his own arm, whereas he may have to think how to use a tool if he is unfamiliar with it.

There are several criteria that are proposed to increase the sense of presence [She92b, Loo92a, Loo92b, Hee92b, SU92, SU93a].

1. The data presented to the senses should be of high resolution.
2. The data should not be obviously from an artificial source. For example the displays should be refreshed at a rate high enough so the user does not see flicker and the displays themselves should not be so heavy that this becomes a source of fatigue.
3. The data presented should be consistent. For example if an object in view makes a sound then the sound should appear to originate from that direction.
4. There should be a wide range of possible interactions that the user can make. For example if there is a virtual table in the environment then you should be able to not just see it, but touch it and feel its weight.
5. The operator can effect changes in the environment.
6. There is a direct visual consequence of each of the user's movements.
7. There should be a obvious mapping between the user's movements and the movements of the virtual body or slave robot.
8. The virtual body or slave robot should be similar in appearance to the operator, so there can be an identification between the user's limbs and those of the representation.
9. Other objects or users in the environment recognising and acknowledging the user in some way (such as a door opening as the user approaches).

Held and Durlach [HD92] and Loomis [Loo92a] also note that the operator's sense of presence can increase over time. Loomis goes on to to suggest that 'distal attribution'[1] and therefore presence might occur when the user can accurately and transparently model the causality of the virtual world.

Having decided what factors contribute to the sense of presence we still have to face the problem of deciding to what degree the user feels actually there, so that we can decide whether changes we have made to our system are an improvement or compare our system to another. Several methods of measuring presence and distal attribution have been proposed [She92a, Loo92b, SU92, Sla93].

1. Users reported sense of presence. This is a complicated process because the process of enquiring the state of the user may indeed change that state. Another complicating factor is that the user's sense of presence may vary over time [SU92].
2. Observations of the users behaviour. This takes observable reactions to certain situations as confirmation of the users presence. For example shying away from looming objects or replying to a welcoming 'hello' message.
3. Performance of tasks in real and virtual environments. This assumes that if a user performs a task in a virtual environment as efficiently and in the same manner as they do in a real environment then they must be present in that virtual environment.
4. Discrimination between real and virtual events. This tests, for example, the users ability to differentiate between sound cues that originate within the virtual environment and originate in the real world.
5. Incorporation of external stimuli. If the user interprets an external event, such as a loud noise, in the context of the virtual environment then they must be present in that virtual environment.

There is some evidence to suggest that each persons sense of presence depends on their psychological make-up [PT93, chapter9] [SU93b]. If this is so then there may not be a virtual reality experience convincing to all users, which may not be surprising if we consider the vastly differing subjective opinions people hold about other immersive media such as films.

[1] That is the attribution of sensory input, which is an internal phenomena, to external sources.

Chapter 3

Real Components

The hardware and software components that go to make up an immersive display system are numerous and come from a wide variety of disciplines, from optics to artificial intelligence. Systems of the type described in the introduction cost several tens of thousands of pounds and consequently have not yet found their way in to the consumer market. However cheaper systems that provide some immersive qualities have been built [Pau91] using such devices as the Mattel Power Glove and Sega Glasses and it has been rumoured that computer entertainment companies, such as Sega and Nintendo, are planning home virtual reality systems.

3.1 Displays

Most virtual reality systems provide some means of providing a 3 dimensional display. This section describes various systems for presenting 3 dimensional images and some rendering issues peculiar to virtual reality. A few issues that are important in virtual reality displays are:

- The update rate of the display. If the frame rate is low then not only will the image not be smooth but task performance may suffer [TLF+92].
- The resolution. Most systems today are present an image whose resolution is below the threshold of legal blindness [Dee92].
- The field of view. Our own vision has a field of view of almost 180^o in the horizontal plane and 120^o in the vertical plane. Currently systems have to choose between high resolution and wide angle of view [PBL+92].

3.1.1 Creating 3D Images

There are many cues that can give the illusion of depth on a stationary single screen image: perspective, occlusion, shadows, shading and texture. When the image is being updated there can be parallax and rotation cues as well. However with displays that do not take into account the viewing position of the user the illusion of space can be destroyed by moving the head left to right parallel to the display to create a situation where parallax and occlusion cues are wrong or by moving towards the display causing the objects to loom incorrectly [vdG86].

The stereo displays described below reinforce the illusion of depth by presenting disparate images to the left and right eyes [Pat92a, Pat92b]. If each eye is considered to be a pin-hole camera then to calculate the correct viewing transforms to apply we need to consider three things: the interpupilary distance, difference in view direction of each eye, and possible distortion that may be caused by the hardware's optics [RR92]. However with these simple considerations there are still conflicts with the actual mechanisms of the eye. Problems have emerged with users accommodating inappropriately [EPC93] and suffering from visual stress due to having to accommodate to a fixed distance to see the screen whilst the convergence of the eyes changes to track objects [RW93].

The simple pin-hole camera model also ignores the effects of motion blur. This can be simulated and often is done in films, however in an interactive system we will also need to track where the eye is looking as the effects of motion blur can be negated by visually tracking the blurred object [vdG86].

Providing a stereo display takes double the time in rendering and it is interesting to note that Liu *et al.* found that either occlusion or disparity provided sufficient cues to perform a simple 3D tracking task, but using both together gave little additional advantage [LTS92]. Also Sollenberger and Milgram found rotational cues more effective than stationary stereo cues [SM91].

3.1.2 Display Hardware

For many years, the flight simulator business drove the development of displays. Early simulator technology used several flat screen displays surrounding the user in a mock up of the cockpit of the aircraft he was training for. This is true to a certain extent today especially when the simulator is of a large craft that has more than one person in the cockpit, but there are a growing number of machines that use head-mounted displays. Military aviation has recently taken up the baton in display development with head-mounted displays that augment the pilot's vision with virtual instruments so the pilot does not have to change his direction of gaze when he needs to consult an instrument. Such a system is the Visualy-Coupled Airborne Systems Simulator (VCASS) built under the direction of Dr Thomas Furness [Gai87, Fur88].

There are four types of display in current use.

- Head-Mounted Display (HMD).
- Binocular Omni-Orientated Monitor (BOOM).
- Through the Window.
- Audio-Visual Experience Automatic Virtual Environment (CAVE).

Head-Mounted Display The original HMD was created by Ivan Sutherland in 1968 [Sut68]. The system, nicknamed 'Sword of Damocles', was a large contraption that was suspended from the ceiling and which could track the user's head position and orientation and provide the corresponding view. The displays were miniature cathode ray tubes and the images could be superimposed over the outside world view using half-silvered mirrors. Unfortunately the display was years ahead of rendering machines and so the views were just simple wire frame images.

Further development has produced systems that are less cumbersome and uncomfortable to use [CHB+89]. The NASA Ames HMD [FWCM88] is a comparatively lightweight system that is tracked by magnetic sensors rather than mechanically. It has miniature liquid crystal televisions as displays and as these require a low voltage they are safer to use than CRTs. The resolution is $620*220$ pixels with 16 greyshades and with the use of LEEP optics this provides a 120^o degree field of view.

Several other systems along a similar line has been built, mostly with colour displays of lower resolution. Kalawsky describes the difficulty of classifying liquid crystal displays as the given resolutions may be misleading or incorrect as they give the overall resolution of the display as the total of red, green and blue cells and the field of view on its own doesn't give any indication of the field of view for each eye and the corresponding binocular overlap [Kal92].

Other approaches to delivering the images to each eye include the fibre-optic helmet-mounted display system (FOHMD) developed by CAE Electronics for use in air combat simulators. In this system four light valve projectors transmit the images through fibre-optics to a wide angle display.

BOOM This system is similar to the HMD, but the display is mounted on an articulated arm that measures the position of the display. The display is counterweighted so that it will remain stationary when released. This enables the user to swap his attention between the BOOM and his computer monitor, though if he wants to change his view point he has to hold the display to his face.

Through the Window Under this section come the systems that have only one display screen, but still provide a 3D display using one of a number of stereo techniques [Sta91].

The first technique involves creating the two views in different colours, usually red and green anaglyphs. The user then has to wear a pair of glasses with one red and one green lens. Through each lens therefore you can only see the view corresponding to the lens colour, so each eye has a separate view. This system has been used not only in computer graphics but in comic books and films.

The second technique also requires the user to wear a special set of glasses, but this time they have liquid crystal shutters as lenses. The software renders the left and right views alternately and also signals the glasses which display it is showing so the other lens is can be turned opaque. This is the principle the Sega Glasses use.

Finally instead of having liquid crystal displays, each lens has a polarising filter, and a filter is put over the display which can change its polarising orientation. This technique has also been used in film cinemas.

The advantage of these systems over HMD or BOOM technologies is that several users can view the same display and get a stereoscopic view, though the perspective will not be correct for all the users.

CAVE The CAVE is a cube with display-screen faces surrounding the viewer [CSD+92]. Each each cube face displays a perspective view correct with respect to the primary user whose position is tracked. Thus the display on each face need not be recomputed when the user turns his head, only when he changes position. Other people may enter the cube, but they will not see a correct view.

Other Displays Some other display technologies do exist, which try to combat the problems of the above mentioned technologies. There is usually a problem with the lag of the display, due to the time needed to sense the new position of the head and compute the image. The CAVE goes some way to eliminating this problem as the whole surrounding view is computed, so as the participant moves his head around there is no need to recompute any of the image. In a teleoperator situation, lag is caused by the time required to get the video signal from the remote scene. A method of reducing this lag has been described [HY92] where the remote camera grabbed a panoramic image, and the local system can decide which area to display to the head mounted display. Again, as long as the user only makes head rotations then this system can reduce the lag, but as in the CAVE, translations require the grabbing of a whole new panorama. Another problem with the usual modes of display, is that the focal length of the image is fixed which can lead to eye strain if the system is used for long periods. Two methods have been proposed to deal with this, volume scanning display and variable focal length mirror. The volume scanning display uses a monitor screen that can moved back and forth whilst the displayed image is altered at each position. By displaying the slices of the object at the correct depth, the illusion of a true 3D image can be created [KOF92]. The varifocal mirror concept involves a similar idea. By projecting the image slices on to a mirror with a variable focal length a 3D image can be presented [FvDFH90, pages 916-917]. Research is also being carried out into developing high resolution displays that use microscopic lasers that can be shined directly on to the retina.

3.1.3 Rendering issues

Considering the actual rendering of the image there are several techniques that can be used to reduce the delay involved with the rendering of each image. The database can be split down into areas each of which has a set of potentially visible areas so that the whole database does not have to be considered when rendering [ARB90, FST92, TS91]. Another technique is to consider the fact that the eye has a greater concentration of rod and cones near the fovea centralis. If the eyeball is tracked then the rendering software can concentrate on rendering the area of the screen at which the eye is looking [LW90]. Also we could use progressive refinement, that is improve the image definition when the head is stationary. This may involve having more than one set of object data, each set having a different number of polygons, the crudest of which we render when moving and the most detailed we draw if the user is stationary for a long period. An associated technique is to choose which set of polygons to draw depending on whether the viewer is within a certain range of the object. This technique has been used in flight simulators where objects may be seen at a wide range of distances [Vin92]. Another technique used in flight simulators is to have texture mapping of a relatively small number of polygons.

3.2 Audio Technology

There are several reasons for having sound in a virtual environment. Firstly, to maintain consistency, objects that usually make sounds should do so. For example if there is a model of a moving car in the environment then engines noises would be appropriate. Secondly sound cues can be used to provide information to supplement the visual cues. In the above example modelling the doppler shift of the engine noise can give the car a rough velocity. Also sound cues can give information that can not be visualised, such as an engine noise that indicates that there is something wrong with the engine. We may also give the user audible cues to signal certain events in the environment. For example if trying to pick an object up then a successful hit could be signalled by a bell. Such cues can compensate for the lack of tactile cues.

These techniques all help to relieve the load on the user's visual sense. Audio cues have the advantage over visual cues that they are more quickly recognised as they have a startling effect and audition is a temporal sense and we are sensitive to small changes in an audio signal over time [Wen92].

When using sound cues a useful technique is to localise the cue, that is give it an apparent location in the virtual environment. Normally when wearing headphones the sounds often seem to originate inside the head. To externalise the sound requires the application of techniques that model the propagation of sound waves through the air and the pinnae of the ear. The basic technique used in the acoustic display is to model the difference in arrival time and amplitude of the signal at the left and right ear which gives the sound a rough horizontal direction. To go further and give sounds an elevation requires the modelling of the distortion caused to the signal as it passes through the outer pinnae. Unfortunately the signal distortion is different for each person so the function used to process the signal needs to be individually generated by using standard cues in various dimensional locations and measuring the signal received with a microphone placed inside the ear. This technique is discussed at length in [Wen92].

But as discussed in [DRP+92] there are other factors to consider that can help localise an auditory cue. The reverberation of the sound might need to be considered, leading to a technique, sound tracing, similar in concept to that of light ray tracing. Also an important factor to consider when there are both visual and audio cues, is that they are synchronised. The problems in this area are inherent in the design of the hardware, with sound and graphics being handled by different processors that may, and probably will, have different rendering speeds. Takala and Hahn note these problems and suggest a sound script which is similar to and coordinated with the animation script [TH92].

3.3 Tracking

Most head mounted displays use a form of tracking to sense in which direction the user is looking so that the correct view can be displayed and a similar system may be used to track the hand position and orientation. There are four tracking methods in current use.

- Mechanical
- Magnetic
- Optical

- Sonic

In evaluating each of the designs, there are several important criteria to consider: Accuracy and resolution, that is the range in which the reported position is correct and the minimum change that the system can detect; Responsiveness, includes the rate at which data is output and the lag with which it is reported; Robustness, the ability to operate in any environment; Registration, the correspondence between reported and actual position; Socialbility, including range of operation and ability to track multiple objects. These criteria are discussed in [MAB92] and an evaluation of current technologies is presented.

Mechanical Sutherland's original HMD was mechanically tracked with a shaft-like assembly that hung from the ceiling [Sut68]. The use of such a system is currently limited to either tracking the head or the hand as two tracking systems would interfere mechanically. However the systems are very accurate and can be combined with some sort of force feedback mechanism as in the GROPE system described later. Another system that uses similar technology is BOOM, described earlier in this paper.

Magnetic Magnetic trackers are the commonest to be found in virtual reality systems, because of their off-the-shelf availability. There are two major competitors in this market, Polhemus whose 3Space Isotrack and 3Space Tracker, made popular by their usage in the VPL eyephone and dataglove, use AC emitters, and Ascension who make trackers based on DC emitters. The principle of both systems is similar. The emitter has 3 orthogonal coils which when a current is passed through each in sequence produce 3 orthogonal magnetic fields. The sensor similarly has 3 coils and by measuring the nine induced currents the relative positions and orientations of the coils can be worked out.

The Polhemus devices use a continuously varying AC current. There are various configurations possible. A single emitter and receiver can update at 60 Hz. Four emitters can be used which reduces the performance to 15 Hz per emitter. However several researchers report worse practical performance, one of only 20 Hz [WKD90, Pau91, WCF90, Wen92].

The Ascension devices use a pulsed DC current. There are three systems. Bird is a single sensor unit, Big Bird is a single unit with a larger sensing volume and Flock of Birds can support up to six sensors. These units are rated at 100 Hz.

The Ascension devices perform better than the Polhemus devices as they have a higher update rate and because the fields are generated by short pulses there is a lesser effect due to eddy currents induced in metal objects in the environment. In fact the noise in the signal of a HMD tracker often leads to image instability, which can degrade the user's sense of presence and lead to a form of motion sickness. To deal with this and the problem of lag some form of predictive filter can be used [LSG91].

Optical There are two approaches to tracking with optical systems [WCF90]. The first, outside-in, involves having several cameras positioned around the working area that track emitters placed on the user's HMD and hand. The second, inside-out, has the sensors on the HMD looking at emitters positioned around the environment. Inside-out systems have the benefit of being able to work in a larger volume. This is because tracking performance degrades rapidly with distance between the emitters and sensors, but with the inside-out system there can be a large number of emitters in the environment so that the sensors can always see the required number within a short range. Inside-out systems have been developed at University of North Carolina and there is a system, *Honeywell Rotating Beam*, that is used within aircraft cockpits. An outside-in system is the SELSPOT system where a stereo-pair of cameras track LEDs. For references to these systems see [MAB92]. Both systems suffer from the problem of occlusion, though with more sensors, and/or emitters, this can be overcome.

There are a couple of other less developed systems. Two Honeywell systems use pattern recognition algorithms to track known symbols, one using an LED array, the other using graphic patterns. Another system uses laser ranging to determine position.

The system used in Krueger's VIDEOPLACE uses realtime processing of a video image of the user. Whilst the video image can give no indication of the direction in which the participant is facing, it can give information about positions of limbs and whole body postures. Deciding the position of the body from an image is difficult and there are many ambiguous poses but this system has the advantage of not needing any special clothing or intrusive machinery.

Sonic Again there are two different approaches to using ultrasonics as a tracking technology, time-of-flight and phase-coherent.

Time-of-flight systems measure the time taken for the sound to reach the sensor from the emitter. Using several emitters which are pulsed in sequence the position and orientation can be judged. Problems occur with ambient noise due to disk drives and lights and there may be a problem if the sensor is occluded with respect to one or more of the emitters. It has also proved difficult to judge the exact moment of arrival of a sound due to its wave characteristics and there must be a short wait after each pulse to allow the emitter to die down and stop sending out sounds waves. However this technique has been used in the *Mattel Power Glove*.

Phase-coherence devices measure distance by comparing the phase of the received wave to that of the emitted wave. Sutherland used this system as an alternative to the mechanical system in his 1968 display. This technique is more accurate than the time-of-flight and suffers from fewer problems due to noise, however when the phase difference is more than one wavelength then errors can occur.

3.4 Tactile and Haptic Input

For many tasks just knowing the position and orientation of the user's head and hand is not enough. In teleoperator systems measuring the position of the whole arm and hand enables the remote or virtual operator to be controlled in a way mirroring the motions of the user, the GROPE system at the University of North Carolina at Chapel Hill is a good example [BOYBK90]. This technique can be extended to measuring the position of the whole upper body and creating a remote robot highly anthropomorphic in design such as the Green Man system [Utt89].

More commonly we wish to measure the finger positions of the hand in detail, as this allows us a natural way of communicating to the computer using gestures and the facility to use tools that have representations close to their real design [KH90]. The most popular device to measure the flexion values of the fingers is the VPL Dataglove (a diagram of one is in [Fol87]). This is a lycra glove that has fibre-optics running along each finger that sense the finger flexion by measuring the loss of light through the cable. This is done by scoring the cable in several positions so that when the finger is bent light leaks out of the cable. This technology has been extended to a whole body suit, the VPL Datasuit. A cheaper alternative to the Dataglove is the Mattel Power Glove that measures voltage through conductive ink strips along the fingers. Both these devices have their limitations, the Power Glove only senses gross finger motion and the Dataglove suffers from sensitivity to the exact position of the glove on the hand and may need reconfiguring for each user.

A more accurate device is the Exos Dextrous Hand Master, that has a series of mechanical joint sensors along each finger [Mar92]. Work is currently going on to provide finger sensing by using a video image and pattern recognition [Kah].

Once the finger positions have been obtained, then if we are using a teleoperator system, we can drive a robot hand directly [Spe92].

Within a virtual environment there are three approaches to whole-hand interaction [SZP89]:

1. Direct Manipulation. The user reaches into the simulation to manipulate and guide the elements as he would manipulate elements of the real world.

2. Abstracted Graphical Input Device. Here the hand acts as various abstract devices such as button, locator and valuator.

3. Gesture Language. The stream of hand positions can be interpreted as tokens of a language such as American sign language.

The distinction between the above three approaches is quite blurred, but one general principle that is used when designing the gesture set is that the gesture should be easy to remember and resemble the corresponding real world gesture as far as possible if such exists [KYB90, ME90, PG92, VB92].

Another interesting input device has been developed that allows the user to actually walk around instead of using a gesture to do so [IM92]. The user is suspended in a frame and has roller skates on his feet so that when he takes a step the device can register so and move the eyepoint correspondingly.

3.5 Tactile and Haptic Feedback

Positioning an object within a virtual environment is difficult and benefits immensely from force cues to indicate when the users virtual hand or tool has come into contact with an object. A simple system to provide such cues is described in [HHK92] along with performance evaluations. The users wears a magnetic ring which is positioned inside a magnetic head. As the user moves his finger the head, sensing the changing magnetic field, can track the finger. When the position of the tracking head intersects an object in the virtual environment, the head stops tracking and the finger is restrained.

Some systems such as GROPE the mechanism can provide forces of various magnitudes. This has proved useful in one of the major applications of GROPE, that of molecular docking where the force feedback gives a useful indication of the attraction between two molecules [Bro88, BOYBK90]. There is evidence to suggest that performance improves significantly when force feedback and compliance cues are provided [DZW+92].

On a smaller scale, to provide a sense of touch, we need to replace the basic glove with one that can provide some tactile feedback. One device that can be added to the Dataglove is described by Burdea *et al.* [BZR+92, BRS+92]. This is a small assembly that straps to the hand with 3 small pistons that can restrict movement of the thumb, first and second fingers, to give simple force cues to indicate for example when an object has been grabbed. Another glove device is the Teletact glove that uses air pockets to provide force display [Sto92]. These two devices currently do not affect whole hand position as well as hand posture. One system that attempts to do both has been designed at the University of Tsukuba [Iwa90]. This device sits upon the desktop and can provide 9 degrees of freedom, 6 in the position and orientation of the hand and 3 for the fingers and thumbs. One example use of this system is to provide virtual handling of prototype objects, so that the designer can get a feel for the weight distribution of his design.

Once we have simulated the forces that can be applied to the hand, we can try to extend our technology to the individual fingers. The fingertips are especially sensitive and have a resolution of 1mm being able to detect vibrations of up to 500 Hz [Rhe91, pages 319&321]. There is therefore a very difficult engineering problem to overcome of getting enough controllable cells to provide the texture simulation. Several systems are in development and the Teletact glove mentioned above is one of them. An interesting approach to the problem is to provide a surface that changes shape by the use of electrorheological liquids, that is liquids that change viscosity when an electric current is applied [Mon92]. Another, created by the TiNi Company, provides an array of pins where each pin can be individually controlled [Rhe91, pages 341-342].

The techniques of tactile display are discussed in [MOYS+90] along with a presentation of the Sandpaper system that can display textures and force fields with the same technology. An analysis of the forces required to perform the simple task of placing a peg in a hole can be found in [KS92].

3.6 System Architecture

Sutherland's original head mounted display required the use of dedicated display hardware, but even still the output was crude wire frame graphics. Today with the introduction of fast graphics chips such as the Intel i860, displays usually consist of filled flat-lit or gouraud shaded polygons. The Silicon Graphics Iris machines have proved popular machines on which to create virtual environments, due in no small part to the systems provided by VPL [BBH+90].

The major problems to overcome when designing a system are the *Multiple Agent Problem* and the *Animation Problem* [RCM89].

The multiple agent problem occurs when the various application and input/output device agents, which will have different time constraints, are competing for processor time. A good example of this problem is given by systems that use a Dataglove for input. One way of getting data from the glove is by 'polling' the glove, where the device interface requests data from the glove and then has to wait until the data arrives before it can carry on. Such a system will be seriously slowed down by the lag rates incurred by the use of the magnetic tracker.

The animation problem then arises when the system is supporting smooth animated graphics and also the various agents described above. The specific problem here is that we wish to have a constant frame rate for the display, but however we don't want to achieve this by slowing the system down so that no matter what demands the applications and devices make, the frame rate will not decrease.

Obviously the solution is to distribute the various processes over several processors, with the graphics rendering being the most obvious choice. The off loading of the rendering to special graphics card is what makes virtual reality on personal computers a possibility. In fact the W Industries Virtuality machines are based upon the Commodore Amiga computer with one, very fast, graphics card for each eye display [W I]. Similarly the RB2 system from VPL has a Machintosh controlling the virtual world and input/output devices and two Silicon Graphics Irises rendering the graphics.

The Provision and Supervision systems from Division provide dedicated parallel architectures that support virtual environments. Each of the various agents (or actors) can be supported on a different transputer. There is an actor to handle the tracking, an actor to calculate collisions between objects and an actor which renders the views. These various actors are coordinated by another actor called the director who handles message passing between actors. Also there may be other actors created within applications to model, say, the physics of the virtual world [Gri92b, Div91]. A similar approach was taken in the design of the AVIARY system from the University of Manchester [WHH92, SWH93].

One problem to consider when constructing the distributed system, is the form and content of the messages that will be passed around the various agents. The messages vary from raw data such as the angles and positions provided by the glove to gestures such as first finger extended to commands such as 'fly in this direction'. A hierarchical and easily extensible structure of messages is described in [ALK+92], where higher levels rules such as commands are device independent. A similar approach to the problem of specifying 3D interactions can be found in [BHV92]. A toolkit designed to enable the construction of network transparent virtual reality applications is described in [SLGS92].

A single user system is useful for many applications, but networked environments are useful, especially in training roles. This then creates the problem of connecting perhaps hundreds of machines together so that each machine see a consistent and interactive world. The creation of the Habitat system provides many useful lessons in minimising the amount of network traffic required to run an interesting and reasonably complex world [MF91]. Habitat was designed with the Commodore 64 home computer acting as a front end connected over a commercial data network to a centralised back-end. The front-end provides a user interface, graphical output and message sending, whereas the back-end maintains the world model and informs up to several hundred currently logged on guests about the state of the world in their area. Players interact with objects and other players in a cartoon like world, where speech appears as balloons above the user's head. The authors use object-orientated data representations, so that a scene description relies on a list of the objects it contains, not the object descriptions; it is up to the the front-end system to decide what to present on the screen. This allows several different front end systems to be used, from text terminals which might just give a scene description to top-of-the-range graphics computer that might the scene out of fractals.

One of the major uses for networked virtual environments is in military training, the SIMNET system, is a network of around 200 tank simulators, all of which interact within the same world [All91, Mos93]. The simulations in the SIMNET system rely on a dead reckoning model. Each object in the virtual world has a host machine which handles its dynamics and every other machine has a simple representation (or ghost) of this object which is assumed to be travelling at a constant velocity in its current direction. When the object's host decides that the positions of the remote ghosts will be significantly different to the actual position then it sends a message to all the other machines to correct the discrepancy. The SIMNET system was designed in 1983 and the wire frame graphics are crude by today's standards, but work is going on to provide the next generation of simulators, using polygon graphics, but the same basic communication network [BHML92, ZPMW92].

Collaborative virtual environments are currently receiving a lot of attention because of their potential for collaborative work and the many applications this would have in the business world. A few that exist include Virtual Environment Operating System (VEOS) being developed at the University of Washington [BC93], BOLIO at the Massachusets Institute of Technology, Distributed Interactive Virtual Environment (DIVE) at the Swedish Institute of Computer Science as part of the MultiG project [Fah91, Fah93], and Rubber Rocks [CJK+92].

Chapter 4

Virtual Components

Once the hardware and software components of our virtual reality system are in place we still have to decide what the participant should hear, see, feel and even smell. In other words we have to design the virtual environment. Should we then choose to make the environment realistic so that the users can interact as they would in real life. Certainly this is one viable option and it is essential in such applications as architectural walkthrough, personnel training and remote operation. But what if we want to explore phenomena usually outside our perceptual range? For example, the GROPE system's display is of molecules, for which there are many representations. In fact virtual reality can and has been used to visualise scientific data in many ways.

Environment Structure What should the structure of our virtual environment be? An obvious approach would be to make the virtual world as similar as possible to the real world. However modelling the physics and constraints of the real world is a time consuming process. Calculating collisions between objects is an $\mathcal{O}(n^2)$ problem in the number of objects in the environment. Then deciding what happens when two objects collide, especially if one or more of them is partially elastic, is also a complex problem[Pen90]. Architectural walkthrough in particular would benefit from these considerations. For example the problem of arranging fittings in a virtual building would be easier if objects could be pushed around and stacked in a realistic fashion [PG92].

Not all virtual environments will need such exhaustive checks, and the AVIARY system supports a hierarchy of world models, where leafs in the hierarchy correspond to specific application worlds and the root is a generic world [WHH92].

Another approach to virtual world design might be to extend the desktop metaphor to 3 dimensions. The NASA Ames system [FMHR86, FWCM88] was designed with the idea of virtual workspaces in mind and a 3D equivalent of windows where various applications surround the user in the virtual space. This could be extended to providing arbitrary data structures that surround the user which could be organised with a fish-eye approach where objects that interest the user are displayed as larger and in more detail than those that don't [FS92].

Some virtual environments are designed to help in the visualisation of data, usually 3d data that can not easily be represented on a traditional display [RCM93]. The design of such an environment need not resemble the real world at all but the metaphor used should provide insight into the nature of the data that can not be obtained otherwise. Stephen Ellis uses the example of air traffic control to illustrate how a well chosen virtual environment display could be easier and more efficient to use than the usual monochrome flat screen display [Ell90, Ell92].

A virtual world concept that has caught many writers imagination is *Cyberspace*. One definition of Cyberspace given by Michael Benedikt [Ben91b] is:

> A new universe, a parallel universe created and sustained by the world's computers and communication lines. A world in which the global traffic of knowledge, secrets, measurements, indicators, entertainments and alter-human agency takes on form: sights,

> sounds, presences never seen on the surface of the earth blossoming into a vast electronic night.

Cyberspace is a space of pure information, around which participants will navigate in order to find specific data. Some argue that the foundations of cyberspace already exist and point out the existence of world spanning data networks such as the Internet. The current Internet is a virtual environment, but one which in Zeltzer's AIP cube has high degrees of interaction and autonomy but a low sense of presence. What is lacking is the immersive visual, auditory and tactile interface so that for example to browse through a library's catalogue, you will travel to its location in cyberspace and there examine it by moving around its various dimensions. What has yet to be decided is what cyberspace will look like and how it will operate. It is quite likely that it will have more than three dimensions so that high dimensional databases can be explored [FB90]. Discussions of the various options for the type of visualisation to be used and the underlying design principles can be found in [Ben91a, Nov91]. Some work has been carried out into the implementation of cyberspace and Cyberterm a protocol for the communication in cyberspace is detailed in [Sno].

Environment Inhabitants The neglected area of virtual reality research is the consideration of what objects should inhabit the virtual environment and how they interact with each other. A world that has objects that appear to be autonomous and which can interact with the user not only makes a more interesting world than one that is static, but may also give the user a greater sense of presence.

The first problem to consider is how to make the objects move. For some objects this is trivial, a rolling sphere has a very simple model controlling its motion. A more complicated problem is that of simulating the motion of jointed figures. McKenna and Zeltzer present a system for modelling the dynamics of motion of a 'roach' within a virtual environment [MPZ90, MZ90]. The roach moves around with a natural gait pattern and it can respond to basic interactions such as come here messages. The gait is modelled using oscillators and reactions to the environment such as over extension of a leg which causes the leg to move immediately. More complex behaviours of actors must be specified in such a way that they can be broken down into motor actions. Zeltzer and Johnson present a system that can model the actor by giving it goals that it can obtain by using various skills it has at its disposal [ZM91].

Once the various actors have been modelled we need a model for the ways in which they interact. In effect we need to consider the content and style of a virtual environment. Joseph Bates identifies three areas of research [Bat92]. *Cognitive Emotional Agents* that appear to be intelligent or rather appear not to be stupid. Agents in the environment need not be sophisticated artificial intelligences, but rather have a wide range of possible interactions. Armed with a large number of shallow rules the agent can hope to get away with the 'Eliza Effect', that is the willingness of the user to attribute sensitivity and emotion to the agent as long as the agent doesn't show explicitly unreasonable behaviour. Agents have a set of goals in the virtual environment and some natural language processing abilities [BLR91]. *Presentation* is a vital component of traditional media but as yet doesn't feature in the design of virtual worlds. Bates suggests that styles will emerge as virtual reality matures and Meredith Bricken suggests that the consideration of how a world will look and interact is the most important part in the design of a virtual reality system [Bri91b]. *Drama* is the imposition of certain structure on the long term development of a virtual world. When there is no interaction with the environment the creator can specify a linear progression of events that will happen in the world. Once the user can alter events the creator's job becomes a lot more complicated and Bates likens it to a game of chess where the user and a creator controlled director take alternate moves. It is this flexible interaction between the computer and participant that is the distinguishing characteristic of virtual reality as a media and Myron Krueger goes as far as to suggest that this interaction is a new art form [Kru83].

Chapter 5

Interaction in Virtual Environments

One of the aims of an immersive system that can fool the senses to a significant degree must be to allow natural modes of interaction with the virtual environment. Thus if a user wishes to walk around inside an architectural model then they will be able actually walk. If any intervening metaphor is used then the experience will not be like walking around the actual building. However this strict adherence to a real world metaphor may need to hold for only a small number of applications in which the ability to judge spatial relationships is of paramount importance. Indeed one of the exciting things about virtual reality is the way in which it can endow the user with new skills such as flying and teleportation.

At a basic level 3D interaction involves a combination of the following tasks [SD91]:

1. Navigation - the user should be able to navigate the scene by controlling the eyepoint.
2. Global Selection - the user should be able to select any object in the the scene.
3. Rigid Body Transformation - this includes the usual requirements of translation and rotation, any transformation which changes the objects orientation and position in space, but leaves its local geometry unchanged.
4. Local Selection - the user should be able to select part of an object which can then be used to deform the object but not its overall position in space.
5. Deformation - these transformations applied to an object cause it to deform either uniformly (for example by scaling) or non-uniformly, by twisting, tapering or bending

These five tasks can be grouped together as *navigation*(item 1) and *object manipulation* (items 2-5). The tasks of object manipulation are often considered cases of object placement, by using a 3D cursor to specify the selection and then using the cursor again to specify the deformation or transformation.

With autonomous objects such as those described in section 2.4 then the interaction becomes much more complicated. Zeltzer [Zel91] identifies four possible levels of abstraction of objects:

1. Structural. The physical attributes of the object and its kinematics.
2. Procedural. Defines the processes of object motion such as collision detection, inverse kinematics, forward dynamics, or elastic deformation that act independently of the object structure.
3. Functional. Defines higher level functions of the object assembly. Such a function may be 'walk' which is the abstraction for the more complicated task of moving each joint of the object to the required position for each step.

4. Agents. An agent is the most complicated form of object. It has a structural definition, a set of possible functions and a means of selecting between various behaviours according to the environment.

Interaction at the structural level involves the basic interaction tasks (2-5) described above such as the translation of an object by selecting it and defining the new position with the 3D cursor. Guiding the procedural or functional units involves providing values to the units that then act on the object. An example of this is a message that is sent to an object to move itself forward. Interaction at the agent level is indirect and involves changing the environment so that the agent reacts.

5.1 Object Placement

The placing of objects is a six dimensional task with three dimensions required for the position and three for the orientation. Some techniques have use 2D devices to perform this task. Nielson and Olsen take the the 2D locator movement and take this to be motion along the virtual world coordinate axis which in the current 2D projection is closest to the direction of movement [NO86]. Thus the plane of motion of the 2D locator is separated into six regions which correspond to motion in the positive and negative directions along the coordinate system axes. This allows placement of a 3D cursor and then rotations can be specified in terms of this and other 3D points. Chen *et al.* give four metaphors for the direct control of the rotation of an object [CMS88]. The most efficient metaphor is that of a virtual sphere where the object is imagined to be contained within a glass sphere. Up, down, left and right motions of the 2D control device roll the sphere in the appropriate direction and rotation about the remaining axis occurs when the user moves the device along a circular path. Specifying 3D points manually is time consuming and difficult to perform accurately so Bier combats this by allowing the cursor to be attached to object vertices, lines or faces. This allows very accurate positioning and most translations and rotations can be described relative to objects that already exist [Bie86, Bie90].

Many systems now use the 6D input devices described earlier in this paper. Badler *et al.* found that using the absolute position of the hand as the cursor position could become tiring for the user [BMB86]. They overcame this by having a button in the user's free hand that would engage the device in a relative mode. They also found that it was extremely difficult to keep the cursor stationary or to move it in just one dimension. Ware and Jessome also considered the problem of object placement with a polhemus device and they found it to be an easy task to perform roughly [WJ88]. Their addition of a 90^o flip about the vertical axis in combination with a disablement of movements perpendicular to the display produced a metaphor that was simple and effective to use. Ware goes on to compare this mode with the all degrees of freedom mode and compare stereoscopic with monoscopic displays [War90].

Immersive virtual reality systems mainly just use hand position for the position of the virtual hand or cursor. The only aid to positioning afforded is the ability to move the eyepoint freely within the virtual environment.

Another device that can be used for object placement is the Desktop Bat [SD91]. This is a 5D input device that consists of a hemisphere connected by three joints to a mouse chassis. Cursor movement in 3D can be effected using a number of techniques, one where the orientation of the dome defines a virtual plane and movements of the Desktop Bat on its surface cause translations in this plane.

5.2 Navigation

Mackinlay *et al.* identify four main types of navigation [MCR90].

1. General movement. Exploration of, say, a building model.

2. Targeted movement. Movement towards a specific point of interest, such as moving into to examine a detail.
3. Specified coordinate movement. Movement to a point specified in a known coordinate system.
4. Specified trajectory movement. Movement along a predefined path.

Navigation is not a problem specific to virtual reality and it has been tackled within the framework of desktop systems. Two basic techniques exist, moving the viewpoint through the workspace or moving the workspace around the user. Essentially the difference is the choice of coordinate system in which to perform the translations and rotations. Ware and Osborne give three navigation metaphors that illustrate the difference [WO90].

- Eyeball in hand. The movements of the input device correspond directly to movements of the eye.
- Scene in hand. The scene itself is slaved to the input device.
- Flying vehicle control. The input device provides the controls for the vehicle such as velocity and rotation.

Ware and Osborne note that the scene in hand metaphor is a poor choice for complex environments because the centre of rotation is fixed and this leads to confusing translations during rotation. Choosing the centre of rotation would be possible using the techniques of cursor placement as described previously.

Once again 2D devices can be used for navigation. An example is the use of a joystick or mouse to indicate the direction of movement by pointing at an item in the virtual environment. The Desktop Bat mentioned earlier is ideal for this task due to the fact that human motion is mostly along the line of sight.

The use of a 6D tracking device allows the view to be slaved to the head position, a technique that McKenna indicates improves the ability to pick locations in space [McK92]. Due to current limitations, real motion is impossible over long distances, so a navigation metaphor must be used. Typical metaphors used in immersive systems are to fly in the direction of pointing or along the direction of gaze. Flying along the direction of pointing has the advantage that you can keep looking at a point of interest and fly around it whilst, flying in the direction of looking allows the separation of the position of the hand and navigation so that for example a virtual tea cup can be kept level during motion. However flying in either of these manners can lead to a conflict of metaphors. This occurs when the user does actually take a couple of steps in a certain direction and they move correspondingly in the virtual world. Soon though they will either get tangled in the wires or bump into a real object which suddenly reminds them that they should have pressed a button. There can also be some confusion as the user is expected to turn by using their body but walk using a button [SU92].

None of the methods above explicitly concerns the situation of choosing the best viewing position to perform an object manipulation. Phillips *et al.* developed a system that automatically chooses a suitable position for the viewpoint when an object is selected for manipulation [PBG92].

Movement over long distances is time consuming so two techniques have been used, velocity control and logarithmic approach [MCR90]. Mackinlay *et al.* show that velocity control, whilst appearing to give greater control, is not efficient when trying to position the eyepoint accurately. The logarithmic allows rapid and controlled approach towards a point of interest by moving the viewpoint towards the point of interest a certain percentage of the remaining distance each frame. This technique extends, with a little alteration, to object manipulation. This technique works when there is a point of interest, for more general velocity control Chapman and Ware's predictor based visual feedback aid can help users learn how to control the viewpoint more effectively [CW92].

Chapter 6

Current Applications

Many applications now exist that use virtual reality technology, from art exhibits to product visualisers [Lau92, Emm92, Hag92]. What follows is a description of just a few systems that illustrate the breadth of current work.

Virtual Wind Tunnel The work of Steve Bryson and Creon Levit has been mentioned before in this paper with regard to their use of the BOOM as a display. The application they developed is an example of the use of virtual reality as a visualisation tool [BL92].

The virtual wind tunnel allows the operator to visualise precalculated fluid flows. The operator can move around within the flow in and has several visualisation tools that he can place by the use of a Dataglove. These tools include *tufts* which are short vanes that indicate the velocity direction at a point and *streaklines* which are similar in concept to the smoke tracers used in traditional wind tunnels. These tools can be placed quickly and can be interactively moved with the glove. Current limitations in the system are mainly due to the vast amounts of data that the visualisation relies upon and the high storage and memory bandwidth required.

Entertainment Virtual reality entertainments systems are the most widespread application of virtual reality technology in current use outside of the academic environment. Various systems are being developed, from single player machines suitable for use in arcades to games based around flight simulator motion platforms for use in theme parks [Coo92, Gri92a, Hee92a].

One company producing entertainment systems is W Industries, whose *Virtuality* systems can have one to four players operating in the same environment [Row92, Wal93]. The games have proved quite popular, maybe due to the media hype that surrounds the virtual reality industry, though the games are quite highly priced and do not last very long. Some example scenarios are a one player flight simulator, a four player game called Dactyl Nightmare that involves the players in shoot-out and a four player Dungeons and Dragons type adventure game.

This system is comparable in price to the systems sold by VPL though all the components are made by W Industries themselves to stand up to the rigours of public use.

Ultrasound Imaging This system is just one application of a technology that is under development at the University of North Carolina. The ultrasound system provides live ultrasound pictures of a human subject overlaid on the actual view of the person. The image that the operator sees is the ultrasound image transformed to be correct as the operator would see it, overlaid on a video image of the subject taken from a small camera attached to the front of the head-mounted display. As the operator moves around the subject he sees an ultrasound image that appears stationary within the subject [BFO92].

The major obstacles with this system are: the acquisition of 3D ultrasound data at real time rates, the display of 3D data at the display update rates and the alignment of the ultrasound image with the live video. Whilst this system does not have any interactive ability and is purely a visualisation tool, the techniques involved could be applied in a large number of situations where data needs to be overlaid on the real world view to either augment the display with extra data or provide visualisation of phenomena that would the eye would not register. Some examples are: a system for fire officers that allows them to see through the smoke in burning buildings using 'radar vision', and the overlaying of complicated machinery plans on the view of the machine to aid in identification of parts that need servicing.

Three Dimensional Modelling The creation of virtual worlds was originally performed within a CAD package external to the actual visualisation software. However the 3DM system developed at the University of North Carolina provides several tools for the creation of objects within the virtual environment [BDHO92].

The system uses a HMD, 6D mouse with 2 buttons and the Pixel-Planes 4 and Pixel-Planes 5 graphics engines. Once inside the virtual world the user can select tools from a toolbox that can either float at waist level or be positioned out of the way. There are tools to create standard shapes, extrude a poly-line or specify triangles by vertex position.

Whilst creation of objects is very easy with this interface, accuracy is not precise enough for engineering design. However with the use of snapping and constraints this may be overcome.

Surgery A lot of research has gone into providing virtual bodies in surgery training because of the advantages over traditional methods [MP92, Sat93b]. A similar amount of research has gone into providing teleoperators for surgery so that operations might be performed in remote or inaccessible locations such as onboard the space shuttle [Sat93a]. Surgeons can now perform on tissues that they can not see and touch directly, but only through the use of a video endoscope and micromanipulators controlled via a teleoperator system. Once such a system is connected to a computer network, remote surgery will take place, with surgeons possibly participating in an operation in a different city or continent.

Chapter 7

Future Directions

Virtual reality is a young subject and is in the unique position of being a well publicised technology that is currently available in the marketplace but about which very few results have been published. The hardware systems used at the moment are cumbersome and of low resolution, hopefully this situation will improve in the not too distant future. Apart from the hardware systems a lot of work has still to be done on the design of the virtual environments themselves [Kal93].

7.1 Future Research

The systems described in this paper far from satisfy the criteria for presence as detailed in section two. Many new techniques will evolve over the next few years and it may be that the current 'goggles and gloves' phase may pass to be replaced by displays that physically surround the user and use video tracking as the interface device. Such a system may be a long way off, but there is room for improvement within our current interface paradigm.

At a recent NSF invitational workshop potential research areas were listed in quite some detail [BF92]. Their recommendations covered the whole spectra of hardware and software issue.

Within the hardware field, there is a need for visual displays, that are light, small and of a high resolution. The current audio and haptic systems are quite advanced but need more development to make their output more realistic. And there is a pressing need for a wireless tracking system which may be possible using gyroscopes as a form of inertial tracking.

There is a general lack of software both for creating the virtual worlds and interacting within them. There is no current standard for the description of the geometry, dynamics and interactions of a virtual world. Pressing needs are for the development of good techniques for navigation, movement and interaction within the virtual world. Another requirements is for functionality within virtual worlds which in turn will require the support of multiple prioritized processes and time and resource management.

The human-computer software interface needs to be improved as currently there are a lot of assumptions made about the specific hardware configuration, possible interactions and the physical model of the particular system the interface has been designed for. An extensible interface is needed, to allow various systems with different hardware set-ups to use the same model. Also required is some consistent theory of presence to classify the differential effects of alternative world designs.

Currently applications are not harnessing the full potential of virtual reality. This is partly due to the lack of a software base from which to develop new applications. There is a real need now to start solving real world problems rather than produce demonstrations or proof of concept models. Research also needs to be done into how people will use a multi-participant system and how this can be supported.

Further consideration is also required of the physical effects of the use of virtual reality. There are two main problems, *simulator sickness* and *cognitive remodelling*. Simulator sickness is a similar condition to motion sickness and is thought to occur when the lag of the system causes discordant visual and kinetic feedback [MS91]. Cognitive remodelling is how William Bricken describes the inclination of people who have just come out of a simulation to interact with the real world as they had been doing in the virtual world [Bri]. For example they may try to open a door by just walking into it! The long term effects of submersion in a virtual environment also need to be studied.

7.2 Future Applications

A future scenario has been envisaged where few people will travel to work in the mornings as they will be able to work from home via virtual reality [And, PB91]. This may be possible, but it is unlikely for some considerable time that virtual reality will provide the depth and breadth of functionality that the real world provides. There is no end to the possible simulations of real life that virtual reality will be able to provide, from virtual cycling to historical re-enactment. What is perhaps more interesting is the ability to create virtual worlds that are totally dissimilar to the real world, such as cyberspace.

One possible use for virtual reality is in education. Meredith Bricken believes that virtual reality can be used in ways similar to the conventional methods of hands on learning, group projects, field trips, simulations and concept visualisation [Bri91a]. Certainly just as a visualisation tool virtual reality is immensely useful as complex systems can be modelled easily. But the real potential lies in the ability to create a scenario in which the pupil can experience events rather than learn about them in a preset and rigid fact based manner.

Carlson and McGreevy both propose the use of virtual reality for planetary exploration. In their scenarios instead of sending a manned mission to Mars an advanced probe would be sent. Once the data was received a virtual model of Mars could be constructed which could be explored by earthbound scientists. They suggest that not only would this be cheaper and less prone to fatal error, but it would be more efficient as a whole host of explorers could venture out on the surface of Mars [McG91, Car].

Taking a different approach, Kellogg, Carroll and Richards imagine that instead of just creating virtual worlds that are separate from the real world, there might also be a virtual world overlaying our own where everyday objects have extra informatic abilities [KCR91]. One example they give is of an augmented kitchen that amongst other things has a cupboard that knows what it contains and that can communicate the lack of an item required by an agent that is planning a meal to an agent that prepares a shopping list.

Bibliography

[ALK+92] P.A. Appino, J.B. Lewis, L. Koved, D.T. Ling, D.A. Rabenhurst, and C.F. Codella. An architecture for virtual worlds. *Presence: Teleoperators and Virtual Environments*, 1(1), 1992.

[All91] Earl A. Alluisi. The development of technology for collective training: Simnet, a case history. *Human Factors*, 33(3):343–363, 1991.

[And] Zap Andersson. A 'realistic' cyberspace scenario. Posting to Usenet *sci.virtual-worlds* newsgroup. E-mail: zap@lage.lysator.liu.se.

[ARB90] J.M. Airey, J.H. Rohlf, and F.P. Brooks Jr. Towards image realism with interactive update rates in computer graphics. *Computer Graphics*, 24(2):41–50, 1990. Special Issue on Symposium on Interactive 3D Graphics.

[Bat92] Joseph Bates. Virtual reality, art and entertainment. *Presence: Teleoperators and Virtual Environments*, 1(1):133–138, 1992.

[BBH+90] Chuck Blanchard, Scott Burgess, Young Harvill, Jaron Lanier, Ann Lasko, Mark Oberman, and Michael Teitel. Reality built for two: A virtual reality tool. *Computer Graphics*, 24(2):35–36, March 1990. Proceedings of the 1990 Symposium on Interactive 3D Graphics (Snowbird, Utah, March 1990).

[BC93] William Bricken and Geoffrey Coco. The VEOS project. Technical Report R-93-3, Human Interface Technology Laboratory, University of Washington, 1993.

[BDHO92] Jeff Butterworth, Andrew Davidson, Stephen Hence, and T.Marc Olano. 3DM: A three dimensional modeller using a head mounted display. *Computer Graphics. Special Issue 1992 Symposium on Interactive 3D Graphics*, pages 135–138, 1992.

[Ben91a] Michael Benedkit. Cyberspace: Some proposals. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[Ben91b] Michael Benedkit. Introduction. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[BF92] G. Bishop and H. Fuchs co-chairs. Research directions in virtual environments: Report of an NSF invitational workshop. *Computer Graphics*, 26(3):153–177, August 1992.

[BFO92] M. Bajura, H. Fuchs, and R. Ohbuchi. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Computer Graphics*, 26(2):203–210, July 1992.

[BHML92] B. Blau, C.E. Hughes, J.M. Moshell, and C. Lisle. Networked virtual environments. *Computer Graphics. Special Issue 1992 Symposium on Interactive 3D Graphics*, pages 157–160, 1992.

[BHV92] Klaus Böhm, Wolfgang Hübner, and Kaisa Väänänen. Given: gesture driven interactions in virtual environments. A toolkit approach to 3D interactions. In *Informatique 92: Interface to Real and Virtual Worlds, Montpellier France, 23-27 March*, pages 243–254, 1992.

[Bie86] Eric A. Bier. Skitters and jacks: Interactive 3D positioning tools. In F. Crone and S. M. Pizer, editors, *Proc. 1986 ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 22-24*, pages 183–196. ACM Press, 1986.

[Bie90] Eric A. Bier. Snap-dragging in three dimensions. *Computer Graphics*, 24(2):193–203, 1990. Special Issue on Symposium on Interactive 3D Graphics.

[BL92] S. Bryson and C. Leith. The virtual wind tunnel. *IEEE Computer Graphics and Applications*, pages 25–34, July 1992.

[BLR91] J. Bates, B. Loyall, and W.S. Reilly. Broad agents. *SIGART Bulletin*, 2(4):38–40, 1991.

[BMB86] N.I. Badler, K.H. Manoochechri, and David Baraff. Multi-dimensional input techniques and articulated figure positioning by multiple contraints. In F. Crone and S. M. Pizer, editors, *Proc. 1986 ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 22-24*, pages 151–169. ACM Press, 1986.

[BOYBK90] Frederick P. Brooks, Jr., Ming Ouh-Young, James J. Batter, and P. Jerome Kilpatrick. Project GROPE — haptic displays for scientific visualisation. *Computer Graphics*, 24(4):177–185, March 1990.

[Bri] William Bricken. Virtual reality: Directions of growth. notes from the SIGGRAPH '90 panel. Posting to Usenet *sci.virtual-worlds* newsgroup. William Bricken, Human Interface Technology Laboratory, University of Washington, FU-20, Seattle, WA 98125, U.S.A. E-mail: william@hitl.vrnet.washington.edu.

[Bri91a] Meredith Bricken. Virtual reality learning environment. *Computer Graphics*, 25(3):178–184, 1991.

[Bri91b] Meredith Bricken. Virtual worlds: No interface to design. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[Bri93] Louis M. Brill. Facing interface issues. *Computer Graphics World*, 15(4), April 93.

[Bro77] Frederick P. Brooks Jr. The computer 'scientist' as toolsmith-studies in interactive computer graphics. In B. Gilchrist, editor, *IFIP Congress Proceedings*, pages 625–634. North-Holland, Amsterdam, 1977.

[Bro86] Frederick P. Brooks Jr. Walkthrough - a dynamic system for simulating virtual buildings. *Proc. 1986 ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 22-24*, pages 9–21, 1986.

[Bro88] Frederick P. Brooks Jr. Grasping reality through illusion. In *Proceedings of CHI, Reading, MA*, pages 1–12. Addison Wesley, 1988.

[BRS+92] G. Burdea, E. Roskos, D. Silver, François Thibaud, and Robert Wolpov. A distributed virtual environment with dextrous force feedback. In *Informatique 92: Interface to Real and Virtual Worlds, Montpellier France, 23-27 March*, pages 255–265, 1992.

[BZR+92] G. Burdea, J. Zhuang, E. Roskos, D. Silver, and N. Langrana. A portable dextrous master with force feedback. *Presence: Teleoperators and Virtual Environments*, 1(1), 1992.

[Car] Shawn Carlson. Virtual mars? Posting to Usenet *sci.virtual-worlds* newsgroup. Shawn Carlson, 50/232, Lawrence Berkeley Laboratory, Berkeley, CA 94720, U.S.A. E-mail: scarlson@csa1.lbl.gov.

[CHB+89] J.C. Chung, M.R. Harris, F.P. Brooks, H. Fuchs, M.T. Kelley, J. Hughes, M. Ouh-Young, C. Cheung, R.L. Holloway, and M. Pique. Exploring virtual worlds with head-mounted displays. *SPIE Vol. 1083 Three-Dimensional Visualization and Display Technologies*, pages 42–52, 1989.

[CJK+92] C. Codella, R. Jalili, L. Koved, B. Lewis, D.T. Ling, J.S. Lipscomb, D.A. Rabenhorst, C.P Wang, A. Norton, P. Sweeney, and G. Turk. Interactive simulation in a mutli-person virtual world. In *Proc CHI'92, May 3-7, 1992*, pages 329–334, 1992.

[CMS88] Michael Chen, S. Joy Mountford, and Abigail Sellen. A study in interactive 3-D rotation using 2-D devices. *Computer Graphics*, 22(4):121–129, August 1988.

[Coo92] Rick Cook. Serious entertainment. *Computer Graphics World*, 15(5):40–41,43,45–48, May 1992.

[CSD+92] C. Cruz-Neira, D.J. Sandin, T.A. DeFanti, R.V. Kenyson, and J.C. Hart. Cave. *Communications of the ACM*, 35(6):65–72, June 1992.

[CW92] Dale Chapman and Colin Ware. Manipulating the future: Predictor based feedback for velocity control invirtual environment navigation. *Computer Graphics. Special Issue 1992 Symposium on Interactive 3D Graphics*, pages 63–66, 1992.

[Dee92] Michael Deering. High resolution virtual reality. *Computer Graphics*, 26(2):195–202, July 1992.

[Dim93] Dimension International. Superscape product literature, 1993. Dimension International, Zephyr One, Calleva Park, Aldermaston, Berkshire, England, RG7 4QZ.

[Div91] Division Ltd. Provision software overview. Technical Specification, 1991. Division Ltd., Quarry Road, Chipping Sodbury, Bristol BS17 6AX.

[DRP+92] N.I. Durlach, A. Rigopulos, X.D. Pang, E.m. Wenzel, W.S. Woods, A. Kulkarni, and H.S. Colburn. On the externalization of auditory images. *Presence: Teleoperators and Virtual Environments*, 1(2):251–257, 1992.

[DZW+92] H. Das, H. Zak, W.S.Kim, A.K. Bejczy, and P.S. Schenker. Operator performance with alternative manual control modes in teleoperation. *Presence: Teleoperators and Virtual Environments*, 1(2):201–218, 1992.

[Ell90] S.R. Ellis. Pictorial communications: Pictures and the synthetic universe. *Leonardo*, 23(1):81–86, 1990.

[Ell92] S.R. Ellis. The design on virtual spaces and virtual environments. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 536–540, 1992.

[Emm92] Arielle Emmett. Down to earth. *Computer Graphics World*, 15(3), March 92.

[EPC93] G.K. Edgar, J.C.D. Pope, and I. Craig. Visual accommodation with virtual images. In *Applied Vision Association Annual General Meeting, Simulated and Virtual Realities*, page 12, 1993.

[Fah91] Lennart E. Fahlén. The MultiG telepresence system. In Yngve Sundblad, editor, *Proceedings of the 3rd MultiG Workshop, KTH, Stockholm, December 17, 1991*, pages 33–57, 1991.

[Fah93] Lennart E. Fahlén. Virtual reality and the MultiG project. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*, pages 78–86, 1993.

[FB90] S. Feiner and C. Beshers. Visualising n-dimensional virtual worlds with n-vision. *Computer Graphics*, 24(2):37–38, March 1990. Proceedings of the 1990 Symposium on Interactive 3D Graphics (Snowbird, Utah, March 1990).

[Fis90] Scott S. Fisher. Virtual interface environments. In Brenda Laurel, editor, *The Art of Human-Computer Interface Design*. Addison Wesley, 1990.

[FMHR86] S.S. Fisher, M. McGreevy, J. Humphries, and W. Robinett. Virtual environment display system. *Proc. 1986 ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 22-24*, pages 77–87, 1986.

[Fol87] James D. Foley. Interfaces for advanced computing. *Scientific American*, pages 83–90, October 1987.

[FS92] K.M. Fairchild and L. Serra. Dynamic fish eye visualisation. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[FST92] Thomas A. Funkhouser, Carlo H. Sequin, and Seth J. Teller. Management of large amounts of data in interactive building walkthroughs. *Computer Graphics. Special Issue 1992 Symposium on Interactive 3D Graphics*, pages 11–20, 1992.

[Fur88] Thomas A. Furness III. Harnessing virtual space. *SID 88 Digest*, pages 4–7, 1988.

[FvDFH90] James D. Foley, Andries van Dam, Steven Feiner, and John Hughes. *Computer Graphics — Principles and Practice.* Addison Wesley, second edition, 1990. ISBN 0-201-12110-7.

[FWCM88] S.S. Fisher, E.M. Wenzel, C. Coler, and M.W. McGreevy. Virtual environment workstations. In *Proceedings of the Human Factors Society 32nd Annual Meeting*, pages 91–95, 1988.

[Gai87] Mike Gaines. Super cockpit. *Flight International*, 131(4052):29–30,32,34, March 1987.

[Gri92a] Jack Grimes. Virtual reality goes commercial with a blast. *IEEE Computer Graphics and Applications*, 12(2):16–17, March 1992.

[Gri92b] C. Grimsdale. Supervision a parallel architecture for virtual reality. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[Hag92] Michael Haggerty. Serious lunacy: Art in virtual worlds. *IEEE Computer Graphics and Applications*, 12(2):4–7, March 1992.

[HD92] R.M. Held and N.I. Durlach. Telepresence. *Presence: Teleoperators and Virtual Environments*, 1(1):109–112, 1992.

[Hee92a] Carrie Heeter. Battletech demographics. *IEEE Computer Graphics and Applications*, 12(2):17, March 1992.

[Hee92b] Carrie Heeter. Being there: The subjective experience of presence. *Presence: Teleoperators and Virtual Environments*, 1(2):262–271, 1992.

[HHK92] M. Hirose, K. Hirota, and R. Kijima. Human behaviour in virtual environments. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 548–559, 1992.

[HY92] M. Hirose and K. Yokoyama. Virtual reality application for transmission of synthetic sensation. In *ICAT 92 Proceedings*, 1992.

[IM92] H. Iwata and K. Matsudi. Haptic walkthrough simulator. In *ICAT 92 Proceedings*, 1992.

[Iwa90] Hiroo Iwata. Artificial reality with force-feedback: Development of desktop virtual space with compact master manipulator. *Computer Graphics*, 24(4):165–170, August 1990. SIGGRAPH '90 Proceedings.

[Kah] David Kahaner. Remarks on VR activities in japan. Posting to Usenet *sci.virtual-worlds* newsgroup. David K. Kahaner, US Office of Naval Research Asia,23-17, 7-chome, Roppongi, Minato-ku, Tokyo 106 Japan.

[Kal92] R.S. Kalawsky. Critical aspects of visually coupled systems. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[Kal93] R.S. Kalawsky. The science and engineering of virtual reality. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*. Meckler, 1993. Not contained in published proceedings.

[KCR91] W.A. Kellogg, J.M. Carroll, and J.T Richards. Making reality a cyberspace. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[KH90] G. Kurtenback and E.A. Hulteen. Gestures in human computer interaction. In Brenda Laurel, editor, *The Art of Human-Computer Interface Design*. Addison Wesley, 1990.

[KOF92] K. Kameyama, K. Ontoni, and Y. Fukui. A virtual reality system using a volume scanning 3D display. In *ICAT 92 Proceedings*, 1992.

[Kru83] Myron W. Krueger. *Artificial Reality*. Addison Wesley, 1983.

[Kru90] Myron W. Krueger. Videoplace and the interface of the future. In Brenda Laurel, editor, *The Art of Human-Computer Interface Design*. Addison Wesley, 1990.

[Kru91] Myron W. Krueger. *Artificial Reality II*. Addison Wesley, second edition, 1991. ISBN 0-201-52260-8.

[KS92] Won S. Kim and Paul Schenker. A teleoperation training simulator with visual and kinesthetic force virtual reality. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 560–569, 1992.

[KYB90] A. Kaufman, R. Yagel, and R. Bakalash. Direct interaction with a 3d volumetric environment. *Computer Graphics*, 24(2):33–34, March 1990. Proceedings of the 1990 Symposium on Interactive 3D Graphics (Snowbird, Utah, March 1990).

[Lau92] Brenda Laurel. Artistic frontiers in virtual reality. *SIGGRAPH 92 Visual Proceedings*, pages 58–63, 1992. ISBN 0-89791-423-7.

[Loo92a] J.M. Loomis. Distal attribution and presence. *Presence: Teleoperators and Virtual Environments*, 1(1):113–119, 1992.

[Loo92b] J.M. Loomis. Presence and distal attribution: Phenomenology, determinants, and assessment. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 590–595, 1992.

[LSG91] J. Liang, C. Shaw, and M. Green. On temporal-spatial realism in the virtual reality environment. In *UIST'91. Proceedings of the 4th Annual Symposium, Hilton Head, South Carolina*, pages 19–25, 1991.

[LTS92] Andrew Liu, Gregory Tharp, and Lawrence Stark. Depth cue interaction in telepresence and simulated telemanipulation. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 541–547, 1992.

[LW90] M. Levoy and R. Whitaker. Gaze-directed volume rendering. *Computer Graphics*, 24(2):217–223, 1990. Special Issue on Symposium on Interactive 3D Graphics.

[MAB92] K. Meyer, H.L. Applewhite, and F.A. Biocca. A survey of position trackers. *Presence: Teleoperators and Virtual Environments*, 1(2):173–200, 1992.

[Mar92] Elizabeth Marcus. Feedback technology and virtual environments. In *ICAT 92 Proceedings*, 1992.

[McG91] Michael W. McGreevy. Virtual reality and planetary exploration. In *29th AAS Goddard Memorial Symposium, Washinton D.C., March*, pages 93–115, 1991.

[McK92] Michael McKenna. Interactive viewpoint control and three-dimensional operations. *Computer Graphics*, pages 53–56, 1992. Special Issue on Symposium on Interactive 3D Graphics.

[MCR90] J.D. Mackinlay, S.K. Card, and G.G. Robertson. Rapid controlled movement through a virtual 3D workspace. *Computer Graphics*, 24(4):171–176, August 1990.

[ME90] Philip J. Mercurio and Thomas D. Erickson. Interactive scientific visualisation: An assessment of a virtual reality system. In *Human-Computer Interaction - INTERACT'90*, pages 741–745. Elsevier Science Publishers B.V. (North-Holland), 1990.

[MF91] Chip Morningstar and F. Randall Farmer. The lessons of lucasfilm's habitat. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[Mon92] G.J. Monkmann. An electorheological tactile display. *Presence: Teleoperators and Virtual Environments*, 1(2):219–228, 1992.

[Mos93] Michael Moshell. Virtual environments in the US military. *Computer*, 26(2):81–82, February 1993.

[MOYS+90] M. Minsky, M. Ouh-Young, O. Steele, F.P. Brooks Jr., and M. Behensky. Feeling and seeing: Issues in force display. *Computer Graphics*, 24(2):235–243, 1990. Special Issue on Symposium on Interactive 3D Graphics.

[MP92] Gloria Molendi and Matteo Patriarca. Virtual reality: Medical researches. Technical Report 1/92, Universita' degli Studi di Milano, Dipartmento di Scienze della Informazione, December 1992.

[MPZ90] Michael McKenna, Steve Pieper, and David Zeltzer. Control of a virtual actor: The roach. *Computer Graphics*, 24(2):29–37, March 1990.

[MS91] Michael E. McCauley and Thomas J. Sharkey. Spatial orientation in virtual reality systems: Lessons from flight simulation. In *Proceedings of the Human Factors Society 35th Annual Meeting*, pages 1348–1352, 1991.

[MZ90] Michael McKenna and David Zeltzer. Dynamic simulation of autonomous legged motion. *Computer Graphics*, 24(4):29–37, August 1990.

[NO86] Gregory M. Nielson and Dan R. Olsen Jr. Direct manipulation techniques for 3D objects using 2D locator devices. In F. Crone and S. M. Pizer, editors, *Proc. 1986 ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 22-24*, pages 175–182. ACM Press, 1986.

[Nov91] Marcus Novak. Liquid architectures in cyberspace. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[Pat92a] Robert Patterson. Factors that affect depth perception in stereoscopic displays. *Human Factors*, 34(6):655–667, 1992.

[Pat92b] Robert Patterson. Human stereopsis. *Human Factors*, 34(6):669–692, 1992.

[Pau91] Randy Pausch. Virtual reality on $5 a day. In *Proceedings of CHI'91*, pages 265–270, 1991.

[PB91] Steve Pruitt and Tom Barrett. Corporate virtual workspace. In Michael Benedkit, editor, *Cyberspace: First Steps*. The MIT Press, 1991.

[PBG92] C.B. Phillips, N.I Badler, and J. Granieri. Automatic viewing control for 3D direct manipulation. *Computer Graphics. Special Issue 1992 Symposium on Interactive 3D Graphics*, pages 71–74, 1992.

[PBL+92] T.P. Piantanida, D. Boman, J. Larimer, J. Gille, and C. Reed. Studies in the field-of-view/resolution tradeoff in virtual-reality systems. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 448–456, 1992.

[Pen90] Alex P. Pentland. Computational complexity versus simulated environments. *Computer Graphics*, 22(2), 1990. Proceedings of the 1990 Symposium on Interactive 3D Graphics (Snowbird, Utah, March 1990).

[PG92] M.J. Papper and M.A. Gigante. Using physical constraints in a virtual environment. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[PT93] Ken Pimentel and Kevin Teixeira. *Virtual Reality: Through the New Looking Glass*. McGraw-Hill, 1993.

[RCM89] G.G. Robertson, S.K. Card, and J.D. Mackinlay. The cognitive coprocessor architecture for interactive user interfaces. In *Proceedings UIST'89*, pages 10–18, 1989.

[RCM93] George G. Robertson, Stuart K. Card, and Jock D. Mackinlay. Information visualisation usind 3D interactive animation. *Communications of the ACM*, 36(4):57–71, 1993.

[Rhe91] Howard Rheingold. *Virtual Reality*. Secker & Warburg, London, 1991.

[Rob92] W. Robinett. Synthetic experience: A proposed taxonomy. *Presence: Teleoperators and Virtual Environments*, 1(2):229–247, 1992.

[Row92] T.W. Rowley. Virtual reality products. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[RR92] W. Robbinett and J.P Rolland. A computational model for the stereoscopic optics of a head mounted display. *Presence: Teleoperators and Virtual Environments*, 1(1), 1992.

[RW93] Simon Rushton and John Wann. Problems in perception and action in virtual worlds. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*, pages 43–55. Meckler, 1993.

[Sat93a] Richard M. Satava. Surgery 2001. a technological framework for the future. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*, pages 101–103. Meckler, 1993.

[Sat93b] Richard M. Satava. Virtual reality surgical simulator. first steps. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*, pages 103–105. Meckler, 1993.

[SD91] Mel Slater and Allan Davidson. Liberation from flatland: 3D interaction based on the desktop bat. In *EUROGRAPHICS '91*, pages 209–221. Elsevier Science Publishers B.V. (North-Holland), 1991.

[She92a] T.B. Sheridan. Defining our terms. *Presence: Teleoperators and Virtual Environments*, 1(2):272–274, 1992.

[She92b] T.B. Sheridan. Musings on telepresence and virtual presence. *Presence: Teleoperators and Virtual Environments*, 1(1):120–126, 1992.

[SIG89] SIGGRAPH'89. Virtual environments and interactivity. *Computer Graphics*, 23(5):7–38, 1989. Panel Proceddings.

[SIG90] SIGGRAPH'90. Hip, hype and hope-the three faces of virtual worlds. *Computer Graphics*, pages 10.1–10.29, 1990. Panel Proceddings.

[Sla93] Mel Slater. Equivalence classes for presence in virtual environments. Discussion Note, April 24, 1993. Department of Computer Science, Queen Mary and Westfield College, University of London.

[SLGS92] Chris Shaw, Jiandon Liang, Mark Green, and Yunqi Sun. The decoupled simulation model for virtual reality systems. In *Proc CHI'92, May 3-7, 1992*, pages 321–328, 1992.

[SM91] Randy L. Sollenberger and Paul Milgram. A comparative study of rotational and stereoscopic computer graphic depth cues. In *Proceedings of the Human Factors Society 35th Annual Meeting*, pages 1452–1456, 1991.

[Sno] Michael Snoswell. Overview of cyberterm, a cyberspace protocol implementation. Posting to Usenet *sci.virtual-worlds* newsgroup. Michael Snoswell, Vision Systems Limited, Second Avenue, Technology Park, The Levels, 5095, South Austrailia.

[Spe92] T.H. Speeter. Transforming human hand motion for telemanipulation. *Presence: Teleoperators and Virtual Environments*, 1(1), 1992.

[Sta91] Micahel Starks. Stereoscopic video and the quest for vitual reality. In Sandra K. Helsel, editor, *Beyond the Vision: Proceedings of the Second Annual Conference on Virtual Reality, Artificial Reality and Cyberspace (San Francisco, September 1991)*, pages 200–225. Meckler, 1991.

[Sto92] R. Stone. Virtual reality: A tool for telepresence and human factors research. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[SU93a] Mel Slater and Martin Usoh. The influence of a virtual body on presence in immersive virtual environments. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*, pages 34–42. Meckler, 1993.

[SU93b] Mel Slater and Martin Usoh. Representation systems, perceptaul position and presence in virtual environments, 1993. Department of Computer Science, Queen Mary and Westfield College, University of London.

[SU92] Mel Slater and Martin Usoh. An experimental exploration of presence. Submitted for publication, 92. Department of Computer Science, Queen Mary and Westfield College, University of London.

[Sut65] Ivan E. Sutherland. The ultimate display. In *IFIP Congress Proceedings*, volume 2, pages 506–508, 1965.

[Sut68] Ivan E. Sutherland. A head-mounted three dimensional display. In *Proc. AFIPS Fall Joint Computer Conference, 33*, pages 757–764, 1968.

[SWH93] David N. Snowdon, Adrian J. West, and Toby L.J. Howard. Towards the next generation of human-computer interface. In *Informatique 93: Interface to Real and Virtual Worlds, Montpellier France, 24-26 March*, 1993.

[SZP89] D.J. Sturman, D. Zeltzer, and S. Pieper. Hands-on interaction with virtual environments. In *Proceedings of the ACM SIGGRAPH Symposium on User Interface Software and Technology, Williamsburg, Virginia, November 13-15*, pages 19–24, 1989.

[TH92] T. Takala and J. Hahn. Sound rendering. *Computer Graphics*, 26(2):211–220, July 1992.

[TLF+92] Gregory Tharp, Andrew Liu, Lloyd French, Steve Lai, and Lawrence Stark. Timing considerations of helmet mounted display performance. In *SPIE Vol. 1666 Human Vision, Visual Processing, and Digital Display III*, pages 570–576, 1992.

[TS91] Seth J. Teller and Carlo H. Sequin. Visibility preprocessing for interative walkthroughs. *Computer Graphics (SIGGRAPH '91 Proceedings)*, 25(4):61–69, July 1991.

[Utt89] William R. Uttal. Teleoperators. *Scientific American*, pages 54–79, December 1989.

[VB92] K. Väänären and K. Böhm. Gesture driven interaction as a human factor in virtual environments- an approach with neural nets. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[vdG86] W.A. van de Grind. Vision and the graphical simulation of spatial structure. *Proc. 1986 ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 22-24*, pages 197–235, 1986.

[Vin92] J. Vince. Virtual reality techniques in flight simulation. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[Vir90] Virtus Corporation. Virtus walkthrough demo 1.3, 1990. Virtus Corporation, 117 Edinburgh South, Suite 204, Cary NC 27511. Email virtus@applelink.apple.com.

[W I] W Industries. Without walls, virtuality leisure magazine, autumn 1992. W Industries Ltd, Virtuality House, 3 Oswin Road, Brailford Industrial Park, Leicester, LE3 1HR, UK.

[Wal90] John Walker. Through the looking glass. In Brenda Laurel, editor, *The Art of Human-Computer Interface Design*. Addison Wesley, 1990.

[Wal93] Johnathan D. Waldern. Virtual reality - the next generation. In *Virtual Reality International 93, Proceedings of the third annual conference on Virtual Reality, London, April 1993*, pages 87–100. Meckler, 1993.

[War90] Colin Ware. Using hand position for virtual object placement. *Visual Computer*, (6):245–253, 1990.

[WCF90] J. Wang, V. Chi, and H. Fuchs. A real-time optical 3D tracker for head-mounted display systems. *Computer Graphics*, 24(2):205–215, 1990. Proc. 1990 Symposium on Interactive 3D Graphics.

[Wen92] E.M. Wenzel. Localization in virtual acoustic displays. *Presence: Teleoperators and Virtual Environments*, 1(1):80–107, 1992.

[WHH92] A.J. West, T.L.J. Howard, and R.J Hubbold. AVIARY - a generic virtual reality interface. In *Virtual Reality Systems, International State-of-the-Art Conference, University of London, 20-21 May 1992*, 1992.

[WJ88] Colin Ware and Danny R. Jessome. Using the bat: A six-dimensional mouse for object placement. *IEEE Computer Graphics and Applications*, pages 65–70, November 1988.

[WKD90] C. Wang, L. Koved, and S. Dukach. Design for interactive performance in a virtual laboratory. *Computer Graphics*, 24(4):39–40, March 1990.

[WO90] Colin Ware and Steven Osborne. Exploration and virtual camera control in virtual three dimensional environments. *Computer Graphics*, 24(2):175–183, March 1990. Proceedings of the 1990 Symposium on Interactive 3D Graphics (Snowbird, Utah, March 1990).

[Zel91] D. Zeltzer. Task-level graphical simulation: Abstraction, representation and control. In D. Zeltzer N.I. Badler, B.A. Barsky, editor, *Making them Move. Mechanics, Control and Animation of Articulated Figures*. Morgan Kaufmann, 1991.

[Zel92] David Zeltzer. Autonomy, interaction and presence. *Presence: Teleoperators and Virtual Environments*, 1(1):127–132, 1992.

[ZM91] D. Zeltzer and M.B.Johnson. Motor planning: An architecture for specifying and controlling the behaviour of virtual actors. *The Journal of Visualisation and Computer Animation*, 2(2):74–80, 1991.

[ZPMW92] M.J. Zyda, D.R. Pratt, J.G. Monaham, and K.P. Wilson. Nspnet: Constructing a 3d world. *Computer Graphics. Special Issue 1992 Symposium on Interactive 3D Graphics*, pages 147–156, 1992.