Head-Mounted Display Technical Report

TR87-015

June 1987

Richard L. Holloway

The University of North Carolina at Chapel Hill
Department of Computer Science
CB#3175, Sitterson Hall
Chapel Hill, NC 27599-3175

This research was supported by NIH grants #RR-02170-04 and #RR-02170-05

Abstract

There are clearly many potential advantages to a head-mounted display system: the 3-D cues and freedom of movement through large models are unparallelled by any existing system.

This report presents a study of the 1986 system at the University of North Carolina at Chapel Hill and lists areas of success as well as those that need improvement.

Finally, the report discusses both predicted and planned improvements to the 1986 system.

Contents

Introduction 2

Chapter 1 : Advantages and uses of a Head-Mounted Display 3

Chapter 2 : The 1986 implementation at UNC 5

Section 1 - Configuration 5
1.1 Head Unit 6
1.1.1 Configuration 6
1.1.2 Adjustments 8
1.1.3 Issues 9
1.2 Host Computer 10
1.3 Graphics 11
1.3.1 E&S PS-330 11
1.3.2 Video Cameras 12
1.3.3 LCD Screens & Optics 12
1.4 Polhemus Tracker 13
1.4.1 Hardware 13
1.4.2 Software 14
1.4.3 Issues 16

Section 2 - Successes and failings of the 1986 System 17
2.1 Successes 17
2.2 Failings 17

Chapter 3 : Future systems and improvements 20

Section 1 - Head Unit 20
Section 2 - Host Computer 21
Section 3 - Graphics 21
Section 4 - Tracker 22

Bibliography 23

Introduction

The purpose of this report is to provide a comprehensive description of existing and future implementations of head-mounted display systems at the University of North Carolina at Chapel Hill. This is not a user's manual, but rather a description of the essential design and implementation issues for a head-mounted display system. Each chapter is divided into sections according to the main components of a head-mounted display system : head unit (the part of the system actually worn by the user), host computer, graphics (hardware and software) and tracker (tracks the user's position and orientation). The document is structured hierarchically, with general information at the higher levels and details at the lower levels. The reader seeking a general understanding of our head-mounted display system may omit sub-sections.

The reader is assumed to have some knowledge of programming and of computer graphics.

This work was supported by NIH Grants RR-02170-04 and RR-02170-05.

The hardware and software for the system was done by : Mike Pique, Richard Holloway, John Hughes, Jim Chung, Ming Ouh-Young, Helga Thorvaldsdottir, Mark Harris, and the UNC Physics Department Machine Shop.

Chapter 1

Potential Advantages of a Head-Mounted Display System

Computer graphics aids in visualizing complex systems, from designed but yet unbuilt objects to large protein molecules. Often the object to be viewed is inherently three-dimensional and cannot be represented adequately by a two-dimensional graphics screen. There are many systems in existence that attempt to give the viewer 3-D cues so that he can "see" a 3-D object on a two-dimensional screen or surface. These cues include :

• Obscuration - A close object obscures another that it is farther away.

• Head-motion parallax (kinetic depth effect) - The change in image when the viewer moves his head to see a different view of an object.

• Stereo - The viewer gets one of the strongest 3-D cues from the slightly disparate images seen by each eye. The brain integrates the two images and fuses them into a single 3-D image.

• Shading - The way light falls on an object is a cue to its shape and orientation.

• Intensity depth-cueing - In this representation, parts of the object that are farther away are dimmer, as if object is in fog.

• Perspective - Scaling an object's height and width as a function of its distance from the viewpoint (eg, receding railroad tracks).

A head-mounted display system can give the user all of these cues as well as allowing the user to move around in a virtual environment.

A head-mounted display (or HMD) is a device worn on the head with screens in front of the user's eyes that present images to a user according to his position and the direction in which his head is aimed. The stereo effect is achieved by displaying a different image for each eye. Head-motion parallax is achieved by tracking the user's head motion and updating the images accordingly. The other cues are created by the computer image-generation hardware and software.

With a head-mounted display, the user is immersed in a virtual environment in which he can move around and examine in detail any object he wishes. Whereas some graphics systems use joysticks or dials to change the user's point of view, a head-mounted display allows the user to change his viewpoint in the most natural way possible: if he wants a closer look at something, he simply walks toward it and moves around it until he sees the point of interest.

Chapter 2

The 1986 Head-Mounted Display System at the University of North Carolina at Chapel Hill

At UNC-Chapel Hill, we have implemented a head-mounted display system with current (market) technology with the design goals that it be as effective and unobtrusive as possible. This chapter describes the 1986 system in detail, followed by a brief discussion of its good points as well as its failings.

Section 1 - 1986 Configuration

The 1986 system at UNC has four main components (see Figure 2-1) :

• Head Unit - This is the actual device the user wears. The 1986 version is a welder's visor which has the screens, optics and the tracker connected to it (see Figure 2-2).

• Host Computer - This is the computer that coordinates and synchronizes the activities of all parts of the HMD system; here we use a Masscomp MC5500.

• Graphics - This is the part of the system (hardware and software) which creates and transmits the images to the user. We use an Evans & Sutherland PS-330 to generate the image, cameras to convert the image to an NTSC signal, and Casio LCD TV screens to display the final image.

• Tracker - This is the hardware and software for detecting the user's position and orientation. We use a Polhemus Navigation Systems 3-Space™ Isotrak™ (Polhemus [3]) tracker, which tracks position and orientation (but does not track eye movement).

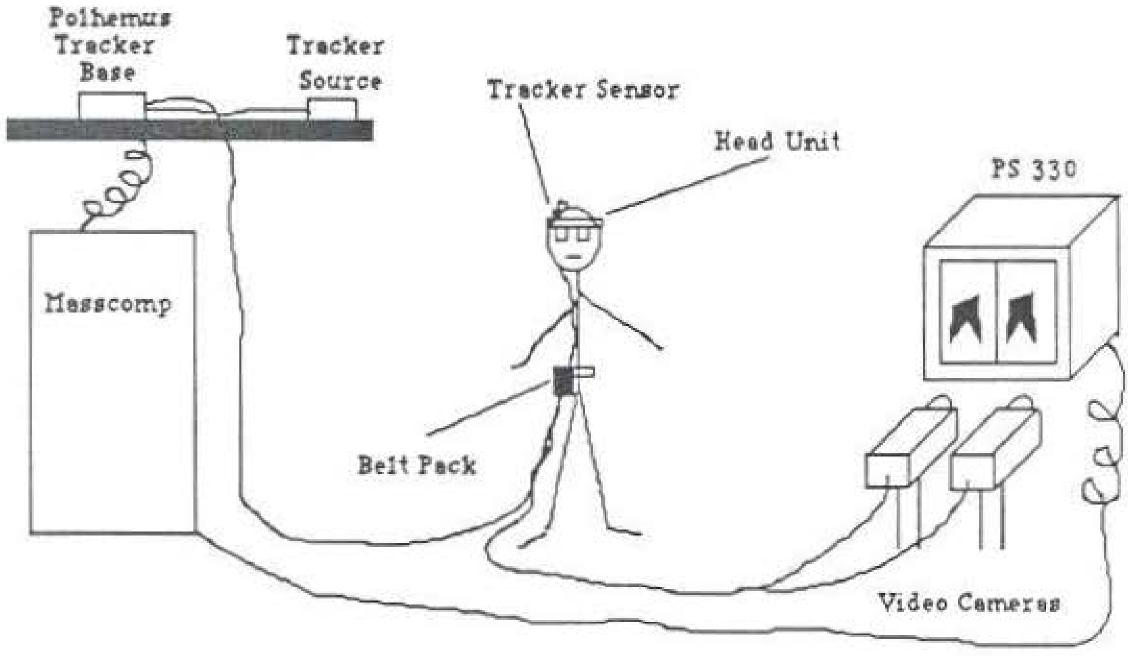

Figure 2-1 : An overview of the UNC HMD System

Now, these four components in more detail.

1.1 Head Unit

1.1.1 Configuration

The 1986 head unit is made up of two parts : the head-unit proper and the auxiliary belt pack. The helmet carries all of the essential parts discussed in Chapter 1 : the tracker sensor, the image display screens, and the necessary optics. The belt pack carries the heavier parts of the image display hardware and is where the cables from the cameras are attached (see Figure 2-1).

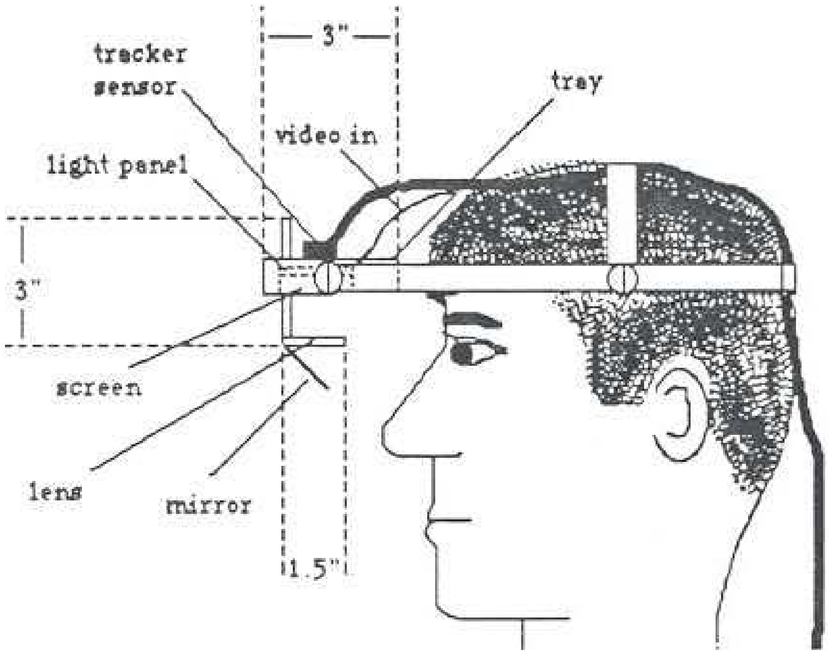

Figure 2-2: The Head Unit

As shown in Figure 2-2, the main physical structure is similar to a welder's visor. The "tray" that projects horizontally from the forehead holds the two LCD screens and is where the Polhemus tracker sensor is attached. The image is projected from the screens downward, through a lens and then reflects off the angled half-silvered mirror into the view of the user. Since the mirror is half-silvered, the user sees the virtual environment superimposed on the real one.

• Tracker - As shown in Figure 2-2, the tracker sensor is attached to the tray, with its cord running along the visor and down past the belt pack back to the main tracker unit (described in Section 4).

• LCD Screens - There are two screens (one for each eye) on which the final image is displayed. These screens are 139 x 110 arrays of monochrome LCDs measuring 2 inches diagonally. The light through the LCD's is supplied by a 6-volt, 16 foot-Lambert, Luminescent Systems flat luminescent panel. The screens are extended from a pair of commercially available Casio miniature television sets. The hardware to take the video camera signal and put it onto the LCD screens is actually in two parts.

The screens themselves and some of the associated hardware are mounted on the tray on the head unit, while the heavier parts of the TV's are kept in the belt pack. The connection between them is made via a ribbon cable.

• Magnifying Lenses - Two Bausch & Lomb 3.5" focal length, 1" x 1.5", doubly convex lenses. These enable the user to focus at infinity, since the screens are too close for normal focusing.

• Half-silvered Mirror - We use two 2"x2" Ealing Optics neutral density filters which reflect 80% and transmit 20% of incident light.

1.1.2 Adjustments

What makes the HMD system different from other software and hardware undertakings is that the user actually wears part of this system. Adjustments are necessary so that different users can wear the device comfortably, with the images in their line of vision :

• **The "Hat-Band" Adjustment** - This is accomplished by a "hat-band" snap similar to those found on baseball caps.

• **Vertical Adjustment of the Image** - Since each person's face is proportioned differently, an adjustment is needed to allow the image to be moved up or down with regard to the user's line of vision. The lens/mirror assembly (shown in Figure 2-2) slides up and down (2.5" of travel), held in place by friction, to do this.

• **Inter-ocular Distance Adjustment** - Since the distance between the pupils ranges from 2.2" to 2.8" (Diffrient [9]), an adjustment is needed for this. On this system, the adjustment is rather crude, since the lens/mirror assembly may move horizontally, but the LCD screens remain fixed. Thus, only minor adjustments (maximum 0.25") can be made since any major adjustment will move the lens/mirror assembly out of line with the LCD screens.

• **Tray Adjustment** - The tray (see Figure 2-2) is also movable to facilitate the above adjustments. It can be rotated slightly, as well as moved up and down to bring the mirrors into the user's line of vision.

1.1.3 Issues

One of the primary requirements for any head unit is that it be unobtrusive, since it may be worn for hours at a time. Here are some of the issues considered in creating the 1986 implementation:

• **Weight** - A heavy device could cause rapid fatigue and possibly neck strain. Therefore, one of the primary concerns in designing the1986 head unit was its weight. All materials for the head unit were chosen not only for function, but also for weight and wieldiness. Moreover, many of the heavier parts of the head-unit hardware were moved to the belt pack to keep the weight of the head unit down.

• **Lack of enclosure** - A device that shuts out the rest of the world may facilitate concentration, but it can also induce claustrophobia and disorientation, particularly for a user who is moving around. The concept of allowing for a "room view" comes from the desire to be unobtrusive. We believe that if the user can see what is going on around him, he will feel less enclosed by the head unit and won't feel such an urge to be free of it. He is less apt to fall over real objects. Moreover, if he is able to see the rest of his environment, it may be possible for him to do other work briefly, without having to remove the head unit and then put it back on. Finally, and most important, we suspect that superimposing an unknown structure into a familiar world-model, such as the laboratory room, will facilitate the user's formation of an accurate world-model of the unknown structure.

• **Restraints** - Since the current configuration requires cables connecting the user to the rest of the system, the cables must be long and flexible so they do not inhibit movement.

1.2 Host Computer

Synchronizing and coordinating all the tasks of the UNC HMD system is done by a Masscomp 5500 workstation. The main control program runs here and is responsible for polling the tracker and sending the updated position and orientation values to the PS-330. A sketch of the main program follows.

```
main()
{

/* init routine connects to tracker, PS-330                */
init();

/* sets up tracker for communication                       */
trackerInit();

/* infinite loop to track user and run system              */
while (TRUE)
        {

        /* polls tracker, updates position and orientation values
         * in array newCoords (a VAR parameter)
         */
        pollTracker(newCoords);

        /* takes new position and orientation information and sends
         * it to the graphics device
         */
        updateView(newCoords);

        } /* end while loop */

} /* end main program */
```

Figure 2-3 Sketch of the main program for the HMD system

The main program itself is rather simple, reflecting the primary tasks of the system: track and update view. Speed is most essential in doing this, since the user must not notice a lag between the time he changes his viewpoint and the time the image changes, or the illusion ceases to be convincing (see Section 1.4.3 for further discussion).

The Masscomp also performs the calculations for generating the two different eye-points to send to the PS-330. The software for the tracking and graphics will be examined in more detail in sections 1.3 and 1.4.

1.3 Graphics

The graphics system is the most complicated part of the HMD system. It can be viewed as a pipeline, beginning with the transformation matrix (a matrix description of the user's position and orientation generated by the tracker) and ending with the image on the user's retina.

1.3.1 Evans & Sutherland PS-330

The PS-330 is the graphics processor for the system as well as the intermediate graphics display device (the final display is on the LCD screens). It is a vector graphics device, capable of drawing colored lines very quickly, but incapable of rendering shaded polygons, etc. The Masscomp sends the transformation matrices to it via an ethernet (a high-speed communications link). The PS-330 has two main functions : Displaying the image of the object(s) to be represented and updating that image.

• **Loading and displaying the object(s)** - The 1986 system displays static objects from a changing viewpoint. Prior to running the main program, the object description is written in software in a special description language specific to the PS-330 (Evans&Sutherland[1]). This object description is then downloaded from the Masscomp to the PS-330 once over the ethernet. The image is then displayed in an initial position, where it stays until the Masscomp main program sends it viewpoint position and orientation updates. Thereafter, the PS-330 transforms the image appropriately for the new viewpoint.

• **Creating a stereo image** - This is done rather simply on the PS-330 by dividing the screen into windows. Each half of the screen corresponds to an eye, and has an appropriately different viewpoint to create the 3-D illusion.

• **Updating the view of the object(s)** - There are library routines (Evans&Sutherland[2]) that can be called from C programs to communicate with the PS-330. The program that runs on the Masscomp sends the position and orientation updates via these routines over the ethernet. To maintain the 3-D effect, a slightly different update matrix is sent to each of the windows corresponding to the left and right eyes. The transformation matrices for each are calculated as if the user's eyes are looking in the same direction (parallel), but from two slightly different positions (2.5 inches apart), as proposed by Lipscomb [10].

1.3.2 Video Cameras

Once the left and right eye images are displayed on the PS-330, they must be transmitted to the black-and-white LCD screens. This is accomplished in the 1986 system by focusing two TV cameras* on each half of the PS-330 screen and running the video signal from them into the LCD TV hardware in the belt pack. The cameras were necessary, rather than taking the signal straight from the PS-330 into the TV hardware, since the PS-330 does not put out a standard video signal (the PS-330 updates its image at a variable rate whereas standard video uses a fixed update rate). See Figure (2-1) for configuration.

1.3.3 The LCD Screens and Optics

The user sees the final image on the LCD screens described in Section 1.1.1. The image signal is fed into the screens via a ribbon cable from the belt pack. The image from these screens is transmitted to the user's eyes through the lenses and mirrors explained in Section 1.1.1.

* RCA 2000 black-and-white closed-circuit TV cameras, model TC2011.

1.4 Polhemus 3-Space™ Tracker

1.4.1 Hardware

The Polhemus has three main components (see Figure 2-3):

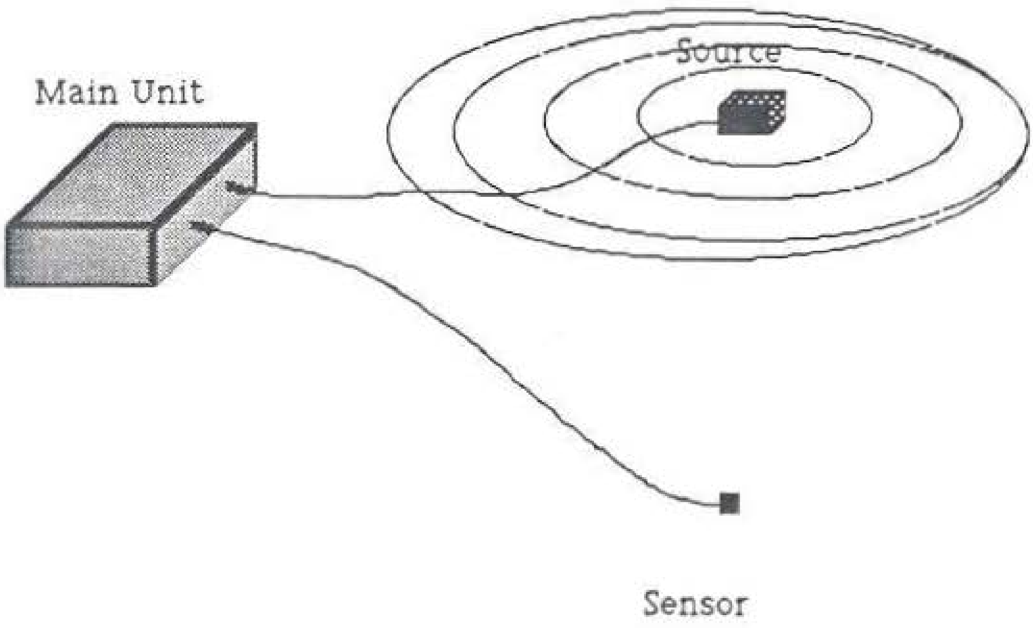

Figure 2-3: The Polhemus 3-Space™ Tracker

• **The Main Unit** - This is the heart of the Polhemus, containing its processor and EEPROM (electrically-eraseable programmable read-only memory). It is connected to a port on the Masscomp via a 9600 baud serial line (RS-232).

• **The Source** - This part remains stationary and serves as the reference point for the sensor. It generates a hemispherical, low-frequency (60 Hz) magnetic field roughly 5 feet in radius, centered on itself. Typically, this is placed 7-8 feet above the floor and the hemisphere is "aimed" down toward the user for maximum working volume. Similarly, the source is positioned 7-8 feet from any wall. Since the Polhemus operates using a magnetic field, care must be taken to keep metal objects as far from the work area as possible, since they may cause field distortion.

• **The Sensor** - This part is attached to the head unit. It is about the size of a thimble and is connected to the main unit via a highly flexible cord. It senses its position within the field emanating from the source and reports this data to the main unit.

• **Specifications** - The Polhemus used in this system has 12 bit precision for its data representation, giving it a resolution of 0.03" for position and 0.1 degree for orientation. The specified accuracy is 0.1" RMS for position and 0.5 degrees RMS for orientation. Specified accuracy is for sensor positions between 4" and 28" from source. Operation is possible up to 60" away from source is possible with reduced accuracy (Polhemus[1]).

The specified update rate is 60 per second. According to engineers at Polhemus Navigational Systems (Anderson[5]), the Polhemus gathers data for 1/60 second, performs calculations on the data for another 1/60 second, and takes a final 1/60 to generate the report. This means that the maximum lag between the request for a report and the generation of that data is 1/20 second.

1.4.2 Software

A simple way to think of the Polhemus is that it is a black box which always knows where the sensor is. When the program needs current position and orientation data, it sends it a request for that information and then reads it from a buffer. For the HMD system, the program does this many times in a single second.

In reality, the Polhemus is a fairly complicated device (for details, see Polhemus[1]). It has several modes and its own command language, not covered in detail here. An overview of the relevant features follows, briefly describing its modes and some commands.

• **Software Connection to Masscomp**

The Polhemus is connected in software by "linking" (in Unix™) the physical port to which it is connected to a virtual device with its name (or something similar). For example, if the RS-232 line is connected to the physical port "tty0", in Unix™, the connection is made via the following command :

ln /dev/tty0 /dev/polhemus

This device can then be configured in software to have the appropriate communications characteristics (baud rate, handshake, etc) via the Unix™ "stty" command. To access the Polhemus, the program calls an I/O routine to get a file descriptor for the Polhemus. In C,

polhemusFD = open("/dev/polhemus", READ_WRITE);

Thereafter, the program can perform C "read"s and "write"s as needed, just as it would an ordinary file.

• Commands

Communication with the Polhemus is done via commands. The program sends a command to the Polhemus whenever it wants to change modes, get a report, or change the Polhemus's field characteristics. The Polhemus reacts to a command by changing mode or returning information via a buffer which the program can read. For the UNC HMD system, the main program typically sends commands to initialize the Polhemus to the desired modes and field characteristics, then it enters a loop in which it only requests and reads position and orientation reports. The commands are short ASCII strings which are written to the Unix™ virtual device described above.

• Modes

The Polhemus has many modes, but only two relevant ones will be described here.

- Binary/ASCII - The Polhemus can return the position and orientation information in either ASCII or binary format. The ASCII is immediately readable but is longer, while the binary is only one-third as long, but requires massaging before it is useful.

Binary mode is used in the 1986 system because the data transmission time is critical.

• Polling/Continuous - In polling mode, the main program sends a request to the Polhemus when it wants a position update. This is easy to synchronize, but slower, since there is a lag due to the time it takes to send the request and wait for the buffer to fill. The 1986 system uses polling mode due to the simplicity of implementation.

In continuous mode, the Polhemus generates position and orientation reports as fast as it can and sends them to the Masscomp. This is faster, but it makes synchronizing the Polhemus and the Masscomp more difficult, since the Polhemus generates reports asynchronously from the host.

1.4.3 Issues

The primary issue in dealing with the Polhemus is speed. In this system, the lag between a user's movement and the update of the image must be minimized. We do not have data on what lag is perceptible. It surely cannot be longer than 1/20 of a second, as James Perkins, McDonnell-Douglas Aircraft Corp, reports that in aircraft flight simulators an update cycle of 50 msec. (20 times/sec) is satisfactory, whereas pilots begin to complain when the cycle is lengthened even to 55 msec. An update rate of 20 times/sec is close to the maximum that we have been able to obtain with the Polhemus.

In polling mode, however, it is possible to request updates faster than the Polhemus can generate them. Therefore, the polling program must wait a fixed period of time for the Polhemus to get a complete report and ensure the accuracy of the data in the buffer. This decreases the update rate since the program must wait even when the Polhemus generates a report quickly.

Section 2 - Successes and Failures of the 1986 UNC Head-Mounted Display System

2.1 Successes

The 1986 system has several successes:

• Weight

The 1986 system is quite light and easy to wear, in contrast with models using motorcycle helmets as the main structure.

• Room View

This system is unusual in allowing the user to see the virtual scene imposed on the real room view.

• Openness

Since the head unit doesn't cover the user's face, the 1986 model doesn't get hot inside or induce claustrophobia.

2.2 Failings

The 1986 system, although a useful prototype, has numerous failings. The most critical among these is the speed of image update, which is caused by many factors.

• Speed

The main problem with the speed of the system is that static objects do not appear to be rigidly located in the virtual space. For this illusion, a shorter update cycle is required. Chung and Harris have measured the lag in this system to be 200 milliseconds.

The delays are in several parts of the system:

- The Polhemus, which is operating near its maximum speed (specified as 60/second, with lag of 50 msec).
- The transmission to and from the Polhemus, which is being done on a slow serial line at 9600 baud.
- The ethernet connecting the Masscomp to the PS-330, which tends to slow down (measured delays on the order of 50 msec) when traffic is high (it is part of a multi-user system).
- The PS-330 itself, which spends a fair amount of its time performing view update calculations (measured at UNC as being 40 milliseconds).
- The Masscomp 5500, which performs the matrix calculations and also has operating system (Unix™) overhead. Chung has measured the delay in this part of the cycle to be 30 msec. We do not yet know how much of that time is due to the operating system and how much is actually calculation.

• Field of View

The field of view in the 1986 system is quite narrow, since the image subtends such a small solid angle of the user's vision ($\pi/8$ steradians). (The normal human field of view is over 2π steradians (Diffrient[9])). This limits the degree to which a user can immerse himself in the virtual world, and thus limits the realism of the system.

• Limited Freedom of Movement

Due to the limited range of the Polhemus (a hemispherical area 5 feet in radius), the user can only move about in a small volume. For visualizations requiring more movement, the system is too limiting.

• Lack of Image Realism

Since the graphics device used (Evans & Sutherland PS-330) is strictly a line-drawing device, the type of objects that can be represented in any realistic fashion is severely reduced. Any task requiring visual realism cannot be done by this system.

• Unwieldiness

The 1986 system has numerous wires connected to it which become obtrusive once the user moves around for any length of time. There are wires from the belt pack to the head unit which can be annoying, as well as wires from the rest of the system to the belt pack which tend to inhibit movement and can trip the wearer.

• Adjustments

The 1986 system has more adjustments than necessary since it was designed to be flexible. These adjustments tend to be difficult to "fine-tune" and are generally difficult to make accurately or easily.

Chapter 3

Planned Improvements

Section 1 - Head Unit

There are several places the HMD can be improved : wide angle optics, better user adjustments, color images and high-resolution images.

• Wide-Angle Optics

The feasibility of using more sophisticated optics to enable wide-angle viewing is being considered. The system would be similar to the 1986 version, but with lenses between the user's eyes and the half-silvered mirror so that the image would subtend a larger solid angle of the user's vision. To retain a coherent world view, corrective lenses would have to be placed beyond the half-silvered mirrors.

• Better User-Specific Adjustments

One improvement would replace the friction adjustments for moving the lens/mirror assembly vertically with thumbscrew adjustments, so that moving this part could be made smooth and easy. Similarly, the interocular adjustment would be a true one, with thumbscrews controlling the movement of both the optics and the screens, so the system could better accommodate physical differences between users. If wide-angle optics were used, then software could be used for line-of-vision adjustments (this is not feasible without wide-angle optics, since the field of view is so small that any movement of an object tends to move it off of the screen).

• Color Images

Another improvement in the 1986 system will be the addition of new (commercially available) Seiko LVD-302 color LCD TV screens for final image display. The addition of color greatly increases the amount of information that can be put on the screen, as well as the realism of the image (for raster images).

• Higher Resolution

In addition to color, the new LCD technology also has higher resolution (the Seiko model has a 2" diagonal screen with 220 x 320 resolution). With this, more complex objects can be represented clearly and thus more realistically.

Section 2 - Host Computer

This part of the 1986 system works adequately; the only improvements would be a faster CPU and perhaps hardware to speed matrix calculations. As stated in Chapter 2, Section 2.2, we do not yet know how much of the delay in this part of the cycle is due to operating system overhead and how much is actually processing time. A faster CPU would help on both of these delays.

Section 3 - Graphics

This is the sub-system that will undergo the most change soonest. Both short and long-range changes in this area will radically improve the system's performance and remove some of the more serious limitations discussed at the end of Section 2.

• Wieldiness

One of the parts of the system that added to its unwieldiness was the heavy coaxial cables from the cameras (or any video source) to the belt pack. These can be removed if that signal is sent into small TV transmitters which broadcast the signal to the (existing) TV receivers in the belt pack. This would greatly increase the freedom of movement and wieldiness of the system. We will still have the tracker cable and an insignificant DC power wire.

• **Realism**

Improvements in raster technology have produced machines (the UNC Pixel-Planes machine, the Pixar machine, etc) that promise complex-scene image update rates fast enough for this system (20-30 per second). This would remove one of the major limitations cited in Section 2: realistic image generation as well as the remove the need for the video cameras in the graphics path.

These raster machines are capable of rapidly generating realistic images of polygon-represented objects with smooth shading, anti-aliasing and shadows. This would radically increase the number of objects that could be represented and viewed by this system.

Since these are raster devices, their output signal can be easily converted to (or is already) a standard video signal which could be patched in or transmitted to the HMD screens, thus obviating the need for the video cameras.

The new higher resolution, color LCD TV's discussed in Section 1 will also improve image quality and, hence, realism.

Section 4 - Tracker

This is a field where the developments are not obviously on the horizon. There are systems on the market (Woltrang[6], United Detector Technology[7]) and on paper (Bishop [4]) which can operate at higher speeds than the Polhemus, but they have numerous other failings such as : lack of support for interactive use (Woltrang [6]), unwieldiness (Bishop [4]), inability to generate orientation information quickly (United Detector Technology[7]).

Some speedups can be made with the existing tracker technology : the Polhemus tracker will soon be used in continuous mode (as discussed in Chapter 2, Section 4.2). A faster baud rate on the communications link between it and the Masscomp will also improve performance. With these improvements, we hope this part of the update cycle will take only 30 milliseconds.

Bibliography

1. Polhemus Navigation Sciences Division, McDonnell Douglass Electronics Company, 3 Space Isotrak User's Manual, August 1985.

2. Evans & Sutherland, PS 300 Programmer Reference, "Graphics Support Routines", September 1984.

3. Evans & Sutherland, PS 300 Graphics Programming, "Graphics Principles", September 1984.

4. Bishop, Gary, "Self Tracker : A Smart Optical Sensor on Silicon", PhD Dissertation, University of North Carolina, May 1984.

5. Anderson, Peter, Engineer at Polhemus Navigation Sciences Division, McDonnell Douglass Electronics Company, personal communication, 1986.

6. Woltrang, H.J., "New possibilities for human motion studies by real-time light spot position measurement," *Biotelemetry*, Vol. 1, 132-146, 1974.

7. United Detector Technology, Trade literature on OP-EYE optical position indicator, United Detector Technology, Inc., Santa Monica, CA, 1981.

8. Sutherland, Ivan, "The ultimate display," Proceedings of the IFIP Congress, Vol. 2, 506-509, 1965.

9. Diffrient, Niels, Tilley, A., Bardagjy, J., Humanscale 1/2/3, MIT Press, Cambridge, MA, 1974.

10. Lipscomb, James, "3-D Cues for a Molecular Computer Graphics System," PhD Dissertation, University of North Carolina, 1981.