


Exploring Virtual Worlds with Head-Mounted Displays

TR89-009

February 1989

*James C. Chung, Mark R. Harris
F. P. Brooks, Henry Fuchs
Michael T. Kelley, John Hughes
Ming Ouh-young, Clement Cheung
Richard L. Holloway, Michael Pique*

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



Appeared in Non-Holographic True 3-Dimensional Display Technologies, SPIE Proceedings, Vol. 1083, Los Angeles, CA, January 15-20, 1989.

UNC is an Equal Opportunity/Affirmative Action Institution.

Exploring virtual worlds with head-mounted displays

J.C.Chung, M.R.Harris, F.P.Brooks, H.Fuchs, M.T.Kelley, J.Hughes, M.Ouh-young, C.Cheung, R.L.Holloway, M.Pique

University of North Carolina at Chapel Hill
Department of Computer Science
Chapel Hill, North Carolina

ABSTRACT

For nearly a decade the University of North Carolina at Chapel Hill has been conducting research in the use of simple head-mounted displays in "real-world" applications. Such units provide the user with non-holographic true three-dimensional information, since the kinetic depth effect, stereoscopy, and other visual cues combine to immerse the user in a "virtual world" which behaves like the real world in some respects.

UNC's head-mounted display was built inexpensively from commercially available off-the-shelf components. Tracking of the the user's head position and orientation is performed by a Polhemus Navigation Sciences' 3SPACE[™] tracker. The host computer uses the tracking information to generate updated images corresponding to the user's new left eye and right eye views. The images are broadcast to two liquid crystal television screens (220x320 pixels) mounted on a horizontal shelf at the user's forehead. The user views these color screens through half-silvered mirrors, enabling the computer-generated image to be superimposed upon the user's real physical environment.

The head-mounted display has been incorporated into existing molecular modeling and architectural applications being developed at UNC. In molecular structure studies, chemists are presented with a room-sized molecule with which they can interact in a manner more intuitive than that provided by conventional two-dimensional displays and dial boxes. Walking around and through the large molecule may provide quicker understanding of its structure, and such problems as drug-enzyme docking may be approached with greater insight.

In architecture, the head-mounted display enables clients to better appreciate three-dimensional designs, which may be misinterpreted in their conventional two-dimensional form by untrained eyes. The addition of a treadmill to the system provides additional kinesthetic input into the understanding of building size and scale.

1. INTRODUCTION

In 1965 Ivan Sutherland¹ first proposed the *Ultimate Display*—a display in which computer-generated images would behave exactly as their real-world analogs do. Computer-generated chairs could be sat upon. Computer-generated apple pies would smell and taste just like Mom's. And computer-generated bullets would be fatal. Fans of the television series "Star Trek—the Next Generation" may recognize that such a display exists on the latest version of the starship Enterprise in the form of the "holodeck." While Sutherland's Ultimate Display may indeed be 400 years away, we in the 20th century can at least begin to investigate more feasible versions of it as our current technology allows.

Even for displays less fantastic than the Ultimate Display, Sutherland recognized the need for as complete sensory input as possible. Most important is kinetic feedback—the response of the computer display to the user's movement. The senses of sight, sound, and feeling lend themselves most easily to this effect, as objects can be moved out of sight, apparent sound sources can shift their relative position when the user's head is turned, and force feedback mechanisms can respond to hand and arm movements. Such display responses are

¹ 3SPACE[™] is a registered trademark of Polhemus Navigation Sciences, Colchester, Vermont.

under complete computer control and may or may not be limited to familiar real-world behaviors. This display with its computer-controlled objects and their computer-generated behaviors comprises what has come to be known as a *virtual world*. It is the basis for a representation-rich approach to problems that previously may have been limited to pencil-and-paper representation.

Here at the UNC-Chapel Hill, the application of the virtual world approach to various problems has become a major research focus, and the use of head-mounted displays (HMDs) is an important component of this research. To be honest, our head-mounted displays are nothing new. The technology we use is all commercially available and has been used in other head-mounted display efforts. What is new, however, is the application of the head-mounted display to the problems of molecular structure, architecture, and in the future, medical imaging. We hope to demonstrate the diversity of the problems in which the head-mounted display can be effectively used.

2. HEAD-MOUNTED DISPLAYS

The use of head-mounted displays in the exploration of computer-generated virtual worlds is a step toward a completely natural interface between man and machine. We observe users' appreciation of complex spatial interrelationships develop more quickly and with less effort with 3-D dynamic displays and 3-D interaction devices. It is much easier to change one's view of a scene by walking around it or stooping to look up at it than to decompose the desired change into a series of axis rotations which are effected by turning knobs (the "Etch-a-Sketch constraint"). And without being distracted by such superfluous tasks, the user is less likely to become confused and lose his orientation. Clearly, the ideal head-mounted display would be a much preferable alternative to conventional displays. Because the head-mounted display is still a relatively new technology, however, what comprises an ideal head-mounted display can not be indisputably defined, nor does one now exist. We will discuss it later on, however, in the context of examining the deficiencies of our current HMD. First, we review the better-known head-mounted displays of the past 20 years.

2.1 Sutherland's HMD

Sutherland himself took the first step towards the Ultimate Display by building a head-mounted display at Harvard University which he took with him to the University of Utah.² This unit used a pair of small CRTs to display stereoscopic images, and also allowed the wearer to see his real surroundings. Special hardware was designed and built to generate the wireframe images presented to the user. Tracking of the user's head position and orientation was accomplished either with direct mechanical linkage between the HMD and an encoding device attached to the ceiling, or with an ultrasonic head position sensor. Sutherland achieved good results with this device. Despite the fact that the transparent wireframe images sometimes led to ambiguous interpretation, most users were able to experience a real three-dimensional effect. Subsequent work by Vickers enhanced the human-computer interface³, and some years later Clark successfully used the unit to design free-form surfaces.⁴

2.2 MIT

In 1983 Mark Callahan of M.I.T.'s Architecture Machine Group produced an updated version of Sutherland's HMD using the then available improved display devices and computing engines.⁵ In its final configuration, Callahan's unit used two Sony black and white, flat CRT television sets (2-inch-diagonal). The television electronics were carried in a belt pack, and the picture tubes were mounted facing downward on the forehead of a bicycle helmet. The wearer of the unit viewed the CRT images through a pair of beamsplitters conveniently mounted on eyeglass frames. Tracking for the MIT unit was provided by a Polhemus Navigation Sciences magnetic tracker. The tracker was used either on the HMD to report the wearer's head position and orientation for computation of a new view, or on the wearer's hand to allow interaction between the user and the scene. Another of Callahan's innovations was the use of optical disk technology to quickly display high-quality pre-recorded images in response to the wearer's head movements.

2.3 NASA

At the NASA Ames Research Center Fisher *et al*⁶ developed the next step up in head-mounted displays, intended for telerobotics and space station information management applications. This unit was capable of displaying computer-generated images or images from remote cameras, and of mixing either kind with frames stored on optical video disk. Breaking tradition with previous systems, the NASA unit positioned its liquid crystal display screens directly in front of the wearer's eyes. While this precluded the superpositioning of computer-generated images onto the user's real surroundings, it permitted the use of the LEEP[™] wide-angle optical system developed by Eric Howlett of Pop-Optix Labs. These optics enabled the 2.7-inch-diagonal screens to cover an extremely large portion of the wearer's field of view—120° horizontally and vertically with a common binocular field of up to 90°. One complication of this system is that some predistortion is required to compensate for the effect of the lenses. Like the MIT unit, the NASA HMD also used a Polhemus tracker. Fisher also enhanced the user interaction with the virtual world through the use of the gesture-sensing DataGlove[™] from VPL Research and through the incorporation of speech-recognition into the system. Both devices provide simple, natural means of communication with the virtual world.

2.4 CAE Electronics

CAE Electronics Ltd. of Quebec has developed a fiber-optic helmet-mounted display system (FOHMD)⁷, intended for use with air combat flight simulators and other such applications as remotely piloted vehicles. In the FOHMD system, four light-valve projectors transmit the two eyes' images through fiber-optic cables to the helmet display, where they are viewed through a wide angle display. The beam splitters on the helmet display permit the pilot still to view cockpit indicators and head-up displays. The FOHMD provides a field of view of 64° vertically by 135° horizontally, including a high resolution inset field (25° by 19°). Head tracking is achieved through two solid state sensors, each capable of reporting the two-dimensional position of an infrared LED within its field of view. By flashing the helmet-mounted LEDs in sequence, the helmet position and orientation can be computed from the information supplied by the sensors.

2.5 VCASS

Under the direction of Dr. Thomas Furness, an experimental head-mounted display was developed at the Armstrong Aerospace Medical Research Laboratory at Wright-Patterson Air Force Base⁸. The Visually Coupled Airborne Systems Simulator (VCASS) was designed as an inexpensive platform with which new cockpit configurations could be evaluated. The VCASS uses miniature television tubes and an innovative optical system to present a 120° three-dimensional scene to the pilot. It also features gesture and voice communication with the host, and three-dimensional sound display.

3. HEAD-MOUNTED DISPLAY RESEARCH AT UNC

3.1 Systems

Head-mounted display research at UNC actually began in the early 1980's, when a small television was mounted on a robot manipulator arm.⁹ A video camera directed at a vector display provided the image for the television. The manipulator reported the television's position and orientation, and the vector image was continuously updated to reflect the current view. This was not a true head-mounted display, but still contained the elements of real-time virtual world exploration on which the HMD is based. Below are described the two head-mounted displays we are currently using.

3.1.1. See-through HMD. Borrowing ideas from the Sutherland and Callahan units, we constructed a see-through head-mounted display cheaply and simply from off-the-shelf commercially available products. (See Figure 1.) The unit was built on the plastic suspension straps from a pilot's instrument-training hood, onto

^{**} LEEP[™] is a registered trademark of Pop-Optix Labs, Newton Center, Massachusetts.

^{***} DataGlove[™] is a registered trademark of VPL Research, Inc., Redwood City, California.

which was mounted a horizontal shelf located at the wearer's eyebrows. Two Seiko color liquid crystal television sets were dismantled to provide the 2-inch-diagonal display screens and driving circuits. The television electronics were carried in a "fanny pack," and the liquid crystal screens were removed and positioned on the horizontal shelf, one above each eye. These screens have a resolution of 220v x 320h pixels. Half-silvered mirrors at a 45° angle enable the wearer to view the screens while still being able to see his physical surroundings. Plastic lenses between the half-silvered mirrors and the screens adjust the focal length to a comfortable value, and an electroluminescent panel backlights the liquid crystal screens. The field-of-view presented by this unit was approximately 25° horizontally.

UNC See-Through Head-Mounted Display




Figure 1. The UNC see-through head-mounted display.

3.1.2. AFIT HMD. We have recently begun collaborating with a group at the Air Force Institute of Technology, located at Wright-Patterson Air Force Base and under the leadership of Major Phil Amburn, a recent UNC student. Amburn's group has designed and constructed a head-mounted display similar to the NASA Ames unit. (See Figure 2.) It is built on a bicycling helmet, and although one is able to strap it on quite securely, it is much heavier than our see-through HMD. The moment of inertia of the unit is very high, and we feel that removing a couple of pounds will greatly reduce user fatigue. Since television technology is constantly improving, Amburn was able to use larger 3-inch-diagonal color television screens, and with simple magnifying optics, these screens provide a horizontal field of view of approximately 55°. The AFIT unit uses fluorescent backlighting, which provides brighter images than our see-through HMD's electroluminescent panel. We have also found that with very minor modification, the LEEP optics can be easily incorporated, and in first experiments, even with no correction for the optical distortion, the wide-angle optics enhanced the visual effect.

UNC/AFIT Head-Mounted Display




Figure 2. The UNC/AFIT opaque head-mounted display.

3.2. Display systems

We use two display systems to generate images for head-mounted displays. The first is an Evans & Sutherland PS300 driven over an Ethernet by a Masscomp 5500. The PS300 was chosen primarily because of its ability to provide real-time response to changes in viewing parameters, and also because of its widespread use in the molecular graphics community. Left and right images are displayed in separate viewports on the PS300, and two television cameras are aimed at the viewports. The camera signals are fed into two television transmitters, with each eye's image being transmitted on a different UHF frequency. From the television electronics in the HMD's fanny pack, a pair of ribbon cables run up the wearer's back and over his head to deliver the image signals to the liquid crystal screens. Though crude, this setup is quite effective in generating color stereo vector images and presenting them to the HMD's wearer.

The second display system involves the transmission of NTSC-encoded images generated by our Pixel-Planes 4 high-speed graphics processor. The Pixel-Planes 4 machine is capable of displaying 30,000 fully-shaded, arbitrarily-sized polygons per second¹⁰, and is thereby capable of displaying moderately complex virtual world scenes with more realism than that provided by the PS300's vector display. In its current configuration, however, the Pixel-planes raster display is incapable of handling stereo image pairs.

3.3. Tracking devices

3.3.1. Polhemus Navigation Sciences' 3SPACE Tracker. For all our head-mounted displays, tracking of the user's head position and orientation has been performed by a Polhemus Navigations Sciences 3SPACE Tracker. The Polhemus provides usable information fairly quickly, but suffers from sensitivity to magnetic perturbations in the environment. In our graphics lab, there are many sources of such magnetic perturbations. While most of these can be placed out of harm's reach so that they do not affect the Polhemus' data, there is no escaping the effect of our metal floor. The metal floor makes itself known by curving the Polhemus' space, so that orientation values reported close to the limits of the Polhemus' spatial range are off by as much as 30°. Close to the Polhemus' radiator, however, the effect is negligible. Since the floor effect is static, it is possible to correct for it in software.

3.3.2. Optical tracking. Bishop¹¹ has proposed a design for an outward-looking optical tracking device, which could sit on the head-mounted display and require no other external equipment. Although Bishop has some doubt that his Self-Tracker could achieve the translational accuracy required by an HMD, the inside-out tracking concept is noteworthy. It is inherently an accurate way to track orientation. Jih-Fang Wang, of our department, is currently developing a reliable inside-out optical tracker that uses fixed fiducial light sources in the environment.

3.4. Input devices

In exploring virtual worlds one needs the ability to interact with and manipulate the scene. We have been using a Polhemus-based, two-button, three-dimensional mouse, constructed from a hollowed-out billiard ball. While it is adequate for selection and manipulation tasks, its weight, shape, and size tend to accelerate user fatigue. We have just recently received a VPL DataGlove, and plan to incorporate it into our system shortly.

4. APPLICATIONS

Two major fields of research at UNC, interactive molecular studies and virtual building exploration, are exploring the advantages of the head-mounted display. A third application area in which we are active, medical imaging, also shows promise of making good use of head-mounted display technology.

4.1. Interactive Molecular Studies

Since 1970, the GRIP project has had as its objective the exploration of molecular graphics tools which may be financially out of reach for today's chemistry laboratory, but not for tomorrow's, as hardware improves and becomes less expensive. As such, we are very interested in the benefits arising from incorporating the head-mounted display into existing molecular graphics tools. Macro-molecules possess complex three-dimensional structures. Understanding these structures is often the key to understanding chemical properties. Various methods of visualizing molecular structure have evolved over time, from physical ball-and-stick models to computer-generated images viewed on two-dimensional screens. Computer-generated images usually rely on kinetic depth effect, intensity depth cueing, and stereo to help the user perceive three-dimensional structure.

We envision a system in which the chemist wearing a head-mounted display sees a room-sized virtual molecule. The chemist walks around the molecule to study its structure, and may even walk through the molecule to examine the internal structure. With appropriate interaction devices, the chemist can more fully explore the molecule and its properties. For example, to study the effects of a different torsion angle on the molecule's structure, the chemist need only reach out and twist the bond in question. If the bond is out of reach, he can pull the molecule down or scale it to a more manageable size. With some animation it would be possible to watch this giant molecule wriggle its way through its calculated vibrational modes in response to changes in temperature. This type of head-mounted display application is still only a vision for us, but future efforts will be devoted to integrating our HMD with the our Grinch electron density map fitting program and Tripos Associates' Sybyl/Mendyl molecular modelling package.

As a first step toward investigating molecular modeling applications of the head-mounted display, we have implemented a demonstration of enzyme-substrate docking using the HMD and the three-dimensional mouse. The docking problem is important both to biochemists attempting to understand how proteins and nucleic acids work, and to chemists doing analytic drug design, attempting to find or synthesize drugs with specified active properties and without undesirable side effects. Our system presents the user with sphere-model renderings of a dihydrofolatereductase (DHFR) molecule roughly two feet in diameter and of a similarly scaled methotrexate molecule, both suspended in mid-air about five feet off the floor. Using the three-dimensional mouse, the user positions the cursor in the methotrexate molecule, and by then depressing the mouse button he is able to manipulate the methotrexate. As long as the mouse button is depressed the methotrexate will follow the mouse's position and orientation. With this control, the task is to insert the methotrexate into the active site of the DHFR molecule. If the active site is not in a convenient position, the

user can release the drug and similarly manipulate the global scene (DHFR and methotrexate together) by depressing the mouse button outside of the drug's "grabbing" volume. In this mode, both molecules move in concert. When the active site is positioned to the user's liking, he can then grab the methotrexate again and continue with the docking.

This system is intended to be an enlightening demonstration of the types of interaction possible with the head-mounted display. In its current state, there are no cues to tell the chemist when correct docking has occurred. Indeed, it is possible for both virtual molecules to occupy the same volume in space, a situation not allowed in the real world. Current research by Ouh-young is addressing the problem of providing the chemist with docking status information through force feedback and visual cues.¹²

4.2. Virtual Building Exploration

The Walkthrough project is aimed at the development of a visualization tool with which a designer and his client can explore an as yet unbuilt building.¹³ Whereas the architect has been trained to visualize unrealized designs, the client may often be unable to translate drawings and descriptions into a complete mental image. With this handicap the client may wind up accepting unsatisfactory designs because of some misunderstanding. Walkthrough has been developed to avoid such pitfalls. With this tool the designer is able to rapidly prototype a solution, and receive from the client meaningful feedback that enables him to iteratively refine his design to the client's satisfaction. There should be no surprises when the building is finished.

The interesting aspect of Walkthrough is that, unlike the virtual molecule environments, it provides a virtual environment with which the users are familiar—they already know how to interact with a building, virtual or real. Everybody knows how to physically walk through and visually explore a building to understand its spatial interrelationships, and everyone is aware that climbing a flight of stairs to reach a balcony will provide a different vantage point from which to see the building. This familiarity is exploited by Walkthrough, which provides modes of interaction with the virtual building that mimic interaction with a real building.

The head-mounted display provides one of these more natural interaction modes. With it, the user is no longer confined to looking at everything through a fixed screen, but is free to crane his neck in any way desired to increase his understanding of the building. Because the Polhemus sensor has a limited effective radius of approximately five feet, however, it is necessary to provide the user with some means of moving about the building while still remaining within the Polhemus' constraints. A vehicle simulator was implemented to allow an HMD-wearing explorer to drive through the virtual hallways and rooms, but we postulated that actively walking through the building would give a greater appreciation of its space. Toward that end, a steerable treadmill has been incorporated into the system, and the user is now able to stroll through the virtual building, all the while receiving subconscious spatial dimension cues from his proprioceptive senses. Interaction with the virtual building is also possible, for the 3-D mouse can easily be used to push virtual furniture around.

5. ISSUES

What issues must be addressed in the design and construction of a head-mounted display? An ideal display is unattainable with today's technology. We do what we can. Below we discuss some of the problems and compromises we encountered in building our head-mounted displays.

5.1. Lag time/Update rate

These two system characteristics should not be confused with each other. Update rate is a measure of how quickly the system can generate and display a new image—the higher the update rate, the smoother the apparent motion perceived by the user. Lag time is a measure of the end-to-end delay incurred by the system in

computing and displaying a new image in response to some event. For example, when a user turns his head, the time it takes for a virtual object to move out of view in response to that movement is the system lag time. Lag time and update rate are different aspects of the illusion we are trying to create, and may require different solutions. Our system runs at about 15 or 16 updates/second, which is sufficient to generate reasonably smooth motion. When the lag is too large, however, virtual world behavior trails noticeably and annoyingly behind the user motion it is responding to. Our early systems had a 400 millisecond lag, and the "swimming" behavior of the virtual objects due to head motion caused by that much delay completely destroyed any chance of creating a successful illusion. Since moving to the Pixel-planes system, we have made preliminary measurements which suggest our lag time is down to around 120 milliseconds.

5.2. Image quality

Image quality is perhaps the aspect of head-mounted displays most limited by current technology. One of the guiding principles for much of the graphics research going on at UNC is not "How can we make realistic pictures?" (after all, who knows what a realistic picture of a molecule would be?), but rather "How can we make useful pictures?" If the pictures are useful, then making them prettier is icing on the cake. In our HMD program, we are often restricted to what is available, what is practical, and what we can afford. Even so, annual upgrades in our display screens, made possible by advances in pocket television technology, have enormously enhanced our system.

5.2.1. Field of view. The limited field of view of our initial systems ($\sim 25^\circ$) was considered by many to be the second most bothersome characteristic (after the long lag time). Peripheral visual cues are very important to creating a sense of total immersion in a virtual world, and these were just not possible with our system. Such wide angle optical systems as the LEEP system provide extremely wide fields of view. The tradeoff is that images must be predistorted to compensate for the radial distortion created by the lenses. This is not such a big problem when dealing with vector graphics systems, but it may cause serious difficulties when creating images with geometric primitives on a raster system. The distortion also makes more difficult the superposition of the virtual image on the real world, for the view of the real world must also be predistorted to achieve correct results. Our next generation unit, currently being built for us at AFIT, will incorporate the LEEP optical system.

5.2.2. Resolution. The limited resolution of our LCD screens makes for somewhat poor image quality, but technology advances will come to our rescue. Use of the small CRT displays available today would have decreased the graininess of our images, but we felt the lower weight and power requirements of the LCDs made them preferable to the CRTs. Unfortunately, low resolution screens degrade images through aliasing. Polygon edges have jaggies, and vectors appear to be more like a string of pearls than a continuous line. Within certain limits aliasing can be lived with. Users may be able to interact effectively with a jagged virtual object, but we always welcome improvements in this area.

5.3. See-through/opaque

We have seen two basic groups of head-mounted displays. One group superimposes virtual objects on the user's real environment, and the other completely replaces the real environment with a computer-generated one. Now that we have one of each type we will be able to compare the two.

The see-through approach has the advantage of placing the user in an environment with which he is familiar. Certain landmarks in the room seem to help him navigate through the virtual world. For example, when dealing with a large, complex molecule it may be useful to note that the active site is over by the door, so that it can be easily located. User comfort is another important advantage of a see-through head-mounted display. We have found that when using an HMD which completely removes the real environment it is very easy to become dizzy and disoriented. Time lags in the system prevent the virtual world from providing the correct visual cues which are so important to maintaining our stability. Some relief is provided by allowing the real world to leak in around the edges of the viewing system, but even with this we have had at least one user become so disoriented as to require 15 minutes of supine recovery.

The opaque approach is more amenable to the use of wide-angle optics. As mentioned above, the distortion introduced by such systems as the LEEP viewer can be compensated for in computer-generated images., but such predistortion for real world objects which are viewed through these systems is more difficult. If the real world is removed and computer-modeled, however, this problem is avoided. Another problem which is avoided by an opaque HMD is inconsistency between visual depth cues, encountered by Schmandt in his work on a three-dimensional work station.¹⁴ His workstation allowed users to view stereo images displayed on a video monitor through the reflective part of a half-silvered mirror. On the other side of the mirror, the transmissive properties of the half-silvered mirror enabled the user to see his hand, which controlled a cursor in three-space. With correct registration between the displayed images and the real objects the user could paint in three dimensions and perform other three-dimensional manipulative tasks. Problems arose because opaque virtual objects failed to visually occlude the user's hand. Interposition cues conflicted with stereo cues, confusing the user. We also have this problem when using our see-through system with shaded polygonal virtual objects displayed on our Pixel-planes system. Again, this problem can be avoided with an opaque HMD, for which the computer can hide the cursor when appropriate to maintain consistency between visual depth cues.

The choice between a see-through or an opaque head-mounted display may ultimately be determined by the anticipated application. For us, it really makes no sense to have a see-through HMD for our Walkthrough system. It would be too confusing to have a virtual building superimposed on the real one. On the other hand, the see-through HMD may be more appropriate for the molecular docking problem, for which it is helpful to place the virtual objects within the context of the user's real environment.

6. FUTURE DIRECTIONS

6.1. Display of choice

We have two goals for our work in head-mounted displays. For the short term we would like to make the HMD the "display of choice" in our graphics lab. This means that when somebody has some three-dimensional data which he would like to examine quickly, he would be able to load it into the head-mounted display system with minimal effort and then explore it with the HMD. In its current state, with its many wires and gadgets and its difficult adjustments, all excursions with the HMD must be supervised by experienced personnel.

6.2. Useful applications

For the longer term, we would like to research the proposition that the head-mounted display can be a useful means of visualizing virtual worlds. This may seem obvious, but it has been our experience that user preferences cannot always be predicted correctly. If a molecular modeling package comes with the option of using a head-mounted display, will the average chemist, who will not necessarily be as thrilled with new whiz-bang gadgets as we are, really choose to use it? In its current state, our HMD is not ready for such an evaluation. Much work is still required to bring the head-mounted display up to a level where it has a fighting chance of acceptance.

In addition to its current usage in molecular studies and architecture, our HMD is slated for future application to medical imaging studies. Our department has already developed computer graphics systems to aid in the positioning of radiation treatment beams¹⁵, and it is possible that a head-mounted display would make such a task easier.

6.3. Lag time

The lag time problem was mentioned above and still requires much work. Our current system configuration should permit us to get down to a delay of 100 milliseconds, which is on the borderline of human perception of "instantaneous." As better computing and tracking systems become available, this problem may not be as troublesome as it has been. Another approach is being taken by our collaborators at AFIT, who are attempting to use predictive tracking techniques to reduce the lag effect.

6.4. Stereo

The Pixel-Planes 4 graphics processor does not yet have the capability of generating two separate stereo images, although development of the hardware has begun. We have been getting by with mono images (no stereo disparity or convergence cues), and this has proven to be satisfactory for our Walkthrough application, where interposition, linear perspective, and head motion parallax provide strong depth cueing. The effect is less satisfactory in our molecule docking application, where disparity and convergence would provide valuable cues.

6.5. Interaction

Following the NASA Ames group, we plan to replace our pool-ball mouse with a VPL DataGlove. As many of our anticipated applications are menu-driven to some extent, we are developing pop-up menus using Bezier-defined fonts which can exist in eye space or in model space. Menu selection could be done manually with the DataGlove or through voice input. Audio feedback will also be added, and should prove extremely useful in bump checking of virtual objects.

6.6. Calibration

The head-mounted display sits differently on different heads. This means that there are interuser variations in gaze direction relative to the Polhemus sensor on the HMD, interocular separation, field of view perceived by the user, and registration between real world objects and their computer-generated counterparts (e.g. hand/cursor). These variations can be adjusted for in software for each user, but simple, effective calibration schemes must be developed.

6.7. Processing power

We eagerly await the completion of the next generation Pixel-Planes 5 graphics engine. The anticipated twenty-fold speed increase and the ability to work with multiple frame buffers will allow us to explore much more complex virtual worlds in real-time.

7. CONCLUSION

Research in our department is aimed at the development of real, working systems. As an intuitive and natural means of exploring virtual worlds, the head-mounted display holds great promise for improving human-computer interaction. In this paper, we have outlined our approach to some of the design issues involved with head-mounted displays, and have discussed the applications in which we are using the HMD. Much work lies ahead, however, before the head-mounted display can become a commonplace tool in the repertoire of problem-solving aids provided by computer graphics.

8. ACKNOWLEDGMENTS

Our appreciation goes to David Lines for his editorial assistance in the preparation of this paper.

This research was supported by the National Institutes of Health (Grant RR 02170-05), and the Office for Naval Research (Contract No. N00014-86-0680).

9. REFERENCES

1. I.E. Sutherland, "The ultimate display," *Proceedings of the IFIP Congress 2*, 506-508 (1965).
2. I.E. Sutherland, "A head-mounted three-dimensional display," *1968 Fall Joint Computer Conference, AFIPS Conference Proceedings*, 33, 757-764 (1968).

3. D.L. Vickers, *Sorcerer's apprentice: head-mounted display and wand*, Ph.D. dissertation, Dept. of Comp. Sci., Univ. Utah, Salt Lake City (1974).
4. J.H. Clark, "Designing surfaces in 3-D," *Communications of the ACM*, 19(8), 454-460 (1976).
5. M.A. Callahan, *A 3-D display head-set for personalized computing*, M.S. thesis, Dept. of Architecture, Massachusetts Institute of Technology, 110 pp. (1983).
6. S.S. Fisher, M.McGreevy, J. Humphries, and W. Robinett, "Virtual environment display system," *Proc. 1986 Workshop on Interactive 3D Graphics*, 77-87 (1986).
7. CAE Electronics, Ltd., *Introducing the visual display system you wear*, 4 pp. (1986).
8. C.V. Glines, "Brain buckets," *Air Force Magazine*, 69(8), 86-90 (1986).
9. H. Fuchs, *private communication*, (1988).
10. J. Eyles, J.D. Austin, H. Fuchs, T.H. Greer, J. Poulton, "Pixel-Planes 4: A summary," *Advances in Computer Graphics Hardware II*, A.A.M. Kuijk and W. Strasser, ed., Eurographics Seminars, Springer-Verlag, 183-208, (1988).
11. G. Bishop, H. Fuchs, "The SELF-TRACKER: A smart optical sensor on silicon," *Proc. 1984 MIT Conf. on Advanced Research in VLSI*, Artech House, Dedham, Mass., 65-73 (1984).
12. M. Ouh-young, M. Pique, J. Hughes, N. Srinivasan, F.P. Brooks, "Using a manipulator for force display in molecular docking," *Proc. 1988 Int'l Conf. on Robotics and Automation*, 1824-1829 (1988).
13. F.P. Brooks, "Walkthrough—A dynamic graphics system for simulating virtual buildings," *Proc. 1986 Workshop on Interactive 3D Graphics*, 9-21 (1986).
14. C. Schmandt, "Spatial input/display correspondence in a stereoscopic computer graphic work station," *Computer Graphics*, 17(3), 253-261 (1983).
15. C.E. Mosher, G.W. Sherouse, P.H. Mills, K.L. Novins, S.M. Pizer, J.G. Rosenman, E.L. Chaney, "The virtual simulator," *Proc. 1986 Workshop on Interactive 3D Graphics*, 37-42 (1986).