AD-A245 819

Advanced Technology for Portable Personal Visualization

Report of Research Progress
September 1990 – March 1991

DTIC
ELECTE
FEB 10 1992
S D

This research is supported in part by
DARPA ISTO Contract No. DAEA 18-90-C-0044

Department of Computer Science
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3175

DEFENSE TECHNICAL INFORMATION CENTER
9203247

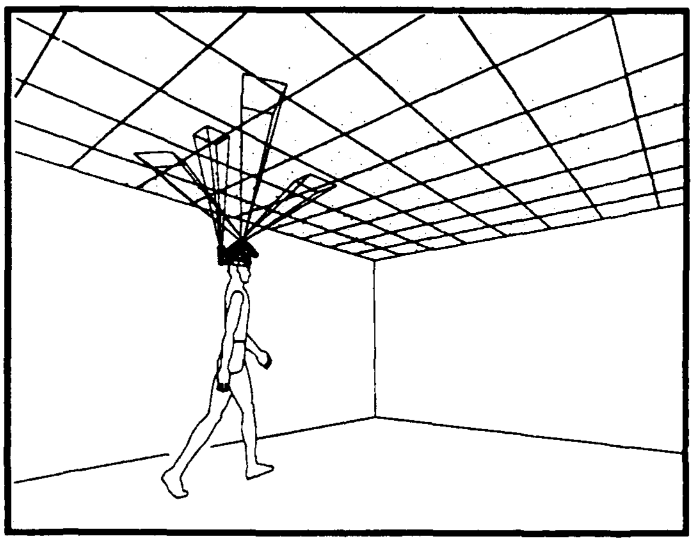

Principal Investigators: Frederick P. Brooks, Jr., and Henry Fuchs
Project Manager: Warren Robinett

Faculty :Steve Pizer, Professor; Vern Chi, Director, Microelectronic Systems Lab; John Eyles, Research Assistant Professor; Gary Bishop, Adjunct Assistant Professor

Staff: Jannick Rolland (Optical Engineer), Mark Ward (Mechanical Engineer)
David Harrison (Hardware Specialist, Video), John Hughes (Hardware Specialist, Force Feedback ARM)
Linda Houseman (Administration), Fay Ward (Secretary)

Research Assistants: Ron Azuma (Optical Tracker), Bill Brown (Sound, Speech Recognition), Jim Chung (Team Leader, Graphics Software), Carney Clegg (Tracker Math Software), Drew Davidson (Sound User Interface), Eric Erikson (Molecular Graphic), Stefan Gottschalk [undergraduate] (Tracker Simulation Software), Rich Holloway (Graphics Library), Phil Jacobsen (Optical Self-Tracker), Russ Taylor (Force Feedback ARM)

92

Advanced Technology for Portable Personal Visualization

1.0 Summary

The user of a Head-Mounted Display is immersed in a computer-simulated or remotely-sensed three-dimensional world. Despite an avalanche of publicity for Virtual Reality in the popular press, many difficult technical problems remain unsolved for this technology, and no application of it has yet been demonstrated to be commercially viable.

We believe that the hardware and software required to create high bandwidth, multi-sensory virtual worlds will in a decade be as cheap and ubiquitous, and that virtual worlds systems will be used in a variety of applications. We are therefore working on three fronts: improving the crucial hardware components required, developing a software base, and demonstrating the usefulness of this technology to solve selected real-world problems.

1.1 Goals of the Head-Mounted Display Project

- Demonstrate the usefulness of the head-mounted display in real applications

- Improve the hardware subsystems which currently limit performance of HMD (tracker, HMD optics and display, real-time graphics generation)

- Design and implement software base to support HMD functions

- Integrate visual, auditory and haptic (force feedback) displays into a working system

- Build new input devices and investigate methods of manual control suitable to a head-mounted display

1.2 Goals of the Optical Tracker Sub-project

Design and build a new real-time position and orientation tracking device for head-mounted display systems that features a large working area (20' x 20'), high resolution (±2 mm translation, .1° rotation), rapid update rate (> 30 Hz), and immunity from environmental disturbances.

Two optical tracking technologies are under development. One, referred to as the optoelectronic tracker, relies on imaging scenes in a structured environment, the other, known as Self-Tracker, does not. The optoelectronic tracker measures the absolute position of the head relative to IR light emitting diode (LED) beacons in a specially prepared ceiling. A cluster of small image sensors are located atop the helmet of a head-mounted display. The sensors form images of LEDs mounted overhead, in a suspended ceiling. The LEDs serve as navigation beacons in that the knowledge of their position in the room, coupled with their image formed atop the head, allows both the position and the orientation of the head-mounted display to be computed. Self-Tracker relies on custom image sensors to measure relative head motion by measuring displacements in images of an unstructured scene. Although a minimum of 6 one-dimensional Self-Tracker sensors are required to compute the six degrees of freedom of the head, we expected to use 10 to 20 for redundancy and noise immunity.

Major challenges facing the project team are: the accurate placement of LED beacons on a suspended ceiling,the calibration of image sensors based on non-metric optics and lateral-effect photodiodes, the real-time computation of head position and orientation, and the management of drift in a Self-Tracker system.

2.0 Summary of Major Accomplishments

2.1 HMD System

• Four models of HMD now operational: EyePhone Model 1, EyePhone Model 2, Air Force Institute of Technology HMD and old UNC see-through HMD

• Three models of tracker now operational: Polhemus 3Space, Polhemus Isotrak and Ascension Bird

• Three HMD stations (HMD, tracker and input device) now operational

• HMD software libraries released and documented: Vlib, Trackerlib and Quatlib

• Most applications now use standard software base (Vlib)

• Sound subsystem functioning and being used in applications

• Integration of force-feedback ARM with HMD demonstrated

• Speech recognition subsystem acquired, programmed and integrated with HMD system

• Commercial 3D modeler purchased and in use to create 3D models for HMD (Sculpt3D on the Macintosh)

2.2 See-through HMD with Custom Optics

• Optical analysis software (Code V) up and in use

• Analyzed optics of the leading commercial HMD (VPL EyePhone)

• Developed model of optical distortion in a HMD

• Design of optics for see-through HMD in progress

• Commercial head-mounting gear acquired (ophthalmologist's headpiece, headgear from night-vision goggles)

2.3 Tracking

• A technique for computing head position and orientation has been developed that appears to offer advantages over Jih-Fang Wang's extension to Church's algorithm. Whereas Wang's solution admitted information from three sensors and at most three LEDs, the new approach allows an unlimited number of sensors to be used and each sensor can image an unlimited number of LEDs. The principle advantage is that measurement error in head position appears to be less susceptible to LED misplacement.

per ltr
A-1

- Hardware has been developed and is currently under consideration to acquire analog data from the head-mounted image sensors and relay the information in digital form to a workstation. This hardware will be worn on the body as a belt pack connected to the workstation host via a 100Mbit/s serial communication link

- Related hardware has been developed that allows a workstation to accept information from the head-mounted sensors and to control the firing of LEDs. The result is a 4 slot wide VME module containing a 68030-based single board computer, a 100Mbit/s serial interface, a parallel interface to the LED controllers, and a power supply.

- LED control circuitry has been developed. The circuit will be used to control up to 32 LEDs mounted in 2'x2' ceiling panels.

- Principal plane locations of a Fujinon lens have been measured using the Siemens SFH487 LED as a light source. Deviation from the Fujinon specification was less than ±50 micrometers.

- Preliminary Self-Tracker sensors have been fabricated by Mosis.

- An interface control board for the cluster of Self-Tracker chips is under design.

- A mechanical engineer, Jack Kite, has been hired for 6 months to design ceiling panels, a support structure, and calibration equipment.

- A new graduate student, Carney Clegg, has joined the effort along with an undergraduate student, Stefan Gottschalk.

- Student competition was sponspored to build mechanical tracker at NC State Mechanical Engineering Department

- Feasibility of inertial tracking investigated (library research, discussions with manufacturers of inertial guidance systems)

3.0 Expected Milestones for the Next 12 Months

3.1 HMD System

- Migrate HMD system onto Pixel-Planes 5 graphics engine

- Implement multi-user shared virtual world on HMD system

- Release Vixen (application for viewing arbitrary 3D databases with the HMD) and get it to function without the need of the workstation screen, keyboard or mouse

- Demonstrate an HMD application that uses visuals, sound and force-feedback

3.2 Optics

- build bench prototype of optics for see-through HMD

- build head-mounted prototype of see-through HMD

3.3 Tracking

- Using the new hardware, calibrate 4 image sensors to 1 part in 1000 accuracy.
- Design and fabricate 100 ceiling panels containing up to 32 LEDs each.
- Design and fabricate a panel support structure that can be adjusted to maintain a tolerance on LED location of ±2 mm over a 20'x20' ceiling.
- Begin work on self-calibration techniques with the goal of developing an algorithm that can compute the position of LEDs and allow less stringent tolerances on absolute LED position.
- Optimize the computation of head position to take advantage of an i860 processor.
- Integrate a four sensor headset with a VPL EyePhone.
- Improve the system's simulation environment.
- Demonstrate a working optoelectronic system.
- Fabricate Self-Tracker interface printed circuit boards.
- Demonstrate Self-Tracker technology using the board-level interface.

4.0 Discussion of Research

4.1 HMD System

The HMD system here at UNC is up and running, supporting a variety of users and applications. We have added several new devices to the system in the last six months (EyePhone Model 2, Bird tracker, another Polhemus tracker, and a Dragon Systems speech recognizer). The following table and diagram show the current configuration of the HMD system hardware.

HMD Hardware Subsystems

	currently using:	*improved version (expected soon)*
• headset -- stereoscopic displays, optics	VPL EyePhone	see-through HMD with custom optics
• high-performance graphics engine	Pixel-Planes 4	Pixel-Planes 5
• tracker for head, hand	Polhemus	optical tracker
• manual input device	billiard ball with push buttons	ring with push buttons
• auditory display	Macintosh and speaker	
• haptic display	Argonne Remote Manipulator (ARM)	

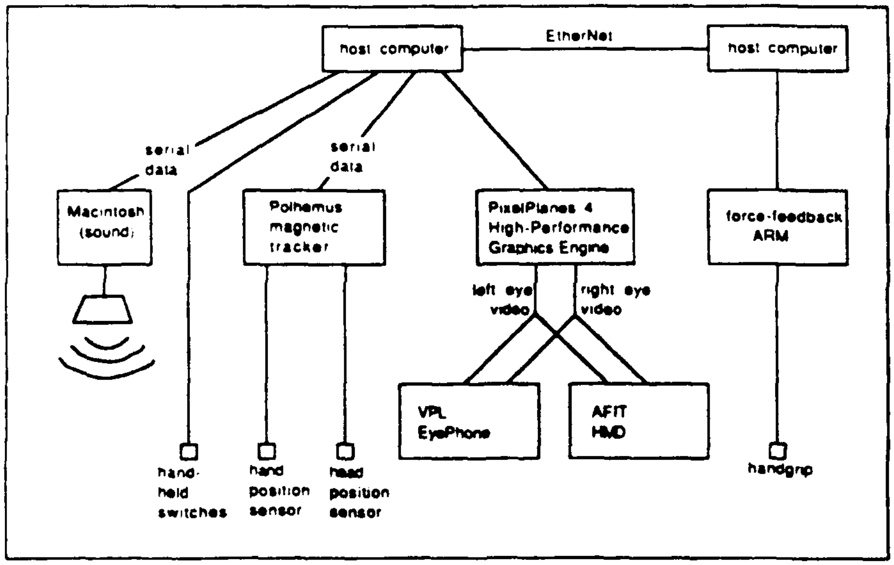

We released an initial version of our standard HMD software library Vlib (virtual worlds library) and some supporting software libraries (Trackerlib and Quatlib). Most applications now use this standard software base. The following diagram and list show the current state and future goals of the HMD software base.

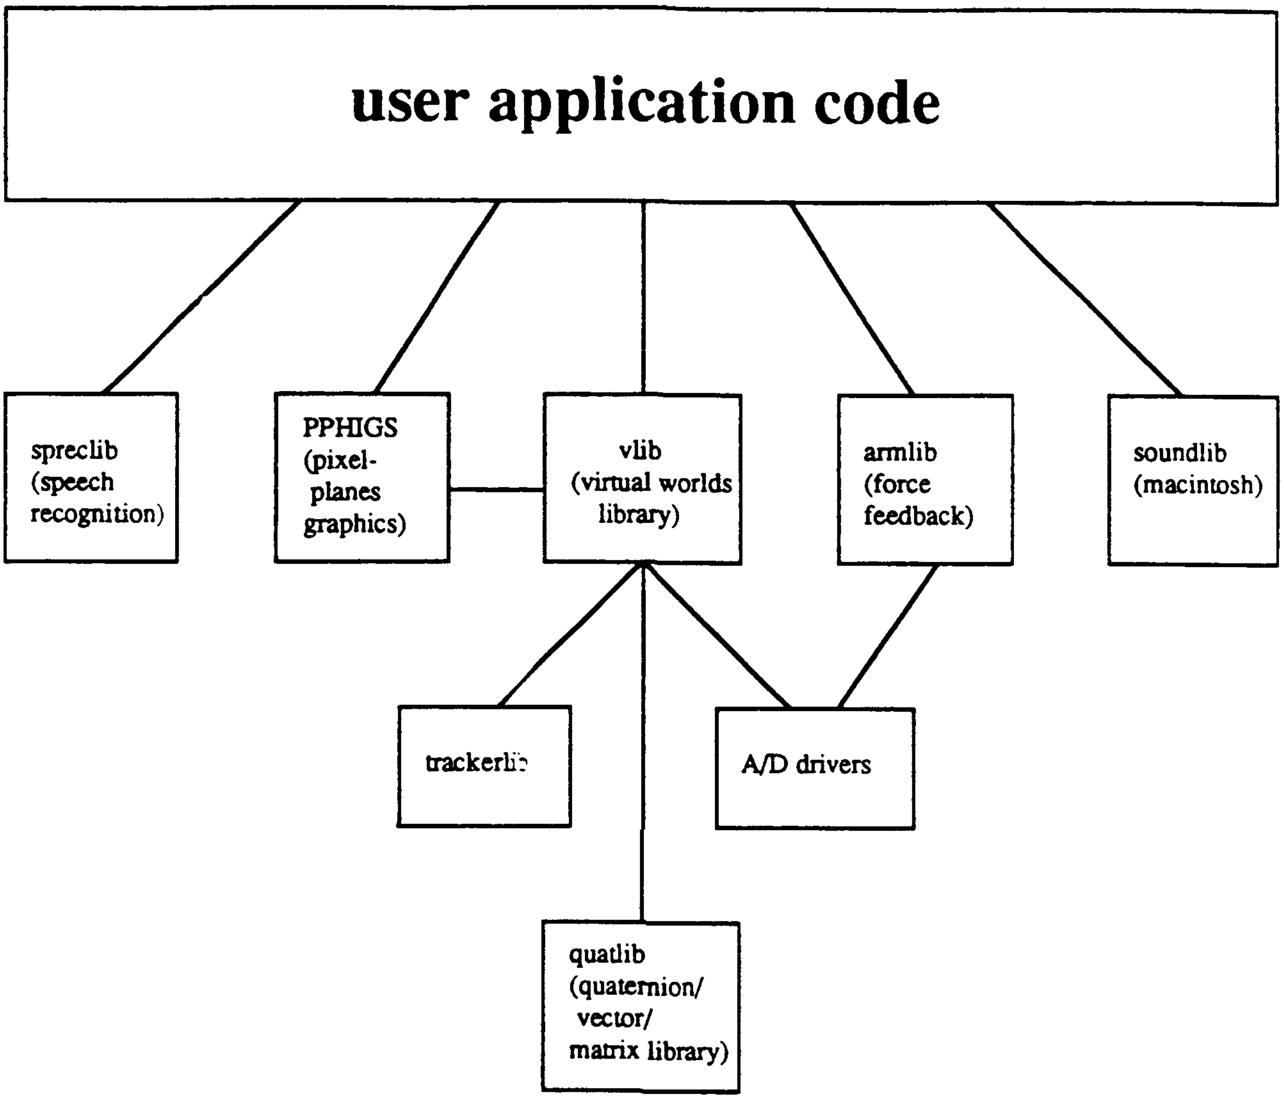

HMD Software Library Goals

- Basic capabilities:
  - -- viewing and walking around in virtual world
  - -- grabbing and moving simulated objects
  - -- flying through the virtual world
  - -- scaling the virtual world
  - -- re-orienting the virtual world
- Standard file format for graphical objects, with hooks for program to refer to objects and their parts.
- Run-time substitution of various trackers, HMDs, and input devices.
- Template program from which to grow new applications.
- 3D user interface.

4.2 Optics

Since last summer when an optical engineer joined our team, we have made substantial progress towards the goal of building a see-through HMD for medical imaging applications which incorporate optics specifically designed for the application. We considered many designs for the HMD optics on a qualitative level. We consulted with optics specialists who have worked on earlier HMDs at Wright-Patterson Air Force Base and the Naval Ocean Systems Center. We acquired the Code V optical analysis software and have it up and running. We developed a computational model for the optics of an HMD, and used this model together with Code V to analyze the optics of the EyePhone, the leading commercial HMD. Using our optics model, we calculated the computational pre-distortion needed to compensate for the optical distortion of the optics used in the EyePhone. Finally, the design is underway for the optics of the see-through HMD we plan to build.

Shown below are diagrams for the monoscopic and stereoscopic optics models we developed.

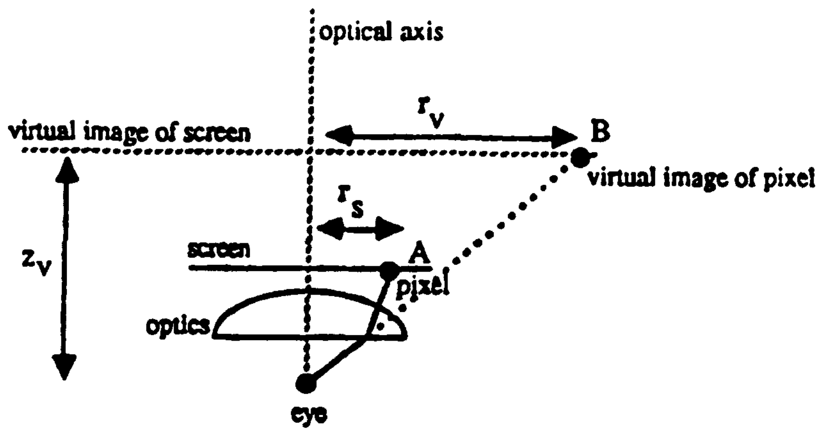

Figure 3.1 Single-eye optics model

object height (of point on screen)	r_s	$r_s = (x_s + y_s)^{1/2}$
normalized object height	r_{sn}	$r_{sn} = r_s / w_s$
normalized image height (3rd order approximation of D)	r_{vn}	$r_{vn} = r_{sn} + k_{vs} r_{sn}^3$
image height (of point in virtual image)	r_v	$r_v = r_{vn} w_v$
angular position of point in virtual image	ϕ	$\phi = \tan^{-1}(r_v / z_v)$

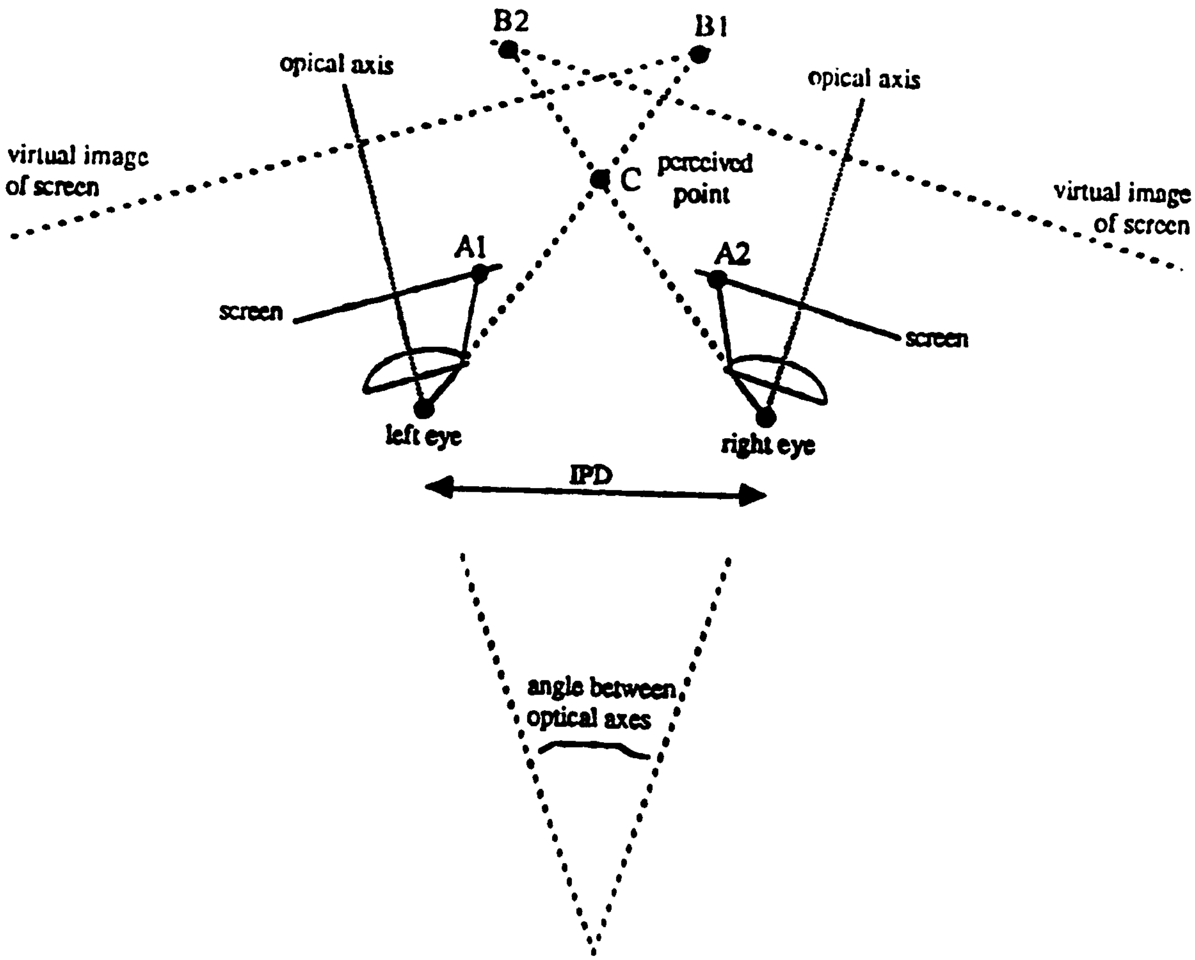
B2
B1
opical axis
opical axis
virtual image of screen
virtual image of screen
C
perceived point
A1
A2
screen
screen
left eye
right eye
IPD
angle between optical axes

The following table shows the results of our analysis of the EyePhone, and the diagram following the table shows the field of view calculated for the EyePhone.

Parameter	Symbol	Value for EyePhone
maximum object field radius	w_s	28.1 mm
maximum virtual image field radius	w_v	271.55mm
transversal magnification	m_{sv}	9.66
eye relief	d_{er}	29.4 mm (nominal)
object plane distance (LCD screen to LEEP lens surface)	d_{ob}	16.4 mm
distance from eye to virtual image plane	z_v	398.2 mm
coefficient of optical distortion for D	k_{vs}	0.32
coefficient of optical distortion for D^{-1}	k_{sv}	-0.18
angle between optical axes	$ø_{axes}$	0 degrees
distance between optical axes (at front surface of optics)	d_{axes}	64 mm
screen center offset from optical axis	(Cx_s, Cy_s)	(6.4 mm, 1.6 mm)
screen resolution	$Res_H \times Res_V$	185 x 139 pixels (color triads)
inter-pupillary distance of user	IPD	varies across users
angular position of virtual image of right edge of LCD	$ø_R$	45.0 degrees
angular position of virtual image of left edge of LCD	$ø_L$	30.3 degrees
angular position of virtual image of top edge of LCD	$ø_T$	31.8 degrees
angular position of virtual image of bottom edge of LCD	$ø_B$	26.6 degrees
single eye vertical field-of-view	FOV_v	58.4 degrees
single eye horizontal field-of-view	FOV_h	75.3 degrees
overlapped field-of-view	FOV_{ov}	60.6 degrees
binocular field-of-view	FOV_{bin}	90.0 degrees

Parameters for the VPL EyePhone, models 1 and 2

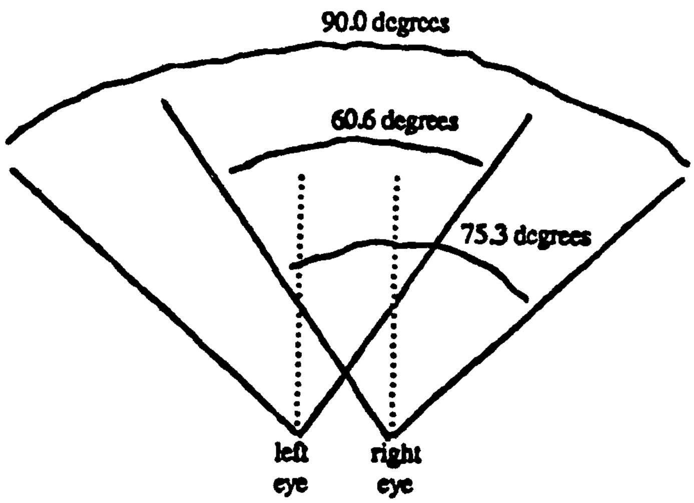

Binocular, overlapped and single-eye FOVs for the EyePhone

The diagram below shows a ray-tracing through a cross-section of the lenses used in the EyePhone. This analysis was done using Code V, and we used this detailed optical analysis to validate the accuracy of our simpler optics model. For more details, see the paper in the Appendix on "A Computational Model for the Stereoscopic Optics of a Head-Mounted Display."

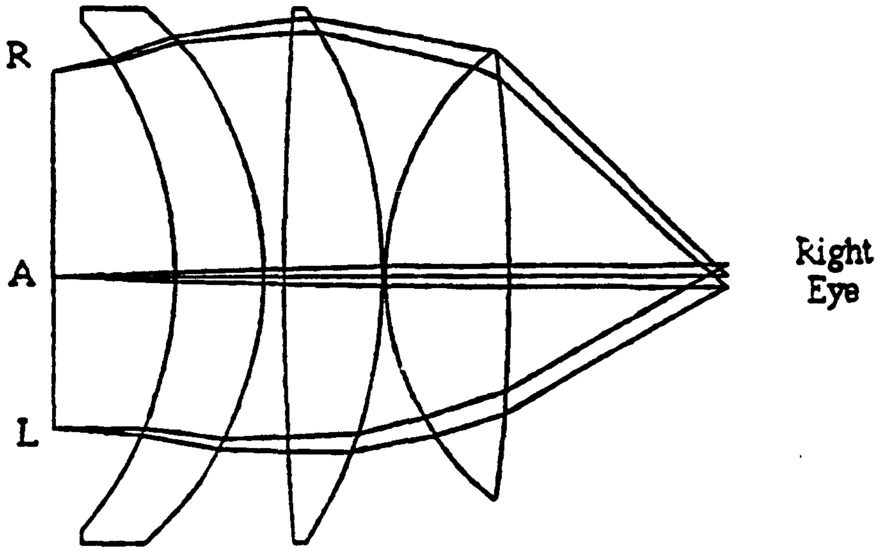

Tracing rays from the EyePhone's LCD screen through the LEEP optics

We used the standard optics representation for lateral field distortion, a third-degree polynomial, to model the distortion of the optics of the EyePhone. This polynomial relates the position (r_{sn}) of a point on the object plane of the optical system with the position (r_{vn}) of the virtual image of that point on the image plane.

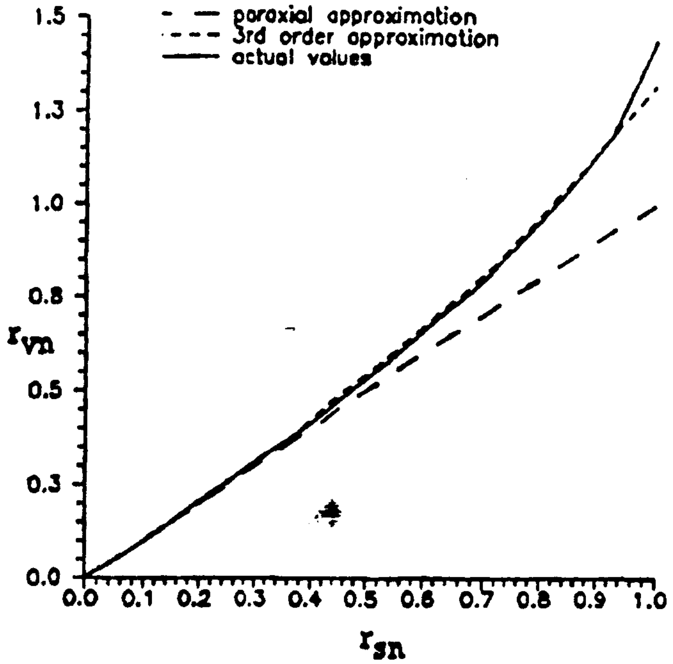

Graph of r_{vn} vs. r_{sn} for D ($k_{vs} = 0.32$)

Then we used the distortion function to pre-distort a square grid so that when viewed through the optics, it would appear straight. We verified that this works by looking through the optics at the pre-distorted grid printed on paper. We have not yet implemented the pre-distortion function for the screens of our HMDs.

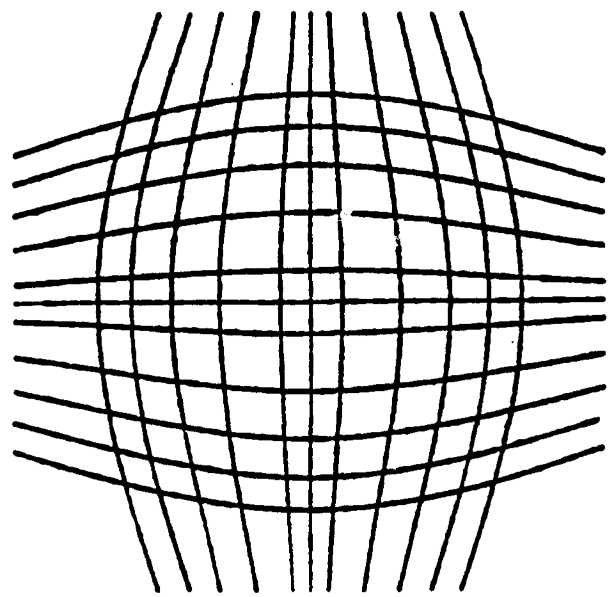

The following drawing shows our old see-through HMD, the point of departure for this work.

UNC See-Through Head-Mounted Display

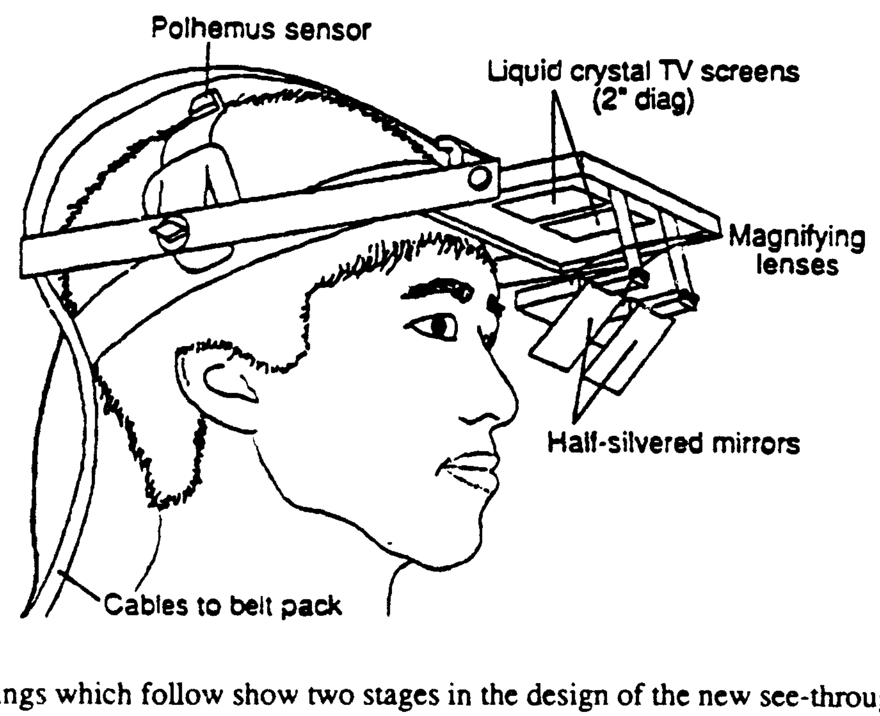

The two ray-tracings which follow show two stages in the design of the new see-through HMD optics.

The two ray-tracings which follow show two stages in the design of the new see-through HMD optics.

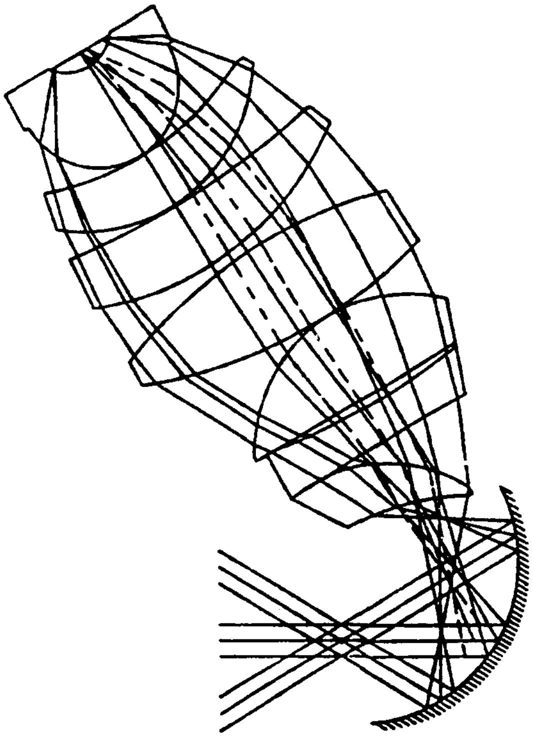

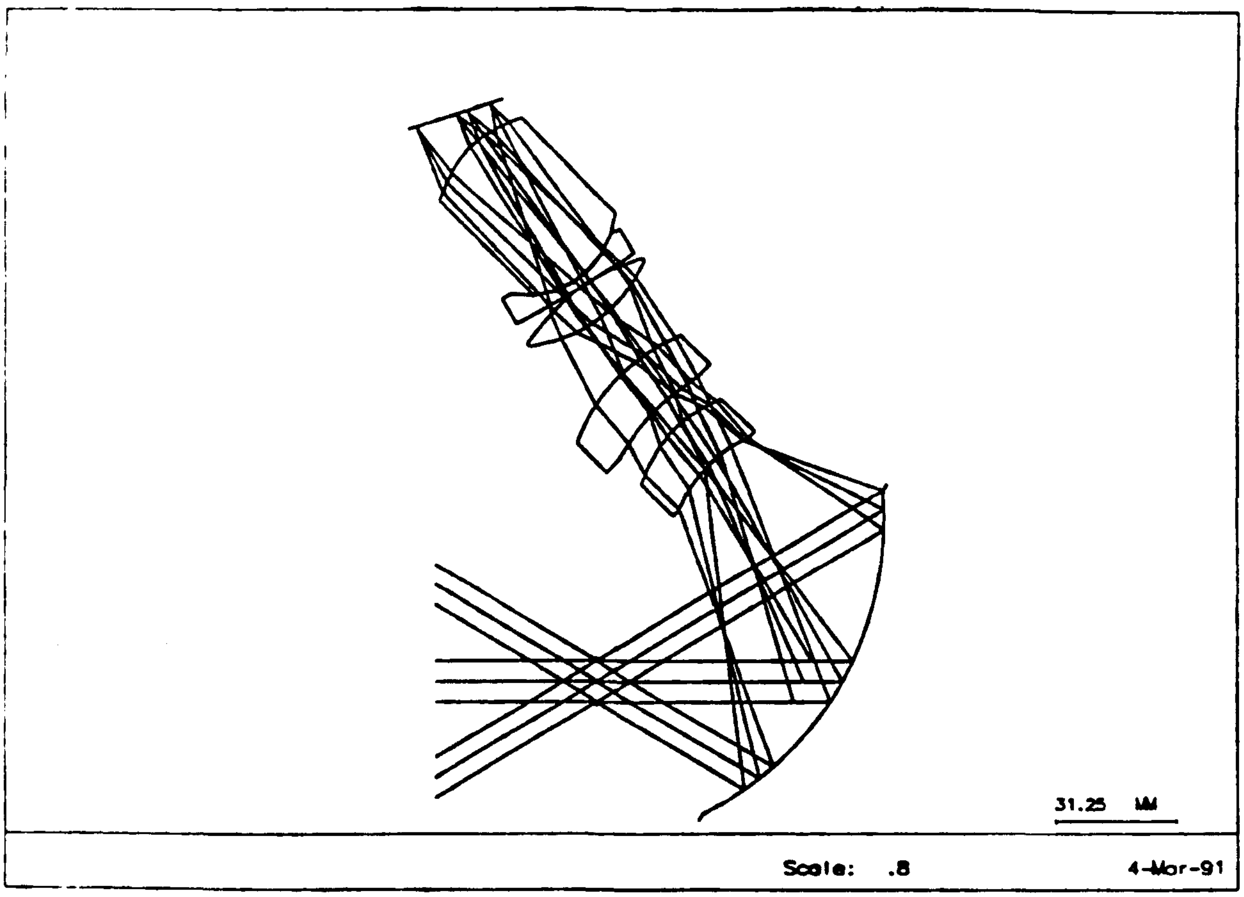

The two diagrams below shows a point scatter point from the optics analysis of the current optics design, and graphs describing the aberrations in those scatter plots.

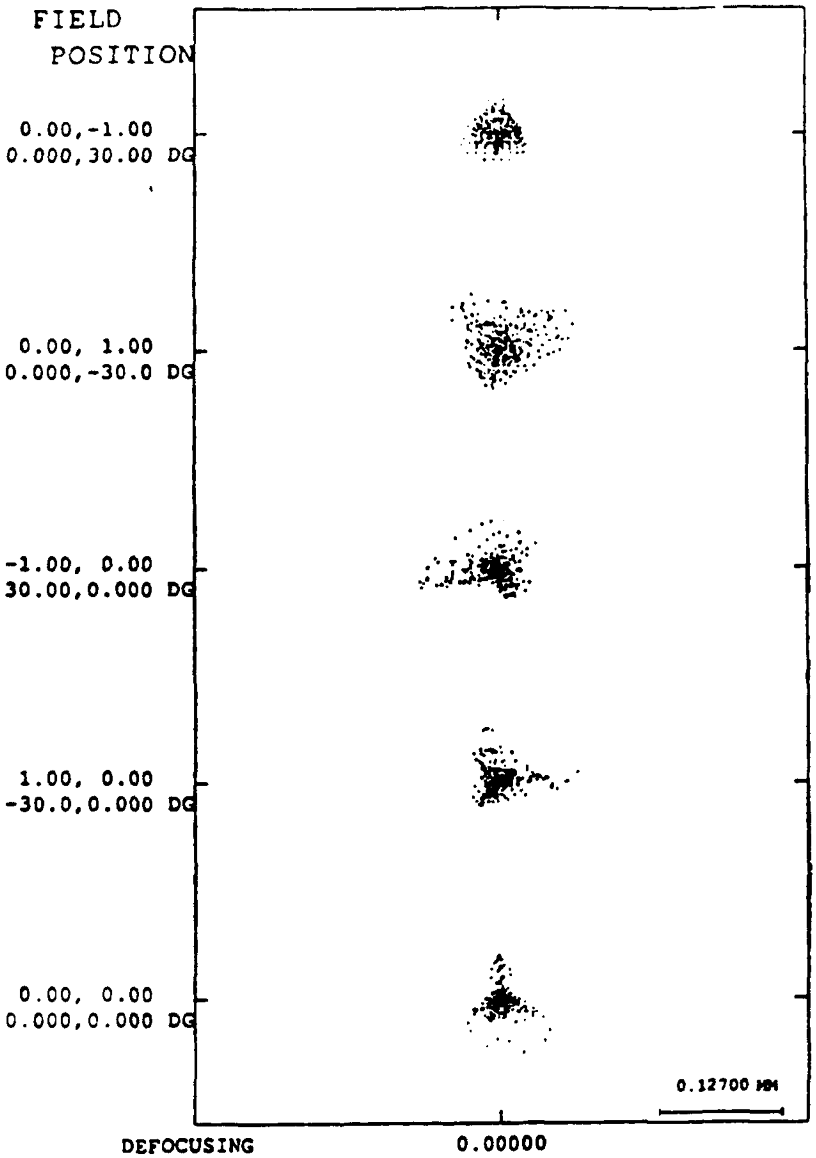

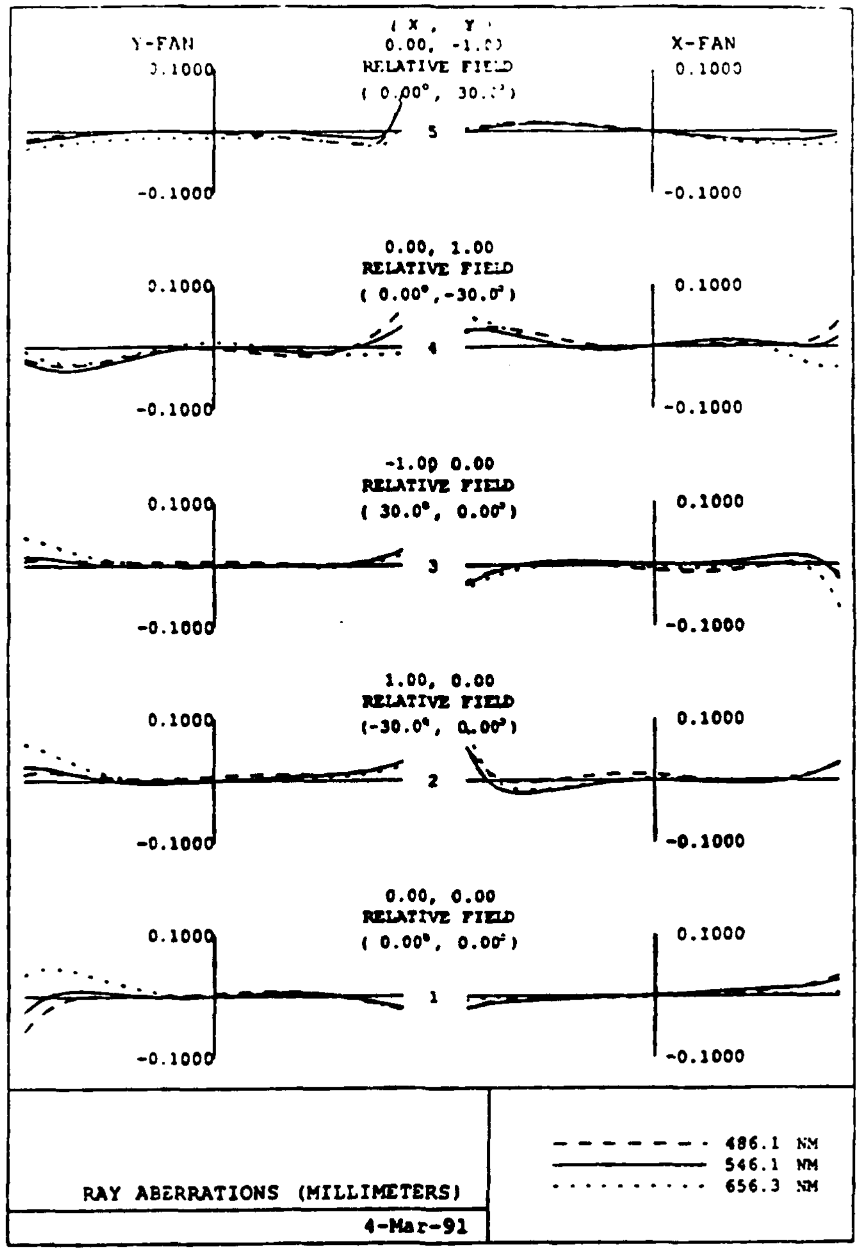
Y-FAN
X-FAN
(X , Y)
0.00, -1.00
RELATIVE FIELD
(0.00°, 30.0°)
0.1000
-0.1000
5
0.00, 1.00
RELATIVE FIELD
(0.00°, -30.0°)
4
-1.00 0.00
RELATIVE FIELD
(30.0°, 0.00°)
3
1.00, 0.00
RELATIVE FIELD
(-30.0°, 0.00°)
2
0.00, 0.00
RELATIVE FIELD
(0.00°, 0.00°)
1
486.1 NM
546.1 NM
656.3 NM
RAY ABERRATIONS (MILLIMETERS)
4-Mar-91

Finally, the drawing below shows the application we are aiming for with this see-through HMD -- X-ray vision for an obstetrician using a HMD and ultrasound scanner.

Possible application of portable personal visualization systems in the near future: Ultrasound Examination

Physician wearing a portable personal visualizer observes a fetus by ultrasound imaging, allowing him to correlate the 3D visual information with that obtained by tactile examination. The hand-held transducer is tracked stereoptically by miniature cameras mounted on the user's visor. Two upward-looking cameras on the head-gear track an array of IR LED beacons on the ceiling as the user moves about the room. Conventional ultrasound equipment can be seen in the background.

4.3 Tracking

Our goal is the development of a tracking system that can report head position and orientation accurately (±2 mm translation, .1° rotation) and rapidly (> 30 Hz) over a large working volume (20' x 20'). Two complementary technologies, each optical in nature, are currently being explored. One, referred tc as the optoelectronic tracker, employs head-mounted lateral-effect photodiodes to image ceiling mounted light emitting diodes (LEDs). Computation of head position relies on the photodiodes' ability to measure the photocoordinates of an LED's image. The other optical tracking technology, referred to as Self-Tracker, employs linear arrays of custom image sensors to measure relative motion. Whereas the optoelectronic tracker relies on a structured environment, i.e., a ceiling filled with LED beacons, Self-Tracker is able to measure image shift from within an unstructured environment. Another tracking technology that can operate in an unstructured environment is that of an inertial guidance system. A preliminary search of the avionics literature has revealed an industry-wide interest in miniaturize inertial guidance systems based on laser ring gyroscopes.

Optoelectronic Tracker

We are implementing an optoelectronic tracking system that adopts an "inside-out" paradigm introduced by Gary Bishop and Henry Fuchs [Bis84]. When complete, the system will feature thousands of infra-red small light emitting diodes (LEDs) mounted in a room's ceiling (see figure below). The LEDs serve as global position references for a head-mounted display that is augmented with four, upward looking image sensors. Each sensor resembles a small camera in that its lens focuses infrared light onto a lateral-effect photodiode's surface. The photodiode, in turn, reports the photocoordinates of an LED that has been momentarily turned on. The photocoordinates associated with 3 or more LEDs, along with their absolute position in the room, allows the position and orientation of the head-mounted display to be computed.

The advantage of this approach lies in its cellular nature. Typically, optical tracking is done with fixed cameras and moving LEDs. That is, the cameras are "outside" looking "in". A fundamental limitation exists with this approach: the working volume is limited to the cameras' field of view. Furthermore, to achieve high resolution, long focal lengths are required and a reduction in working volume follows. Reversing the situation, i.e., moving the cameras and fixing multiple LEDs, solves this dilemma. Long focal focal lengths can be used to achieve high resolution and as the cameras move, the set of LEDs in a camera's field of view is constantly changing. Jih-Fang Wang demonstrated the feasibility of "inside-out" tracking in his dissertation **[Wan90].**

Using three Hamamatsu [Ham85] camera bodies and three LED beacons, Wang showed that the lateral-effect photodiode could be used in an outward looking image sensor to measure head position with high accuracy and resolution. To compute the location of the three camera cluster, Wang modified Church's algorithm, a technique developed in 1945 by Earl Church [Chu45] to compute the location of a camera used in aerial photogrammetry.

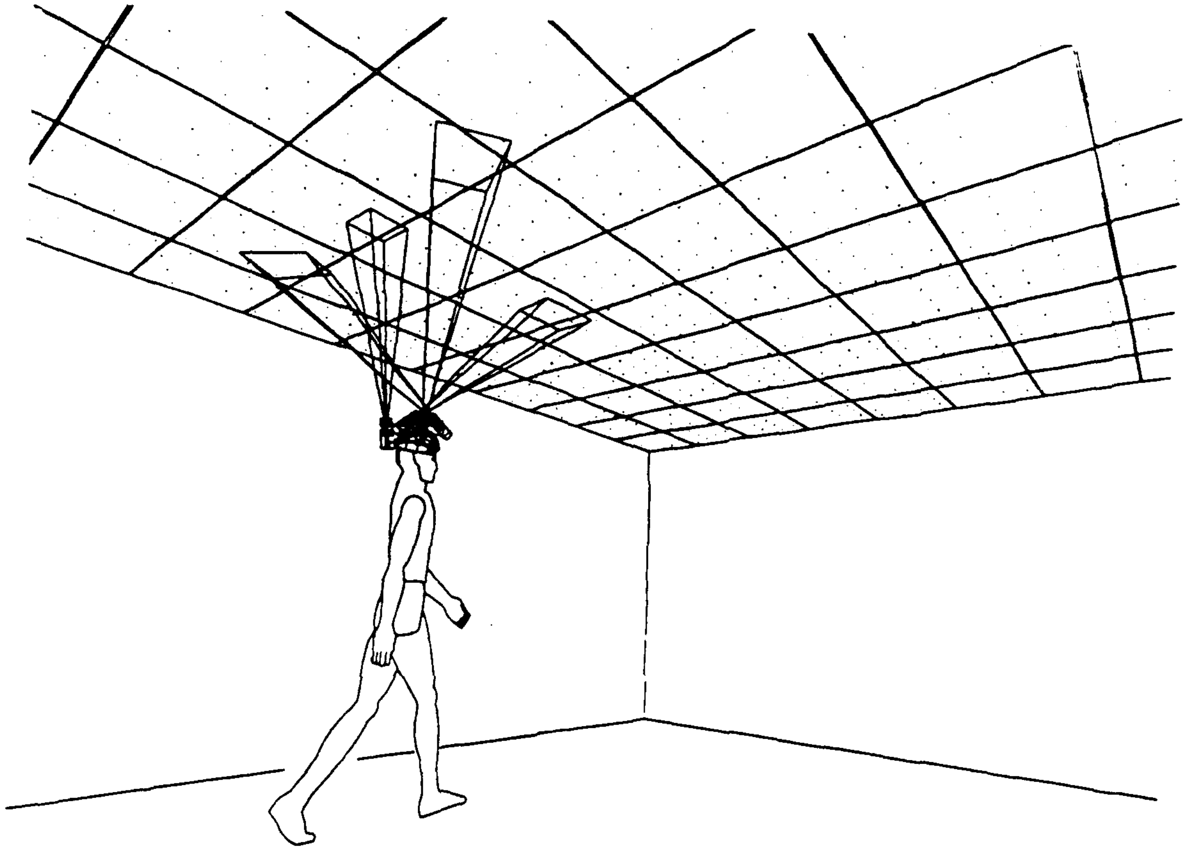

Figure 1. With "inside-out" tracking, multiple head-mounted image sensors view LED beacons suspended in ceiling panels. Each sensor's field of view is mapped onto the ceiling to show that more than 3 LEDs are typically visible.

Presently, we are extending Wang's work with the goal of developing a room-sized optoelectronic tracker for researchers and users in our graphics laboratory. Several efforts are underway to realize this goal. One involves the development of rugged electronic hardware to acquire image data. This involves managing multiple image sensors, selectively flashing one LED in a ceiling of thousands, and getting data in and out of a suitable workstation. Another task involves the creation of a 20'x20' ceiling structure that precisely locates LEDs. To ensure compatibility with our building's 2'x2' ceiling tiles, the structure will initially resemble a grid of 2'x2' panels suspended from a structure just below the existing ceiling. This structure will have provisions for leveling each panel to achieve our goal of ±2mm absolute position tolerance over a 20'x20' ceiling. In the future, we hope to develop techniques that will allow us to relax our tolerance on LED placement. This would allow the 2'x2' panels to be installed in the existing suspended ceiling grid. Another task involves the real time computation head position. Image data from a variable number of sensors and LEDs must be converted into head position in 5-10ms. A final task involves camera calibration. As received from Hamamatsu, the image sensors are limited by a 1% measurement error. For the system to work, this error must be reduced to .1%.

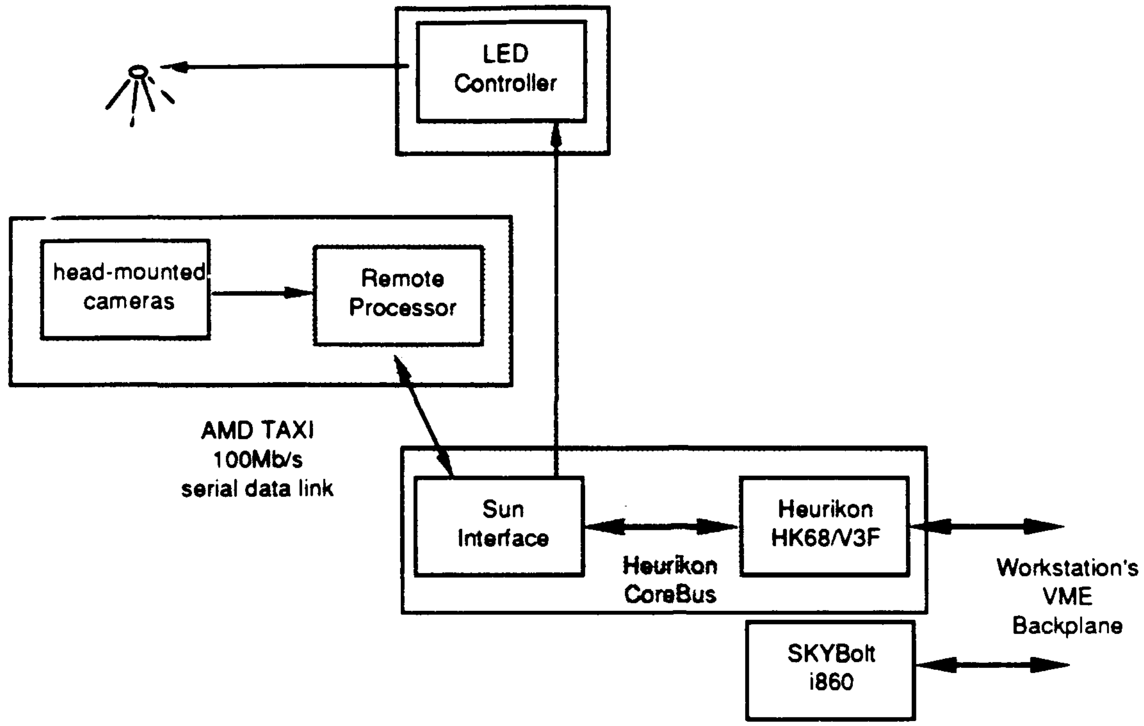

Figure 2. LED controllers, a remote analog processor, a 68030-based single board computer, an an i860-based array processor constitute the tracker controller's 4 major subsystems.

The tracker controller involves four subsystems: each ceiling panel has an LED control module, the user wears an analog processor (Remote Processor) that converts sensor data into a digital format, a 9U VME board (Sun Interface) connects a 68030-based single board computer to the Remote Processor and to the LED controllers, an i860-based numeric processor accepts global LED positions and photocoordinates from the 68030 and computes the position and orientation of the head. Shaded blocks in Figure 2 are nearly complete.

Up to eight Hamamatsu image sensors can be attached to the Remote Processor. Worn by the user, it consists of a bank of analog multiplexers, 4 parallel A/D converters, and a serial communication port. Analog voltages emerging from the sensors are digitized and are then relayed to the Sun Interface via a 100 Mbit/s serial data link. This link is based on the Advanced Micro Devices 7968/7969 products.

The Sun Interface consists of a 68030 microprocessor, a serial interface to the Remote Processor, and a parallel interface to a daisy-chain of ceiling panel controllers. The Sun Interface flashes an LED by writing its address and current level to the parallel interface. Software running on the board is responsible for flashing LEDs and for preprocessing the image data acquired by the Remote Processor. The result of this preprocessing is a list of LED absolute coordinates and the associated photocoordinates of each LED that is imaged.

The list is transformed into head position and orientation by an i860-based array processor. A board-level product has been selected and is currently being procured. During this reporting period, two vendors of i860-based array processors benchmarked samples of our code for computing position and orientation. Each reported an execution time of 40ms. For the next several months, Carney Clegg will be optimizing existing C routines to take advantage of the i860's architecture. His goal will be to achieve an execution time of 5-10ms. Each vendor believed this to be a reasonable goal. If so, our update rate will be 100-200Hz.

Ceiling Structure

We have begun the design of a suspended ceiling (see figure below) structure that can provide absolute position accuracy of ±2mm on each LED beacon This is necessary for LED misplacement maps almost directly into head position measurement error. Our concept calls for each panel to measure 2'x2' and to be constructed of a composite material. Two panel types are being considered: aluminum sheet over an aluminum honeycomb and fiberglass cloth over a Nomex core. An interested fabricator has been located in Jacksonville, Florida that can process such materials.

LED to LED spacing on a given panel, as well as panel to panel relationship, is established during the manufacturing process. In one operation, LED holders and inter-panel locating features will be bonded into a composite sheet. Panels are connected to one another by small spiders that are hung from a subceiling and mate with features bonded into each panel's corners. Panel leveling is achieved by jacking spiders up or down. A Spectra-Physics LaserLevel will be rented to assist in the leveling process.

Software for computing head position

An improved algorithm for computing the head-position and orientation is being developed that will accommodate more LED placement error than Wang's modified version of Church's algorithm. The algorithm is a generalized version of the technique known as *space-resection by collinearity* **[Gho88]**. For the problem of computing a camera's position and orientation given the photocoordinates of 3 or more landmarks (known as space-resection), this technique has several advantages over Church's method. This significant feature of this technique is that if more than 3 landmarks can be observed, the additional information can be used in the solution. Furthermore, the 6 unknowns are solved for in one step, whereas Church's algorithm directly solves for only the linear degrees of freedom. Least-squares techniques are used to accommodate the over-determined matrix equations that result from more than three landmarks.

Space-resection by collinearity has been generalized to allow the use of multiple image planes, as well as multiple landmarks per image plane. This will allow us to use more than three cameras, and if each camera is capable of imaging more than one LED, then the additional information is used to compute the camera assembly's position and orientation. Simulations have shown that when this approach is used, the effects of beacon placement error are less than with the generalized version of Church's algorithm.

The chief disadvantage of space-resection by collinearity is computational complexity. As with Church's method, the technique is iterative in nature. However, whereas in Wang's work the number of beacons was fixed at 3, space-resection by collinearity admits as many beacons as can be imaged. The dimensions of matrices which must be manipulated increases by two for each LED imaged. Fortunately, the one matrix inversion required involves a matrix of fixed size: 6x6. Prototype C code has been benchmarked by two manufacturers of numeric processor products based on the Intel i860 microprocessor. Computation times of roughly 40ms were reported by each. For the remainder of this year, one student, Carney Clegg, will be optimizing this code with a goal of 5-10ms as an execution time for each sample.

Calibration

In preparation for the full-scale system, considerable effort is being directed towards camera calibration. Wang showed that the image plane of each camera must have an accuracy of 1 part in 1000 to satisfy the overall system requirements. Antonsson, formerly of MIT's Biomechanics Laboratory and now with the Division of Engineering and Applied Science at the California Institute of Technology, found that this was impossible with off-the-shelf, optoelectronic equipment **[Ant89]**. For example, the lateral-effect photodiode possesses a systematic error that is as great as 1 part in 40. When this sensor is combined with a non-metric lens and a camera body whose tolerances are no better than .5mm, the result falls far short of our goal. Antonsson experimentally derived a two-dimensional calibration table for Selspot cameras [Sels88] that are nearly identical to ours and achieved the goal of 1 part in 1000 accuracy.

Fixtures for camera calibration are currently being constructed. An existing motor-driven x-y positioning table is being installed on an optical bench in our laboratory. The x-y table will be used to position an LED in the object space of each Hamamatsu camera. A calibration table will be constructed for each using techniques similar to those of Antonsson.

Self-Tracker

The self-tracker is one possible inside-out tracking system which may replace, or augment, the optoelectronic system. Its goal is to eventually eliminate fixed beacons and perform tracking using features found naturally in the environment. The self-tracker could then be taken to any room, hallway, or even outdoors.

The self-tracker will use a cluster of custom made VLSI chips which will be mounted on the head-mounted display and look out on the environment in multiple directions. The cluster of chips will report the apparent shift of the image seen by each chip during successive clock cycles. A host computer will collect the information and calculate the distance the user has travelled relative to his previous position. Approximately 1000 frames per second will be required to track the fastest possible movements reported in human motion literature. This system may not work by itself because of accumulated drift. For the initial versions, a small number of beacons may be placed in the environment to allow for error correction.

Each self-tracker chip will contain a one-dimensional array of photosensors and supporting image processing circuitry. A one-dimensional array will be used so that a sufficiently large number of pixels can be placed on the chip and so that image correlation can be performed in real time. The correlation between successive frames will be accomplished by converting the image to a digital representation and then successively shifting the position of one frame and reporting the shift which corresponds to the closest match to the other frame.

The tracker chips will be mounted in pairs so that stereopsis can be used to measure the distance to the observed surface in the scene. Each tracker chip will report the observed shift in terms of the number of pixels of shift, so the distance value will be required to convert the shift from pixels to a physical measurement. One chip in each pair will transmit the its image's digital representation to the other chip. The second chip will then calculate the correlation between the two images in the same manner that it calculated the correlation between subsequent images on the same chip.

The diagram below is a block diagram of a self-tracker chip. A one-dimensional photosensor array generates an array of analog voltages which represent the incident light on the array. The edge detection block produces a digital representation of the image. Then, the image comparison block compares the current image with the image seen during the last time frame. The shift between the two images is reported, and the current frame is stored for comparison with the next frame. Also, each image is sent to its stereo partner to calculate the distance to the observed image.

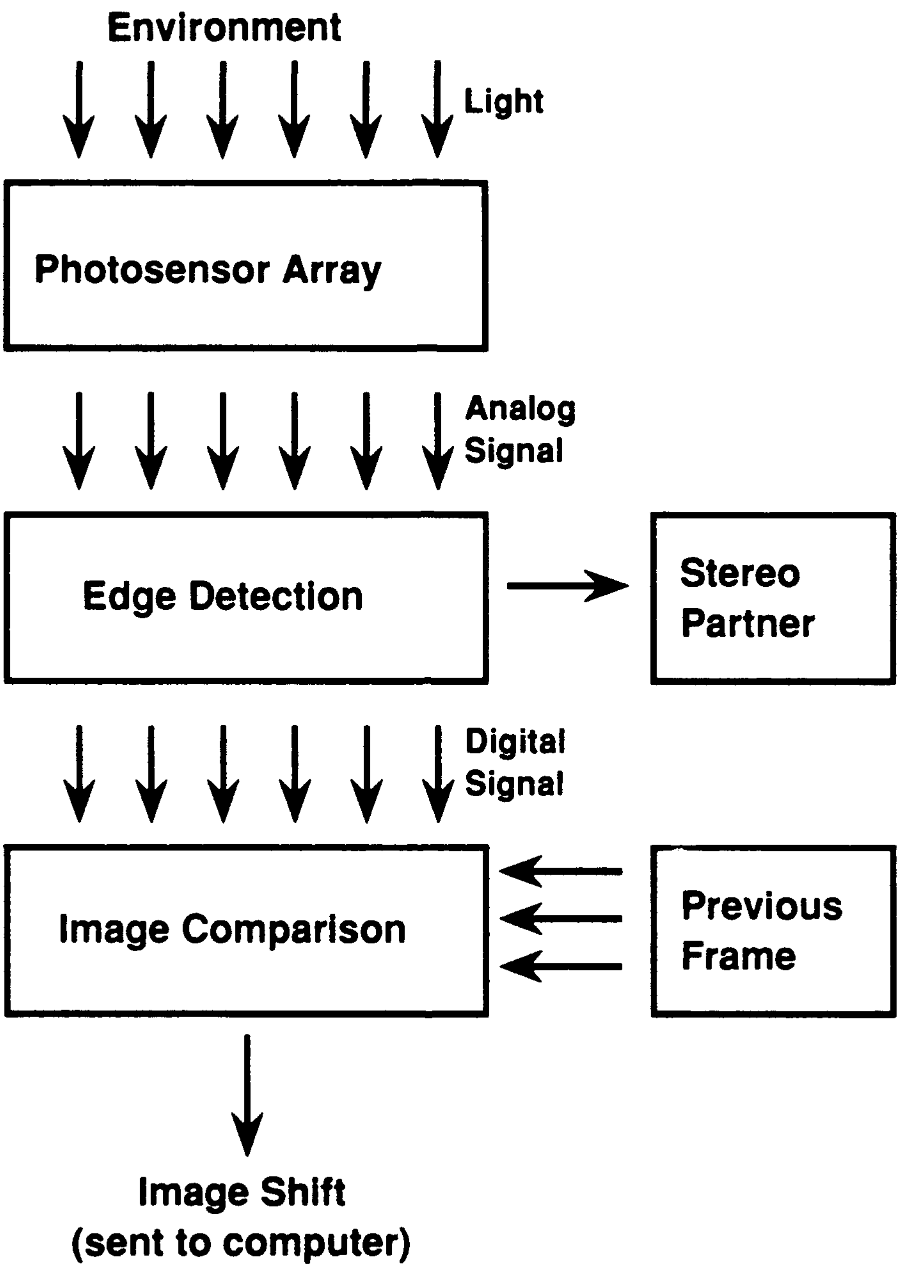

Figure 3. Architecture of the Self-Tracker custom VLSI sensor.

Over the past six months, research has progressed in two areas. Supporting hardware is being constructed which will be placed on the user and will control the communication between the self-tracker chips and the host computer. The tracker chips will communicate with the host computer via a serial link being controlled by AMD's TAXI chips. Also, a computer software model of the self-tracker chip has been constructed. The software model simulates the operation of the self-tracker chip using video images captures from a TV camera. The software model is being used to refine the image processing and correlation circuitry to reduce the level of errors.

Inertial Guidance Systems

Relative motion sensors based on accelerometers and gyroscopes may offer an alternative to Self-Tracker. Several companies are currently developing extremely small inertial navigation modules based on either laser ring gyroscopes or fiber optic gyroscopes and miniaturized servo-accelerometers. Combined with a low bandwidth absolute tracking technology to compensate for drift, a hybrid approach employing inertial techniques could possibly provide the same level of performance as the optoelectronic tracker.

The optoelectronic tracker, degraded by removing ceiling panels, is envisioned as a candidate for the role of the absolute position reference. That is, LED beacons would be arranged such that intermittently, a sufficient number of LEDs would fall into the sensor's field of view to compute the absolute position and orientation of the head. In between these absolute position fixes, inertial sensors would be used to measure displacements.

References

[Ant89] E. K. Antonsson and R.W. Mann. Automatic 6-D.O.F. kinematic trajectory acquisition and analysis, J. Dynamic Systems, Measurement, and Control, 1(2), March 1989, pp. 31-39.

[Bis84] T.G. Bishop and H. Fuchs. The self-tracker: A smart optical sensor on silicon Proc. Conf. on Advanced Research in VLSI, MIT Press, 1984.

[Chu45] Earl Church. Revised Geometry of the aerial photograph. *Bulletins on Aerial Photogrammetry*, No. 15, Syracuse University, 1945.

[Gho88] Sanjib K. Ghosh. *Analytical Photogrammetry,* 2nd. Ed., Pergamon Press, New York, 1988.

[Ham85] Hamamatsu Photonics K.K., Image Processing and Measuring Systems. C2399 Position Sensor Instruction Manual, no. 5210-005, Hamamatsu City, Japan.

[Sels88] SELCOM. SELSPOT II HARDWARE and MULTILab Software, Southfield, Michigan, 1988.

[Wan90] J. F. Wang. A real-time optical 6D tracker for head-mounted display systems. PhD. thesis, University of North Carolina at Chapel Hill, 1990.

5.0 Dissemination of Research

5.1 Publications

Chung, J. C., M. R. Harris, F. P. Brooks, Jr., H. Fuchs, M. T. Kelley, J. W. Hughes, M. Ouh-Youn C. Cheung, R. L. Holloway, and M. Pique. "Exploring Virtual Worlds with Head-Mounted Displays," reprinted in *Proc. Ars Electronica,* 2, Linz, Austria: Virtuelle Welten, Sept. 1990, 138–14

Holloway, Richard L. "Art-Related Virtual Reality Applications at the University of North Carolina at Chapel Hill," *Proc. 2nd International Symposium on Electronic Art,* Groningen, Holland, 12–16 No 1990.

Robinett, Warren. "Head-Mounted Display Project Summary," *Proc. Ars Electronica,* 2, Linz, Austria: Virtuelle Welten, Sept. 1990, 119–122.

Robinett, Warren, and Jannick P. Rolland, "A Computational Model for the Stereoscopic Optics of a Head-Mounted Display," *Proc. SPIE Conf on Stereoscopic Displays and Applications*, **1457**, San Jose, California, Feb. 26, 1991

5.2 Presentations

By:	Warren Robinett
Topic:	Head-Mounted Display Project Summary
Event:	Ars Electronica 1990
Place:	Linz, Austria
Date:	Sept. 1990

By:	Warren Robinett
Topic:	Synthetic Experience
Event:	Mondi Virtuali Symposium
Place:	Venice, Italy
Date:	Nov. 1990

By:	Richard Holloway
Topic:	Art-Related Virtual Reality Applications at UNC at Chapel Hill
Event:	2nd International Symposium on Electronic Art
Place:	Groningen, Holland
Date:	12–16 Nov. 1990

By:	Warren Robinett
Topic:	A Computational Model for the Stereoscopic Optics of a Head-Mounted Display
Event:	SPIE Conference on Electronic Imaging
Place:	San Jose, California
Date:	Feb. 26, 1991

5.3 Visitors who observed HMD demos

- Bob Andron, architect
- Dieter Breade and photographer, Stern Magazine, Hamburg, Germany
- Charles Brownstein, National Science Foundation
- Mark Cutter and Pete Litwinowicz, Apple Computer
- Max Donath, Director of Robotics and Human Motion Labs, Univ. of Minnesota
- Glen Emory, Insight Magazine
- J. Folda, UNC Art Department
- Senator Al Gore
- Klaus Hoog, graphic designer
- Diane Jessen and colleagues, Rensselaer Polytechnic Institute
- Drs. R. Kikinis and F. Jolesz, Brigham and Women's Hospital, Boston
- Tom Kominski, Astronautics Corp., Madison, WI
- Edouard Launet, Les Echos, Paris
- Pete Meenan, Bill Lorensen, Nelson Corby, and Art Chin, General Electric Research and Developme
- Tatsuo Miyazawa, IBM Tokyo Research Lab
- Mehran Moshseghia and Benjamin Zhu, Philips Medical Research
- Doug Schiff et al, Sun Microsystems
- Steve Squires et al, DARPA
- Fred Symington, DEC
- Larry Yaeger, Vivarium Project, Apple Computer
- Armed Forces Institute of Pathology
- DARPA Contractors Meeting
- Disney Imagineering
- East Rowan High School students, Salisbury, NC
- Gifted and talented 6th graders, Glenwood School
- IBM Science Advisory Committee
- Japanese Public Television video crew
- NC State University computer science undergraduates
- NC State University mechanical engineering undergraduates

Appendix A

Paper: A Computational Model for the Stereoscopic Optics of a Head-Mounted Display
by Warren Robinett and Jannick P. Rolland

Appendix B

Proposal to the SIGGRAPH 1991 Virtual Reality Applications Gallery
(7 hands-on exhibits proposed)

A Computational Model for the Stereoscopic Optics of a Head-Mounted Display

Warren Robinett
Jannick P. Rolland
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA

ABSTRACT

For stereoscopic photography or telepresence, orthostereoscopy occurs when the perceived size, shape, and relative position of objects in the three-dimensional scene being viewed match those of the physical objects in front of the camera. In Virtual Reality, the simulated scene has no physical counterpart, so orthostereoscopy must be defined in this case as constancy, as the head moves around, of the perceived size, shape and relative positions of the simulated objects.

Achieving this constancy requires that the computational model used to generate the graphics match the physical geometry of the head-mounted display being used. This geometry includes the optics used to image the displays and the placement of the displays with respect to the eyes. The model may fail to match the geometry because model parameters are difficult to measure accurately, or because the model itself is in error. Two common modeling errors are ignoring the distortion caused by the optics and ignoring the variation in interpupillary distance across different users.

A computational model for the geometry of a head-mounted display is presented, and the parameters of this model for the VPL EyePhone are calculated.

1. INTRODUCTION

1.1 The Problem: Computing the Correct Stereoscopic Images in Virtual Reality

As you move through the world, images of the objects that surround you fall onto your retinas. As you move past a fixed object, seeing it from various angles, the size and shape of the images on your retinas change, yet you effortlessly and unconsciously perceive the object to have a stable position, shape and size. This innate perceptual ability, honed by your daily experience ever since infancy, is so fundamental and habitual that it seems almost absurd to talk about objects which could change their position or shape or size depending on how you moved your head.

Yet the current state of the art in Virtual Reality (VR) gives us simulated objects which change their position, size and shape as the head moves. The location of these objects appears to change as the head moves around, and their size and shape appear to change depending on whether they are being viewed directly in front of the user's head or off to the side.

In Virtual Reality, a Head-Mounted Display (HMD) and a head-tracker are used to rapidly measure head position and create an image for each eye appropriate to its instantaneous viewpoint. The HMD-user can then see simulated objects from different points of view as the head moves. However, it is difficult to correctly calculate the images to be painted onto the display screens of the HMD. The user's eyes and brain (and vestibular system) are very sensitive to inconsistencies.

The computational problem of calculating the correct stereoscopic images in VR -- getting the perceived objects to have the right position, size and shape -- is the same problem that faces the designers of stereoscopic photography and telepresence systems. For these systems, orthostereoscopy occurs when the perceived size, shape, and relative position of objects in the three-dimensional scene being viewed match those of the physical objects in front of the camera. In Virtual Reality, the simulated scene has no physical counterpart, so orthostereoscopy must be defined in this case as constancy, as the head moves around, of the perceived size, shape and relative positions of the simulated objects. To calculate orthostereoscopic images, the display code must precisely model the geometry of the HMD system upon which the images will be viewed. This includes the relative positions of the display screens, the optics and the eyes. The relationship between the screen and the virtual image of it must also be modeled.

This paper addresses only the static image generation problem. To simulate objects that are spatially stable, temporal problems must also be solved, but those problems are outside the scope of this paper.

1.2 Prior Work

Since Ivan Sutherland built the first HMD in 1968, several HMD systems have been built. The display code for each system defined an implicit model for the particular geometry of each HMD. The authors of the display code for each system structured the code as they judged appropriate, and it is difficult to know the precise details of their display code from what has been published. It appears that most HMD systems treated their optics simply as a magnifier, ignoring distortion introduced by the optics.

In Sutherland's HMD, tiny half-inch monochrome CRTs were the display devices, and the virtual images seen through the optics subtended an angle of 40 degrees and appeared to be at a distance of 18 inches in front of the eyes [SUTH65] [SUTH68]. Half-silvered mirrors superimposed the computer graphics onto the user's direct view of the real world. Later versions of the HMD were stereoscopic. The stereoscopic HMD had both a mechanical adjustment for inter-pupillary distance (IPD) and a software adjustment for the virtual eye separation. This HMD system was moved to the University of Utah, and essentially the same system was used by several students there [VICK74] [CLAR76].

In 1983, Mark Callahan at MIT built a see-through HMD similar to Sutherland's [CALL83]. It used half-silvered mirrors mounted on eyeglass frames, 2-inch monochrome CRTs, and a bicycle helmet. An optical disk was used to rapidly display pre-recorded images in response to head movements.

In 1985 at NASA Ames Research Center, Mike McGreevy and Jim Humphries built a non-see-through HMD from monochrome LCD pocket television displays, a motorcycle helmet, and the LEEP wide-angle stereoscopic optics [HOWL83]. This HMD was later improved by Scott Fisher, Warren Robinett and others [FISH86]. The display code for this system treated the LEEP optics as a simple magnifier. The LEEP optics system has very large exit pupils and therefore no mechanical IPD adjustment.

At Wright-Patterson Air Force Base in the Seventies and Eighties, Tom Furness directed a program which developed prototype HMDs for use in military aircraft [GLIN86]. The system was called Visually-Coupled Airborne Systems Simulator (VCASS), and several prototype see-through HMDs with custom-designed optics were developed there [KOCI86].

CAE Electronics of Quebec has developed a fiber-optic helmet-mounted display system (FOHMD), intended for flight simulators [CAE86] [HEND89]. Four light valves drive the HMD through fiber-optic cables and pancake optics allow the user to see-through to the flight simulator's control panel. There is a mechanical adjustment for IPD. The binocular field-of-view (FOV) is 135 degrees horizontally by 64 degrees vertically, and it also has a 25 x 19 degree high resolution inset field.

Several prototype HMDs have been constructed here at the University of North Carolina at Chapel Hill [HOLL87] [CHUN89]. In 1985, a see-through HMD was made from color LCD displays, half-silvered mirrors, magnifying lenses, and a pilot's instrument training hood. The FOV was approximately 25 degrees horizontally. A later model, built at the Air Force Institute of Technology (AFIT) with UNC collaboration, was made from color LCDs, very strong reading glasses, and a bicycle helmet. Its FOV is about 55 degrees horizontally and it is not a see-through HMD. We are currently designing a see-through HMD for medical imaging applications. It will incorporate custom-designed optics.

In 1988, VPL Research of Redwood City, California, began selling the EyePhone, the first commercially available HMD [VPL89]. It uses color LCD displays and the LEEP optics. It attaches to the head with a rubber diving mask and fabric straps. It is not see-through.

In 1989, Eric Howlett, the inventor of the LEEP optics, put together a commercial HMD, the LEEPvideo System I. It used monochrome LCD displays, the LEEP optics, and a head-mounting apparatus designed by Howlett. Howlett subsequently introduced improved models which use color LCDs and have a wider FOV [LEEP90].

In 1989, Reflection Technologies of Waltham, Massachusetts, produced a product called Private Eye, a single eye monochrome HMD [TIME89]. It uses a vibrating mirror and an LED linear array to produce a 2D image. The horizontal FOV is about 25 degrees.

1.3 Remaining Problems

Generating correct stereoscopic images for an HMD is a difficult task. The display code for each HMD system embodies an implicit computational model of the geometry of the HMD, and there are many sources of error which must be compensated for. In current practice, most of these models are inadequate because they ignore certain sources of error. Also, since these models are embodied only in the display code of the HMD systems, they are difficult to comprehend, and are not accessible to most people. It is difficult to compare the computational models of different HMD systems.

The display software often ignores the system optics. But because the optics actually do affect the images seen by the eyes, some of the parameters in the display software are tweaked to get the convergence and FOV of the HMD to be roughly correct. This type of measurement of the parameters of the HMD system by subjective calibration by the users is inaccurate compared with calculating the model parameters from the specifications of the optics and display screens.

Another problem is that most current HMD systems ignore the variation in IPD across different users. In this case, wide and narrow-eyed users will have different size perceptions of the same simulated objects.

This paper presents an explicit computational model for generating orthostereoscopically correct images. Implementing display software which follows this model will produce stereoscopic images which are orthostereoscopic -- simulated objects will be perceived as three-dimensional and will be undistorted and correctly sized.

We first survey the various sources of error which cause incorrect stereoscopic images to be generated. We then introduce a computational model of the geometry of an HMD which models the optics, the distance between the user's eyes, and the relative positions of the eyes, optics, and display screens. This model allows correct orthostereoscopic images to be calculated. Finally, we calculate the model parameters for the VPL EyePhone.

2. SOURCES OF ERROR

There is a very precise correlation between the movements of one's head and the images of an object that are formed on the retinas from moment to moment. This correlation can be described by simple geometry: the object's images are projected onto the retinas, and the retinal images depend only on the object's shape and the relative position of the two eyes with respect to the object. An HMD system attempts to mimic this geometry, painting images onto display screens in front of the eyes to fool the eyes and brain into perceiving three-dimensional objects. If the wrong images are painted onto the screens, the user is not able to perceive the simulated object correctly. The object will either be distorted, or there will be no perception of a 3D object at all.

The wrong images are painted onto the screens either because of errors and inaccuracies in the head-tracking, or because of errors in the software which controls the image generation. While the tracking hardware can introduces significant error, it is the display software which is the subject of this paper.

We will introduce a **computational model for a Head-Mounted Display**, and say that the display software implements this model. For the display software to generate the images required to give the HMD-user the perception of undistorted objects, the software must take into account the physical geometry of each hardware component that affects the final image seen by the eyes. This geometry includes the display screens, the optics used to image the displays, and the placement of the displays with respect to the eyes. Before introducing the computational model, we discuss some common errors in the display code for HMDs.

2.1 Incorrect convergence

Both eyes are necessary for stereoscopic vision. When the eyes are focussed on a distant object, the lines of sight are roughly parallel, and when focussed on a near object, the lines of sight converge. The nearer the object, the greater the convergence.

A stereoscopic HMD has, for each eye, a display screen and an optical system through which the screen is viewed. If the optical axes were parallel for the two optical systems and if the optical axis passed through the center pixel of each screen, then by illuminating those two center pixels, the user would see a stereoscopic image of a point of light at infinity. This would be more or less like looking at a single star in the night sky. However, many HMDs do not satisfy those two conditions: turning on the center pixels would either produce a percept of a not-so-distant point of light in front of the user, or else be too divergent to fuse at all into a stereoscopic percept.

In addition to the horizontal alignment problem related to convergence and divergence, there can also be a vertical misalignment between the two eyes. This is called dipvergence.

Creating stereoscopic images with the correct convergence, when the optical axes are not parallel or the centers of the displays are offset from the optical axes, requires corrective transformations in the computational model. Neither of these properties is a mistake in the design of an HMD, they just make the computational model a little more complicated. In fact, many current HMDs have non-parallel optical axes and off-center screens.

2.2 Accomodation not linked to convergence

The eyes converge on a nearby object by rotating inward; the lenses of the eyes simultaneously accomodate to the distance of the object to bring it into focus. Thus convergence and accomodation are normally linked. However, in an HMD system, each eye sees the virtual image of a display screen. With respect to focus, the entire virtual image appears at a fixed distance from the eye. (In a physical scene, different parts of the scene will be in focus at different accomodation depths.) Hence, the HMD user must learn to decouple accomodation and convergence. This problem cannot be overcome with any currently-used display device. Until a display device with variable focus is developed, this problem must be accepted as a limitation of HMDs.

2.3 Incorrect field-of-view

The display for a single eye in an HMD has an FOV, which is the angle subtended by the virtual image of the display screen, as seen by the eye. There is a horizontal and a vertical FOV corresponding to the left-to-right and top-to-bottom angular sweep across the virtual image of the display. Let's call these angles the physical FOV. To be accurate in our definition of the FOV, we shall consider that a point in object space contributes to the physical FOV if the chief ray defined as the ray passing through that point and the center of the exit pupil of the viewing system (here the pupil of the eye) is not obstructed. We are somewhat conservative therefore in our definition of the FOV since a point could be said to belong to the FOV as long as at least one ray reaches the image plane. This FOV is usually referred to in the optics literature as the total FOV.

An FOV is also specified in the display code. This computational FOV determines how far away the center of projection is from the screen rectangle in the perspective transformation that is used to project the 3D virtual world onto the 2D screen. There is both a horizontal and vertical FOV in this case, also. These are specified in many graphics systems by giving the aspect ratio, the ratio of horizontal to vertical FOV angles, and then giving the vertical angle.

Unfortunately, the computational FOV angles may not match the physical FOV angles. As with convergence, the software designer may not know the optical specifications and may be forced to measure the FOV empirically. To get it right, the position of the center of projection with respect to the screen should be exactly at the entrance pupil of the eye.

2.4 Failure to use off-center projection when required

If the display screen is perpendicular to the optical axis and off-center from the axis, then the eye is off-center with respect to the screen, and the computational center of projection should be off-center, too. This situation requires an off-center perspective projection, in which the left, right, top and bottom edges of the screen rectangle are specified independently of one another, rather than having left-right and top-bottom be symmetrical as usual. An off-center perspective projection transformation can be set up with the standard computer graphics hardware using the standard 4 x 4 homogeneous matrix representation of the transformation. Many people even in the computer graphics profession are unfamiliar with the off-center perspective projection -- it is never needed in normal computer graphics because, since the physical eye position is unknown, a convenient one directly in front of the screen may as well be assumed.

Another mistake is to simply rotate the whole scene left or right to produce images for the two eyes. In general, the computational model needs to include transformations that take into account the position and orientation of the screens with respect to the user's eyes. This is likely to be a combination of translation, rotation and off-center perspective projection, and very unlikey to be just a pure rotation [SAUN68].

2.5 Inter-pupillary distance ignored

Among male and female adults, there is a fairly wide variation in the distance between the eyes, called the inter-pupillary distance (IPD). The range is roughly 53 to 73 mm, with the average IPD being about 63 mm. Children have even smaller IPDs.

The variation in IPD imposes a requirement on the HMD optics and display hardware. Either the exit pupil of the optics must be large enough to be used by the widest-eyed and the narrowest-eyed people, or else there must be a mechanical adjustment such as is found on binoculars. Both of these approaches have been used. For example, the LEEP optics, used in many current HMDs, has a very large exit pupil. The CAE fiber-optic HMD used for flight simulation requires a mechanical adjustment for each new user.

The distance between the eyes is the baseline from which a person judges the distance to objects in the physical world. The convergence of the lines of sight from the eyes to a point of fixation can be measured as an angle. As Figure 2.1 shows, a narrow-eyed and a wide-eyed person will have different convergence angles when looking at an object at the same distance, say, half a meter away.

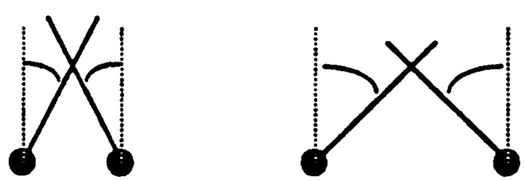

Figure 2.1 Eyes with narrow versus wide IPDs looking at an object from the same distance

These two people have different convergence angles yet both perceive the object to be half a meter away. Each person is calibrated to his or her own IPD.

With the mechanical IPD adjustment, the images get piped into the user's eyes and his or her physical IPD has no effect on the images seen. With no mechanical adjustment but a large exit pupil, if the virtual images are at optical infinity, then a lateral change in eye position has no effect on the angles at which an object appears -- in other words, in this case, too, the user gets the same one-size-fits-all images regardless of the physical distance between the eyes. If the virtual images are not at optical infinity, the situation is more complicated, and people of varying IPD are still not going to see images matched to their own IPDs.

The solution to this problem is to measure the IPD for each user, and have the IPD as a user-specific constant in the computational model. If this is done correctly, then each user can see a simulated object half a meter away with a convergence matched to his or her own IPD.

2.6 Optical distortion ignored

The display screens in an HMD are too close to the eyes to focus on directly, so an optical system is interposed between the eye and the screen. The main purpose of the optics is to provide an image to the user at a comfortable distance of accomodation and with as large a magnification as possible without altering the image. The eye, looking into the optics, sees a virtual image of the display screen. The virtual image is distant enough from the eye to focus on easily, and large enough to cover a large swath of the user's FOV. But the optics also distort the image non-linearly, causing lines that were straight on the display screen to appear as curved in the virtual image.

Optical aberrations are defects of the image. They may be described in terms of the amount by which a geometrically-traced ray misses a specified location in the image plane formed by the optical system. The displacement of the ray is referred to as the transverse ray aberration. Most often, the specified location for a ray in the image plane is that inferred from first-order laws of image formation [HECH74] [LONG73]. Rays which do propagate not only near the optical axis but also at shallow angles with respect to the optical axis are known as paraxial rays. Under the paraxial approximation, the formation of images is referred to as first-order, paraxial, or Gaussian optics. The most common aberrations are spherical aberration (SA), coma, astigmatism (AST), field curvature (FC), distortion, and chromatic aberrations. It should be noted that, while SA, coma, AST, FC, and chromatic aberrations all affect the sharpness of the image points being formed, distortion distinguishes itself from the others since it causes the image points to be displaced transversally in a non-linear fashion across the FOV but does not alter the sharpness of the image.

Transversal aberrations can be expressed mathematically as a polynomial expansion of third-order and higher-order terms in both the image height and the height of strike of the ray on a reference sphere centered on the ideal image point and passing through the exit pupil of the optical system. The ideal image point is often chosen to be the paraxial image point. The sum of the exponents of the aperture and field terms indicates the order of the aberration represented by that term. Depending on how open the optical system is and how large the angles of incidence of the rays on the different surfaces of the optical elements are, a system is best described by a third-order or a higher-order approximation. The complexity of the optics used is usually such that the sharpness of the images formed through the optical system is good enough for the display resolution available on the market today. The distortion of the images, however, is often disturbing if it has not been corrected for optically. The first-order polynomial is linear and describes an ideal magnifier with no distortion, and since there are no even terms appearing in the expansion, the third-order polynomial is the simplest model of distortion.

Non-linear field distortion causes straight lines on the screen to appear curved. This can be corrected for in the graphics system by predistorting the images written onto the screens. A straight line in the virtual image would be created by writing a curved line onto the screen, such that its curvature exactly balanced out the optical distortion. This would require that the inverse of the screen-to-virtual-image distortion function be stored in the graphics system, and that each pixel be remapped to a new location on the screen. This is computationally expensive.

Most current HMD systems just ignore the optical distortion, because they may not have access to the distortion function for the optics, and because of the performance penalty even if they did do the correction. This is not an unreasonable choice in the early stages of development of an HMD, because the system is usable even with the optical distortion.

However, one side-effect of non-linear field distortion is that there is no single correct value for the FOV. Non-linear distortion causes the magnification to vary across the FOV. If the computational FOV is set to match the physical FOV, then objects in the center of the field will appear to be the wrong size. But if the computational FOV is set to make small central objects appear to be the right size, the objects in the peripheral field will be positioned wrong. The only way to avoid this unpleasant choice is to predistort the image to correct the optical distortion.

If the optics are designed specially for an HMD, then specific types of aberrations can be minimized. But if optics designed for another purpose are used, then you take what you can get. The LEEP optics, used in many current HMD systems, were designed for stereoscopic photography, and purposely incorporate substantial field distortion and chromatic aberrations.

2.7 Transforming only polygon vertices in the presence of non-linearities

Computer graphics has traditionally gained much efficiency by representing graphics primitives such as lines and polygons as collections of points, running only the points through the transformation pipeline, and then drawing the lines or polygons between the transformed points. This works only if all the transformations in the pipeline are linear. It is tempting to run only the polygon vertices through the predistortion function and let the very efficient hardware fill in the polygons. But then only the vertices would be in the right place -- the polygon edges would still be curved by the optical distortion. Figure 2.2a shows a simple case of this. Edges of polygons that cross a large fraction of the screen would be most noticeably curved.

Another problem with this approach is that continuity would be lost. A vertex that touched an edge before predistortion and scan-conversion would not be guarranteed to do so in the final virtual image. Gaps and holes would open up everywhere. Figure 2.2b shows this.

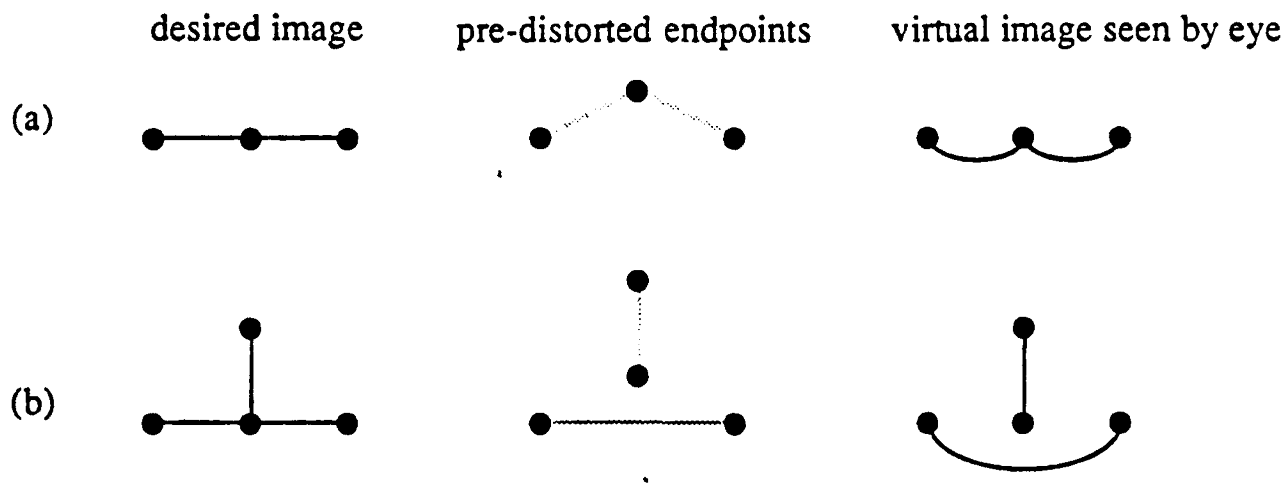

Figure 2.2 Problems with predistorting only vertices

3. OPTICS MODEL FOR A HEAD-MOUNTED DISPLAY

The purpose of the optics model is to specify the computation necessary to create orthostereoscopically correct images for an HMD and indicate the parameters of that system which need to be measured and incorporated into the model.

3.1 Single-eye optics model

To achieve orthostereoscopy, the non-linear optical distortion must be corrected by remapping all the pixels on the screen with a predistortion function. Linear graphics primitives such as lines and polygons are written into a Virtual Screen image buffer, and then all the pixels are shifted according to the pre-distortion function and written to the Screen image buffer for display. The predistortion function is the inverse of the field distortion function for the optics, so that the virtual image seen by the eye matches the image in the Virtual Screen buffer. A straight line in the Virtual Image buffer is predistorted into a curved line on the display screen, which is distorted by the optics into a line which is seen as straight.

Figure 3.1 shows the optics model for a single eye.

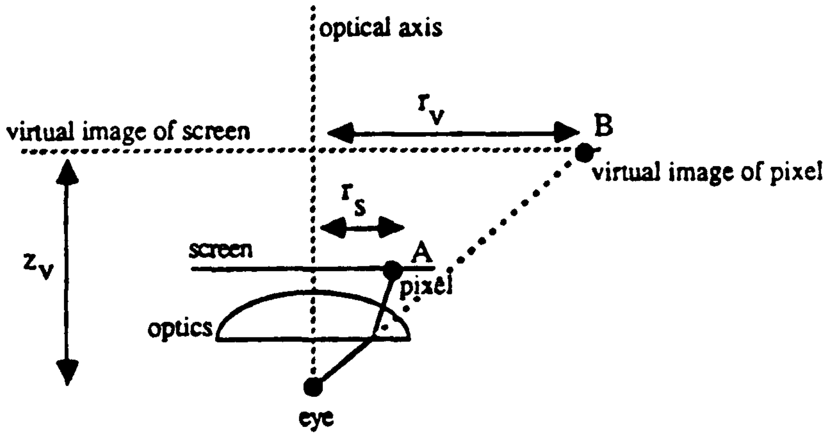

Figure 3.1 Single-eye optics model

The mathematical representation of the optical distortion will depend on the nature of the optical system. For optics such as LEEP, a third-order polynomial approximation is adequate. We want to relate the radial position r_s of a pixel on the screen to the radial position r_v of the virtual image of that pixel. These two quantities are measured in mm with respect to the optical axis. Dividing r_s and r_v by the object field width w_s and the image field width w_v, respectively, we get the normalized position of the pixel on the screen

$$r_{sn} = r_s / w_s$$

and the normalized position of the virtual image of the pixel

$$r_{vn} = r_v / w_v .$$

(The paraxial magnification of the system is thus $m_{sv} = w_v / w_s$.)
The distortion is modeled with a third-order polynomial approximation

$$r_{vn} = r_{sn} + k_{vs} r_{sn}^3$$

in which the coefficient k_{vs} describes the amount of distortion present. This can be rearranged algebraically to

$$r_{vn} = (1 + k_{vs} r_{sn}^2) r_{sn}$$

and then expanded to rectangular coordinates using

$$r_{sn}^2 = x_{sn}^2 + y_{sn}^2$$
$$r_{vn}^2 = x_{vn}^2 + y_{vn}^2$$

to give the position of the virtual image of the pixel

$$(x_{vn}, y_{vn}) = ((1 + k_{vs}(x_{sn}^2 + y_{sn}^2)) x_{sn}, (1 + k_{vs}(x_{sn}^2 + y_{sn}^2)) y_{sn})$$

from the position (x_{sn}, y_{sn}) of the pixel on the screen. These positions are measured in the screen plane and virtual image plane, respectively, with respect to the optical axis.

The distance from the eye to the virtual image plane z_v is a constant. By combining the above equations, we have a function that gives the three-dimensional position of the virtual image of a pixel on the screen (x_v, y_v, z_v) in terms of its screen coordinates (x_s, y_s) and some constants. We will ignore z_v from here on since it is constant.

The expressions for x_{vn} and y_{vn} can be thought of as single-valued functions of two variables. If we name the distortion function D, then

$$(x_{vn}, y_{vn}) = D(x_{sn}, y_{sn})$$

and what we need to predistort the image on the screen is the inverse D^{-1}. There are various ways this inverse function could be represented in the computer. An exact closed-form expression is not feasible, but a polynomial approximation of the inverse is possible.

$$r_{sn} = r_{vn} + k_{sv} r_{vn}^3$$
$$(x_{sn}, y_{sn}) = D^{-1}(x_{vn}, y_{vn})$$
$$(x_{sn}, y_{sn}) = \left((1 + k_{sv}(x_{vn}^2 + y_{vn}^2))\, x_{vn},\ (1 + k_{sv}(x_{vn}^2 + y_{vn}^2))\, y_{vn} \right)$$

Note that these two functions D and D^{-1} are each third-order polynomial approximations and are not exact inverses of one another. The coefficients k_{sv} and k_{vs} will be opposite in sign.

Another possibility for representing D^{-1} on the computer is a two-dimensional table lookup for each of the output variables x_{sn} and y_{sn}. Using this approach, limits on table size would probably make it necessary to interpolate between table entries.

Distortion causes non-uniform magnification across the field, which causes the brightness also to vary across the field. This could be compensated for on a pixel-by-pixel basis with a brightness correction function $B(x_s, y_s)$, but limitations of space prevent us from going into this here.

3.2 Stereoscopic optics model

Figure 3.2 shows the stereoscopic optics model. One pixel is illuminated on each screen (points A1 and A2) and a line-of-sight is drawn from each eye to the virtual image of its corresponding pixel (points B1 and B2). These two lines-of-sight intersect at the three-dimensional point perceived by the user (point C). The IPD is the baseline from which the user makes distance judgements based on the convergence angles to perceived points.

If the specifications for the optics are known, and the relative positions of the display screens, the optics, and the eyes are also known, then it is possible to accurately calculate several important parameters needed in the computational model. The horizontal and vertical FOVs can be calculated by starting from the known positions of the left, right, top and bottom edges of the screen with respect to the optics and tracing the rays through the optics system back to the eye.

The screen centers may be offset from the optical axes by a certain distance. The two optical axes may be rotated with respect to one another. The screen center offsets and axis divergence angles can be used to set the convergence properly with no need for subjective calibrations. From the known position and orientation of the virtual image of the screen relative to the eye, the Eye-To-Virtual-Image transformation can be calculated to be the correct mix of translation, rotation, and off-center perspective projection. Translation is needed because the eyes are in different positions in space; rotation is needed if the optical axes are not parallel; and the projection is off-center if the screens are off-center from the optical axes.

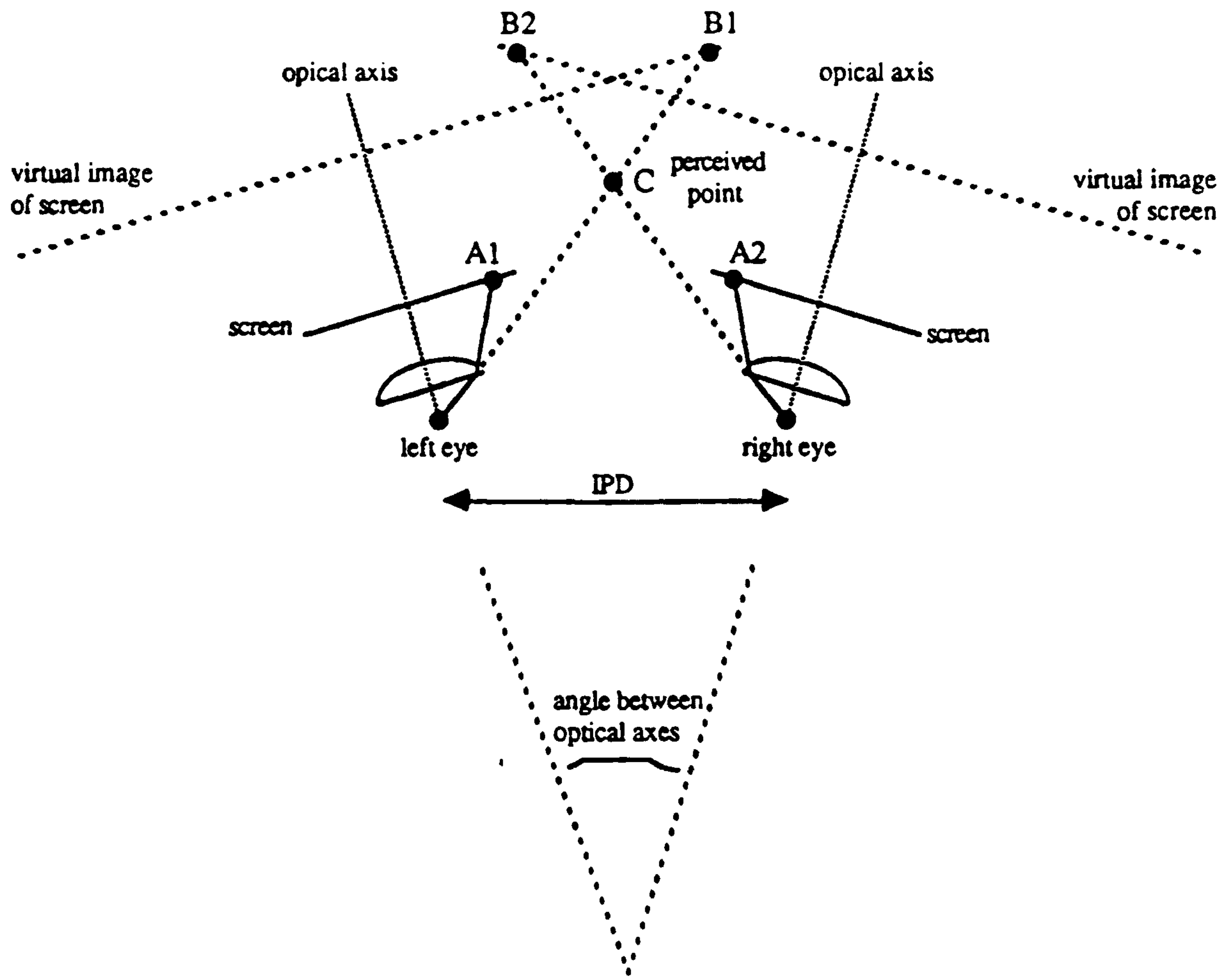

Figure 3.2 Stereoscopic optics model

Table 3.1 shows the sequence of starting from the specifications of the optics and displays, measuring certain parameters of the HMD, and then calculating other parameters needed by the computational model. The left and right halves of the optical system are assumed to be symmetrical. The calculations are done for the right side, so the left edge of the screen is on the inside beside the nose and the right edge is on the outside.

Except for the IPD which varies among users, every other parameter necessary to specify the Head-To-Eye, Eye-To-Virtual-Image and Virtual-Image-To-Screen transformations for the left and right eyes can be derived from the specifications of the optics and the relative positions of the eyes, optics and screens. Calculating these parameters is much more accurate than relying on subjective calibration procedures in which the parameters are adjusted until the image looks right. However, the subjective measurements provide a nice check against mistakes in the model, the measurements or the calculations.

We have defined a computational model for the graphics computation without reference to a particular HMD system. Having done that, we now turn to calculating the model parameters for a specific HMD, the VPL EyePhone.

Parameter	Symbol	Where It Comes From
screen resolution	$Res_H \times Res_V$	from display spec
angle between optical axes	ϕ_{axes}	from optics spec or measure
distance between optical axes (at front surface of optics)	d_{axes}	from optics spec or measure
eye relief	d_{er}	measure
maximum field-of-view	ϕ_{max}	calculate from d_{er} and optics spec
object plane distance	d_{ob}	measure
distance from eye to virtual image plane	z_v	calculate from d_{ob} and optics spec
transversal magnification	m_{vs}	calculate from d_{ob} and optics spec
coefficient of optical distortion	k_{vs}	from optics spec
maximum object field radius	w_s	from optics spec
maximum virtual image field radius	w_v	$w_v = m_{vs} w_s$
inter-pupillary distance of user	IPD	measure user
screen center offset from optical axis	(Cx_s, Cy_s)	measure
position of left screen edge	(Lx_s, Ly_s)	measure
position of right screen edge	(Rx_s, Ry_s)	measure
position of top screen edge	(Tx_s, Ty_s)	measure
position of bottom screen edge	(Bx_s, By_s)	measure
object height (of point on screen)	r_s	$r_s = (x_s + y_s)^{1/2}$
normalized object height	r_{sn}	$r_{sn} = r_s / w_s$
normalized image height (3rd order approximation of D)	r_{vn}	$r_{vn} = r_{sn} + k_{vs} r_{sn}^3$
image height (of point in virtual image)	r_v	$r_v = r_{vn} w_v$
angular position of point in virtual image	ϕ	$\phi = \tan^{-1}(r_v / z_v)$
angular position of left edge (inner edge of screen)	ϕ_L	from formula for ϕ
angular position of right edge (outer edge of screen)	ϕ_R	from formula for ϕ
angular position of top edge	ϕ_T	from formula for ϕ
angular position of bottom edge	ϕ_B	from formula for ϕ
single eye vertical field-of-view	FOV_v	$FOV_v = \phi_T + \phi_B$
single eye horizontal field-of-view	FOV_h	$FOV_h = \phi_L + \phi_R$
overlapped field-of-view	FOV_{ov}	$FOV_{ov} = 2\phi_L - \phi_{axes}$
binocular field-of-view	FOV_{bin}	$FOV_{bin} = 2\phi_R + \phi_{axes}$
translation part of viewing transformation	M_{trans}	$(\pm IPD / 2, \; 0, \; 0)$
rotation part of viewing transformation	M_{rot}	$(\pm \phi_{axes} / 2)$ around Y-axis
perspective projection	$M_{perspec}$	use FOV_v, FOV_h, and offset $(\pm Cx_s / (Rx_s - Lx_s), \; Cy_s / (Ty_s - By_s))$

Table 3.1 Calculating parameters in the optics model

4. CALCULATING THE MODEL PARAMETERS FOR THE VPL EYEPHONE

4.1 Description of EyePhone components

We have used the optical specifications of the LEEP optics and size and positioning of the LCD screens inside the EyePhone to calculate the model parameters for the VPL EyePhone, Model 1 [VPL89]. Model 2 of the EyePhone has identical optics, LCD displays, and positioning of the parts.

The LEEP optics [HOWL83] are wide-angle, stereoscopic viewing lenses. It consists of a three-lens system in front of each eye, encased in a molded plastic mount, with a cut-out for the nose. It was designed for a single transparency in the object plane, upon which are two side-by-side stereoscopic photographs of a scene, each one a square approximately 64 mm on a side. For an eye relief distance of 29.4 mm, the FOVs for each eye are approximately +45 to -45 degrees horizontally, and +45 to -45 degrees vertically. The distance from the center of the rearmost lens surface to the object plane is approximately 16 mm. The optical axes for the two eyes are parallel and are 64 mm apart. The two optical systems are bilaterally symmetrical with respect to each other, and each optical system is radially symmetrical around the optical axis, except for the cut-outs for the user's nose in the front lenses.

The construction of the EyePhone is rigid enough to keep the eyes, optics, and display screens in a fixed relationship to one another, so it is possible to make accurate measurements of these relative positions.

The EyePhone uses the LEEP optics with two color LCD display screens positioned in the LEEP object plane. The two displays are positioned symmetrically with respect to the left and right eye optical axes. Figure 4.1 is a diagram of the positions of the LCD screens in the object plane.

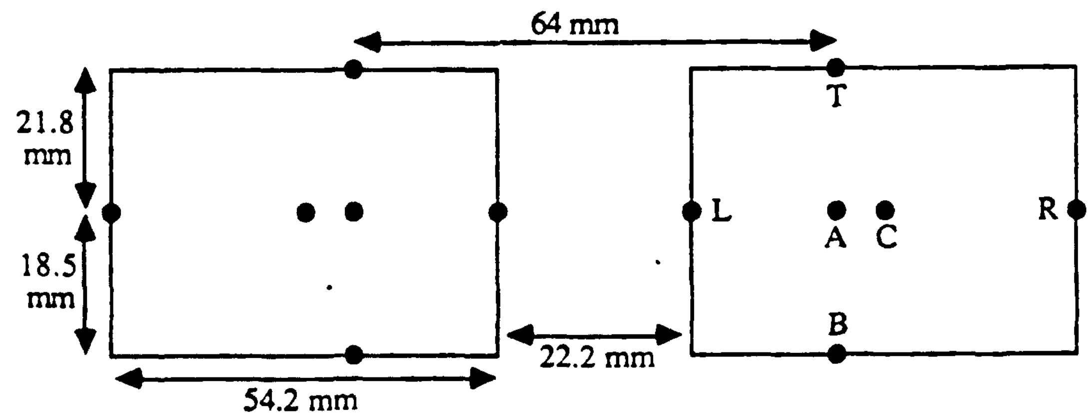

Figure 4.1 Position of the EyePhone's LCD displays in the LEEP object plane

Six important points are labeled in Figure 4.1. Only one side need be analyzed since the LCDs and optics are symmetrical.

4.2 Calculation of EyePhone field-of-view

We use two different methods to compute the FOVs for the EyePhone. First, we use the optics model with the parameters specific to the EyePhone to calculate the angles at which the points in Figure 4.1 are seen. Then, for comparison and validation, we compute these same angles by tracing rays through the LEEP optics. We do the ray-tracing with Code V, a commercial optical analysis software package [ORA90]. To do the ray-tracing, we use the detailed optical specifications of the LEEP optics, including the position, curvature, and index of refraction for each lens surface in the optical path.

First, we use the optics model. The dimensions given in Figure 4.1 are sufficent to determine the coordinates of the points A, L, R, T and B in the LEEP object plane (with respect to the optical axis A). From these coordinates, the value of r_s for each point can be calculated. We feed the known positions of the edges of the LCD screens into the optics model to predict the positions of the screen edges in the virtual image, and from these positions calculate the FOVs for the EyePhone. We shall see that the chief rays corresponding to some of the object points on the LCDs are obstructed for d_{er} = 29.4 mm. In this case the FOV covers only part of the LCDs. The unseen part of the LCD is said to be vignetted. The model parameters for the LEEP optics as used in the EyePhone are

$$d_{er} = 29.4 \text{ mm}$$
$$w_s = 28.1 \text{ mm}$$
$$w_v = 271.5 \text{ mm}$$
$$z_v = 398.2 \text{ mm}$$
$$k_{vs} = 0.32$$

and the equations from the optics model

$$r_{sn} = r_s / w_s$$
$$r_{vn} = r_{sn} + k_{vs} r_{sn}^3$$
$$r_v = r_{vn} w_v$$
$$\varnothing = \tan^{-1}(r_v / z_v)$$

can be used to calculate the angle ø for each point, as shown in Table 4.1.

The second method of computing the angles is ray-tracing. Figure 4.2 shows a horizontal cross-section of the right-eye LEEP optical system, with the rays from the points L, A and R traced back to the eye. The ray-tracing was done for each of the points A, L,R, T and B in Figure 4.1.

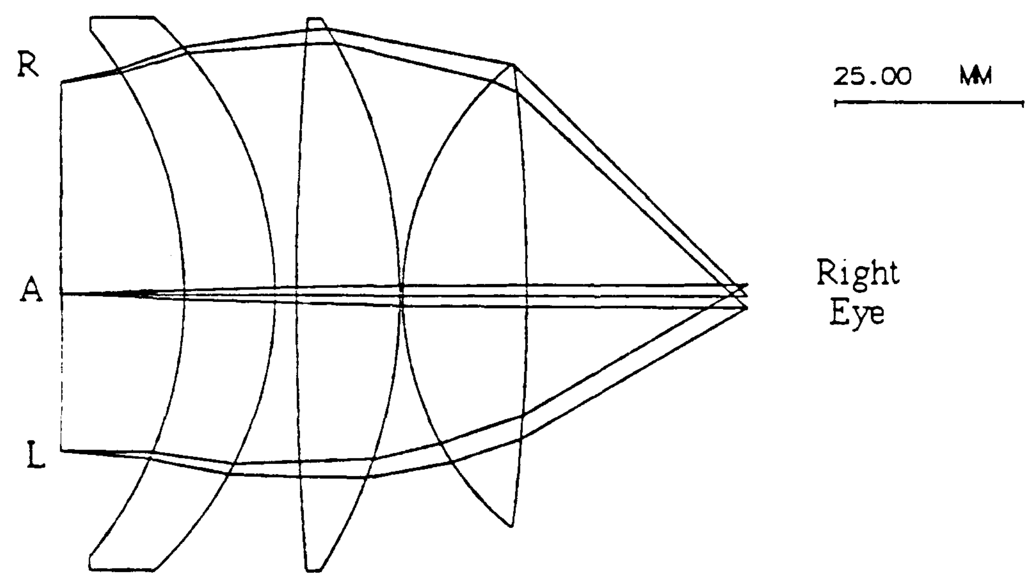

Figure 4.2 Tracing rays from the EyePhone's LCD screen through the LEEP optics

Table 4.1 shows the angular positions of the virtual images of the edges of the LCD screen as predicted by the optics model, and as predicted by ray-tracing. The comparison of the angles calculated by the two methods shows that the third-order approximation used in the optics model is adequate for the LEEP optics.

Point	x_s (mm)	y_s (mm)	r_s (mm)	r_{sn}	r_{vn} 3rd order	r_v (mm)	angle ø from model (degrees)	angle ø from ray tracing (degrees)
A optical axis	0	0	0	0	0	0	0	0
R right edge of screen	33.3	0	28.1 (33.3 vignetted)	1.000	1.32	358.5	42.0	45.0
L left edge of screen	-20.9	0	20.9	0.744	0.876	237.9	30.9	30.3
T top edge of screen	0	21.8	21.8	0.776	0.926	251.5	32.3	31.8
B bottom edge of screen	0	-18.5	18.5	0.658	0.749	203.4	27.1	26.6
C center of screen	6.2	1.65						

Table 4.1 EyePhone FOV calculation, assuming d_{er} = 29.4 mm

From these calculations, we can see that, for a single eye, the horizontal FOV is 75.3 degrees (45.0 + 30.3) and the vertical FOV is 58.4 degrees (31.8 + 26.6). These are the physical FOVs for the EyePhone -- the physical angles at which the virtual images of the edges of the LCD screen are seen by the eye. To make the graphics calculation match the physical FOV, these angles must be incorporated into the calculation. Here at UNC, we are using the Pixel-Planes graphics engine [EYLE88] to generate images for the EyePhone. Like many graphics systems, the Pixel-Planes graphics software accepts an angle for the vertical FOV, here 58.4 degrees, and an aspect ratio, here 1.289 .

The graphics calculation must also take into account the fact that in the EyePhone, the center of the LCD screen (point C) is off the optical axis (point A). How this off-center perspective projection is specified to the graphics software varies somewhat among graphics systems. For the Pixel-Planes graphics software, off-center projection is specified as a horizontal and vertical offset in pixels from the screen center. For the horizontal offset, 512 pixels across a 54.2 mm screen gives 0.106 mm/pixel, so the 6.2 mm horizontal offset is 58.5 pixels. For the vertical offset, 512 pixels across a 40.3 mm screen height gives 0.079 mm/pixel, so the 1.65 mm vertical offset is 21.0 pixels.

These calculated values for FOV and screen-center offset together with the IPD of the user is enough to specify precisely the perspective projection for the two eyes. Using these projections, the convergence and FOV are guarranteed to be correct, with no need for adjustment or calibration by the user. Our experience with the EyePhone is that using the calculated parameters gives very solid stereo perception for all users who can see in stereo, with no need for tweaking these parameters. Getting this all to work depends on getting several subordinate things right -- the specifications for the optics must be correct, the analysis of the optics must be correct, the measurements of the LCD positions and dimensions must be correct, and the graphics software for setting up the projections must be correct.

Before these calculations were done, the FOV and off-center parameters had been tweaked, through a long process of trial and error, trying to get the image to look right. The most effective test was to try to get the image of a 5.75 centimeter red sphere to be the right size and stay on top of a physical 5.75 centimeter red 3-ball (from the game of pool) with a Polhemus sensor inside as the 3-ball was moved around. This 3-ball, which has two pushbuttons on it, is one of our manual input devices. The FOVs for a single eye which were finally arrived at through these subjective tests were a horizontal FOV of 80 degrees and a vertical FOV of 60 degrees. Since the EyePhone does not allow the user to see through to the real world, the subjective tests had to be done by repeatedly raising the EyePhone to see how the position of the physical 3-ball compared with the remembered position of the simulated 3-ball. We estimated the accuracy to be about 2 degrees. For a see-through HMD, superimposing a 3D stereoscopic image onto a physical object of the same known shape is an extremely accurate test, because the eye can simultaneously compare the two scenes and detect tiny discrepancies. The acuity of the human eye is considered to be roughly one minute of arc.

Figure 4.3 shows that the horizontal FOVs for the two eyes in the EyePhone partially overlap. The region of overlap (60.6 degrees) is wide enough for strong stereoscopic perception, and the binocular FOV (90.0 degrees) is wide enough to provide a feeling of immersion within the scene. We believe, however, that a wider FOV will make the feeling of immersion stronger.

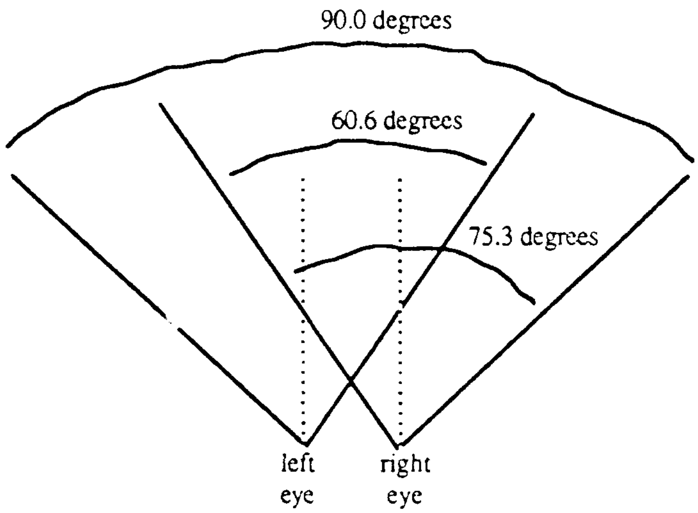

Figure 4.3 Binocular, overlapped and single-eye FOVs for the EyePhone

4.3 Correction of EyePhone optical distortion

To correct the image on the screen from optical distortion, the screen image must be predistorted as specified by the function D^{-1}. This would radially shift all the pixels in the image by some amount. This has not yet been implemented on the UNC HMD system, but we plan to use a pair of two-dimensional tables $T_X[x_v, y_v]$ and $T_Y[x_v, y_v]$ in this implementation.

We have a 512 x 512 pixel screen to cover, but a table size of 512 x 512 is impractical. We expect to use a reduced table size, such as 64 x 64, and interpolate bilinearly between table entries. The table values will be computed off-line using the formula for D^{-1}.

The optics map object points of height r_s to image points of height r_v. This mapping is monotonic and so its inverse exists. We approximate the mapping with a third-degree polynomial D and approximate its inverse with another third-degree polynomial D^{-1}

$$r_{sn} = r_{vn} + k_{sv} r_{vn}^3$$

in which, for the LEEP optics, k_{sv} is -0.18 .

Figure 4.4 shows the graph of the normalized virtual image position r_{vn} versus the normalized screen position r_{sn} for the LEEP optics. It shows that the third-degree polynomial approximation of the distortion function D is quite close to the more accurate graph of D calculated by ray-tracing through the optics. The second graph compares the third-degree polynomial approximation of the inverse D^{-1} with the values calculated by ray-tracing.

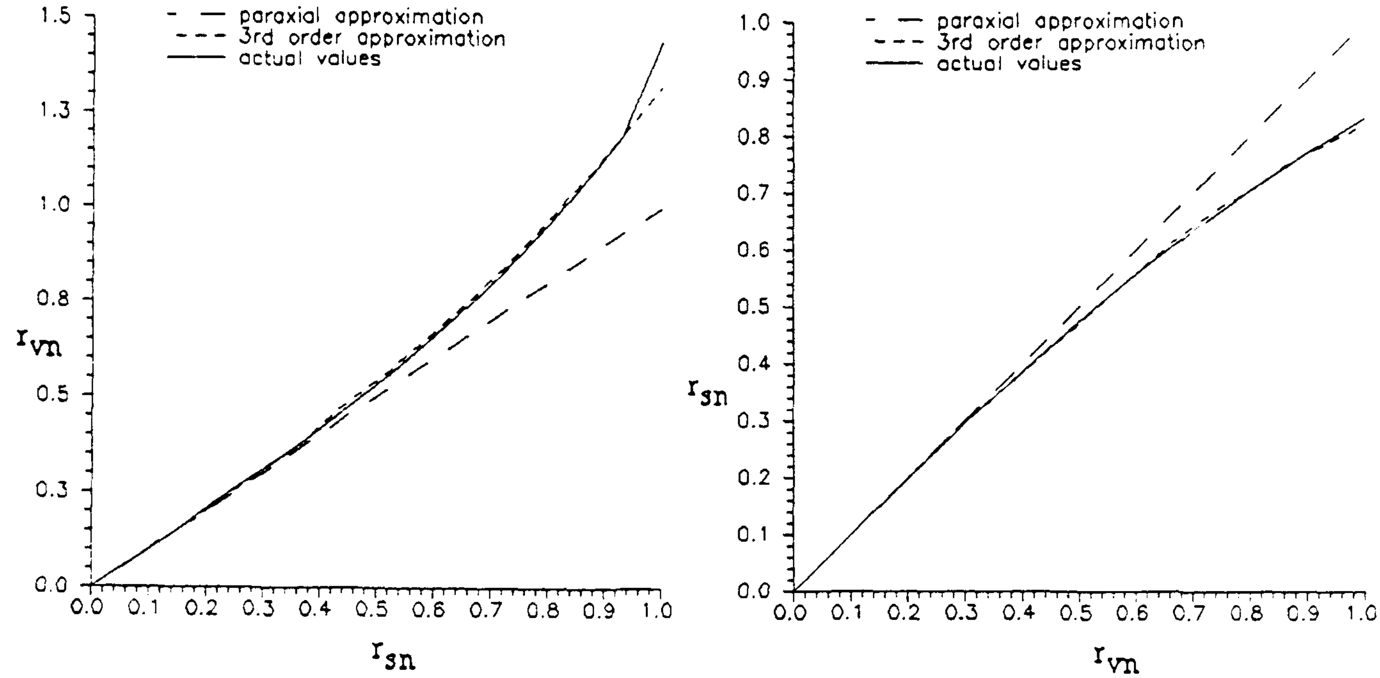

Figure 4.4 Graph of r_{vn} vs. r_{sn} for D (k_{vs} = 0.32) and ray-tracing for the LEEP optics;
graph of r_{sn} vs. r_{vn} for D^{-1} (k_{sv} = -0.18) and ray-tracing

Figure 4.5 shows a grid which has been predistorted by the function D^{-1}. When viewed through the LEEP optics, the lines of the grid appear straight.

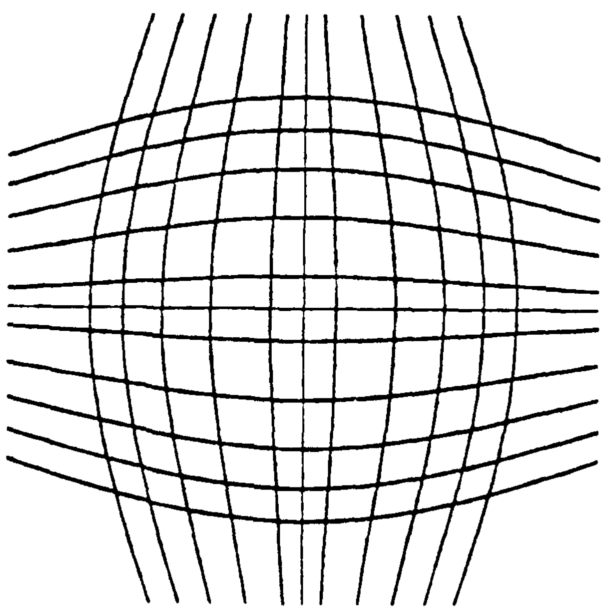

Figure 4.5 Grid predistorted for the LEEP optics

Although the predistortion function D^{-1}, also known as the Virtual-Image-To-Screen transformation, has not yet been implemented on the UNC HMD system, we have looked through the LEEP optics at the predistorted grid printed on paper to verify that the grid lines do appear straight. The predistorted grid in Figure 4.5 is full-sized (79 mm square) and may be viewed through the LEEP optics. The lines will appear straight when the grid object is in close contact with the lens. Note that in the EyePhone, the LCD screen is separated by a gap of a few millimeters from the LEEP lens.

4.4 Other EyePhone and LEEP parameters

The graph of Figure 4.6 shows how the LEEP optics FOV varies with the eye relief distance d_{er}, the distance between the eye and the nearest lens surface. With the nominal eye relief of 29.4 mm, the FOV for the LEEP optics is +45 to -45 degrees. (The EyePhone's FOV is less than this because the EyePhone's LCD screen does not fill the LEEP object field.) If the eye was able to get closer, the FOV would increase, but an eye relief of 25 to 30 mm is necessary to allow people with spectacles to use the system.

When moving the eye closer to the lens, two factors contribute to an increase of the FOV : first, each point on the virtual image is seen over a larger angle, and second, more of the LCD screens can be perceived. Especially for the EyePhone, moving the eye closer to the lens does increase the FOV because some of the LCD display screen is vignetted for a pupil distance of 29.4 mm.

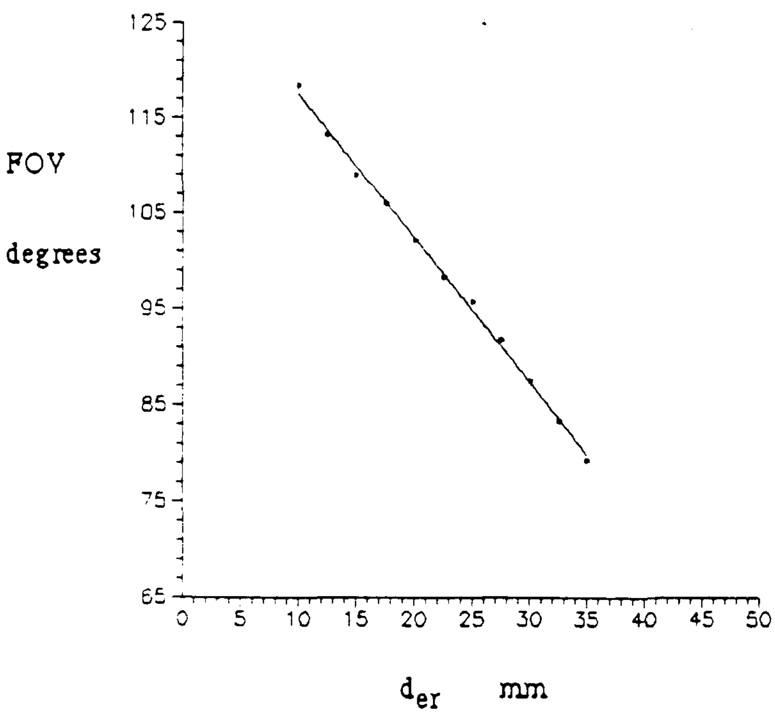

Figure 4.6 Graph of FOV vs. d_{er}

The virtual image of the screen is formed at some distance z_v from the eye, and the eye must accomodate to this distance. The graph of Figure 4.7 shows that z_v is very sensitive to changes in the distance d_{ob} of the object plane to the the LEEP optics. The EyePhone positions the LCD screen at d_{ob} = 16.4 mm from the nearest lens surface (measured along the optical axis). This value of d_{ob} results in an image distance of z_v = 398.2 mm and a magnification of m_{vs} = 9.66 . As the object approaches the object focal point of the lens (d_{ob} = 20.7 mm) the image distance goes to infinity. However, such a positioning seems undesirable because of the conflict between convergence and accomodation.

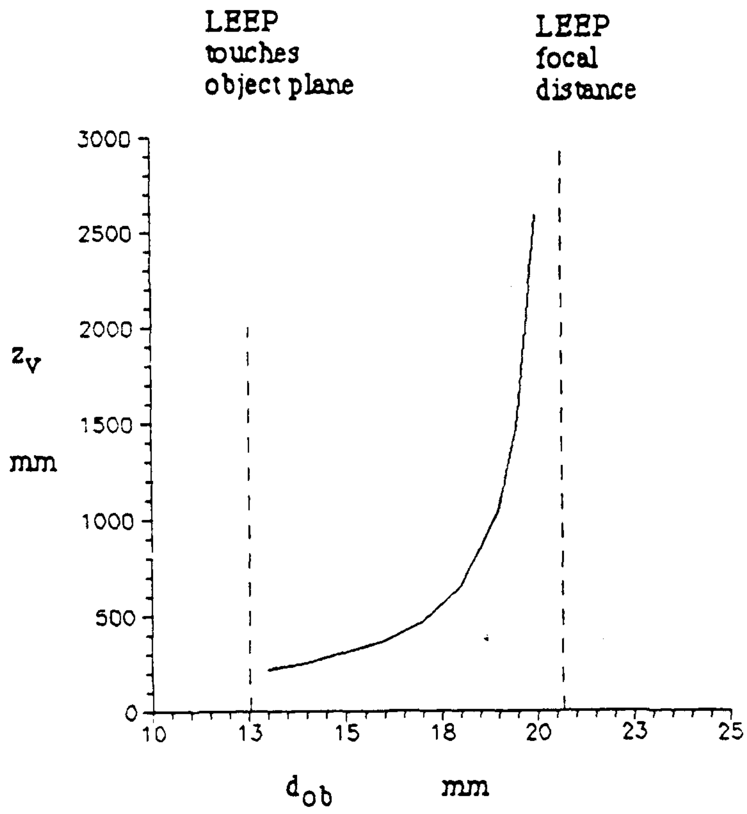

Figure 4.7 Graph of z_v vs. d_{ob}

The discussion of the resolution of the LCD display screens in the EyePhone is complicated by two competing ways of specifying resolution in a color display -- by color pixels or by the RGB component cells of the pixels. The LCD used in the EyePhone has a monochrome array of 320 x 240 individually controllable light-producing cells, with red, green and blue filters overlaid to divide the cells into three equal-sized groups. A triad of one red, one green and one blue cell makes up a color pixel. There are 76,800 cells (320 x 240), and therefore 25,600 color pixels. The resolution of the EyePhone in terms of color pixels is thus approximately 184.7 x 138.6.

Table 4.2 lists all the parameters for the EyePhone with respect to the model developed in this paper. All the angles (ø and FOV parameters) were calculated from the other parameters as described above.

The LEEP optics are used in other HMDs besides the EyePhone, and a new model of EyePhone with a different LCD display is also being developed. Some of the parameters in Table 4.2 depend only on the LEEP optics ($ø_{axes}$, d_{axes}). Other parameters with specific values in the table describe the specific configuration in the EyePhone (Models 1 and 2) of the LCD screen's size and position in front of the LEEP optics.

The rubber diving mask of the EyePhone holds the face and eyes in a fairly constant position with respect to the LEEP optics, although some individuals have deeper-set eyes than others. Decreasing d_{er} by moving the eyes closer to the LEEP lens would cause more of the object field to be seen and the value of w_s which was chosen to be the highest value of the field would then have to be increased accordingly. If the same nominal eye relief used in this paper (29.4 mm) is assumed, then the distortion model can be applied to an HMD using the LEEP optics and a different screen using the coefficients calculated (k_{vs} = 0.32, k_{sv} = -0.18).

Parameter	Symbol	Value for EyePhone
maximum object field radius	w_s	28.1 mm
maximum virtual image field radius	w_v	271.55mm
transversal magnification	m_{sv}	9.66
eye relief	d_{er}	29.4 mm (nominal)
object plane distance (LCD screen to LEEP lens surface)	d_{ob}	16.4 mm
distance from eye to virtual image plane	z_v	398.2 mm
coefficient of optical distortion for D	k_{vs}	0.32
coefficient of optical distortion for D^{-1}	k_{sv}	-0.18
angle between optical axes	$\varnothing_{axes}$	0 degrees
distance between optical axes (at front surface of optics)	d_{axes}	64 mm
screen center offset from optical axis	(Cx_s, Cy_s)	(6.4 mm, 1.6 mm)
screen resolution	$Res_H \times Res_V$	185 x 139 pixels (color triads)
inter-pupillary distance of user	IPD	varies across users
angular position of virtual image of right edge of LCD	$\varnothing_R$	45.0 degrees
angular position of virtual image of left edge of LCD	$\varnothing_L$	30.3 degrees
angular position of virtual image of top edge of LCD	$\varnothing_T$	31.8 degrees
angular position of virtual image of bottom edge of LCD	$\varnothing_B$	26.6 degrees
single eye vertical field-of-view	FOV_v	58.4 degrees
single eye horizontal field-of-view	FOV_h	75.3 degrees
overlapped field-of-view	FOV_{ov}	60.6 degrees
binocular field-of-view	FOV_{bin}	90.0 degrees

Table 4.2 Parameters for the VPL EyePhone, models 1 and 2

To determine the model parameters for other HMDs that use the LEEP optics, the distance d_{ob} from the LEEP optics to the display screen must first be known. This will determine the distance to the virtual image z_v and the magnification m_{sv}. An eye relief distance d_{er} must also be measured or assumed. This will determine the radius of the object and image fields w_s and w_v. The object and image fields describe what could be seen if the object field was completely filled, regardless of how completely the display screen does fill the object field. To compute the FOVs for the virtual image of the display screen, the positions of the edges of the display screen (points L, R, T and B) must be measured in the object plane with the optical axis as the origin. Cranking these measurements through the model will give the angular positions of the edge points ($\varnothing_R$, $\varnothing_L$, $\varnothing_T$, $\varnothing_B$), and thus the horizontal and vertical FOVs (FOV_h, FOV_v) for a single eye. To find the binocular field-of-view FOV_{bin} for both eyes, the angles between the optical axes $\varnothing_{axes}$ must be taken into account. The position of the display screen's center with respect to the optical axis must also be measured in order to properly set up the perspective projection.

5. SUMMARY AND CONCLUSIONS

The optics model presented in this paper, if implemented correctly, will generate undistorted orthostereoscopic images for the user's two eyes.

To calculate the display parameters needed by the model for the particular HMD being used, it is necessary to know or measure the specifications of the optics, and the relative positions of the eyes, optics, display screens, and head position sensor. The construction of the HMD must be rigid enough that these values will not vary from day to day. If these parameters for the HMD are known, then several important derived parameters can be calculated -- the FOVs, the screen-center offset for the perspective projection, the angle between the optical axes, and the coefficients for the optical field distortion function. Calculating the values of these parameters is much more accurate than attempting to measure them subjectively with users.

6. ACKNOWLEDGEMENTS

We would like to thank many people for their contributions to this work. Eric Howlett, designer of the LEEP optics, kindly provided the optical specifications. Mike Teitel of VPL provided needed measurements of the EyePhone. Jack Goldfeather helped with the mathematical representation of the distortion.

Fred Brooks, Henry Fuchs and Steve Pizer helped to initiate and guide the HMD project at UNC, and various parts of the UNC HMD system were built by each of the team members: Ron Azuma, Bill Brown, Jim Chung, Drew Davidson, Erik Erikson, Rich Holloway, Jack Kite and Mark Ward. Fay Ward and Linda Houseman helped to hold it all together and Vern Chi, David Harrison and John Hughes provided essential technical know-how. The Pixel-Planes team provided us with the high-powered graphics engine needed to drive the HMD.

This work builds upon earlier work by one of us (Robinett) at NASA Ames Research Center, and we would like to acknowledge the contributions of Scott Fisher, Jim Humphries, Doug Kerr and Mike McGreevy.

This research was supported in part by the Defense Advanced Project Research Agency, contract # DAEA18-90-C-0044 and also by the Office of Naval Research, contract #N00014-86-K-0680.

7. REFERENCES

[CAE86] CAE Electronics, Ltd., Quebec, *Introducing the visual display system that you wear*, 4 pp. (1986)

[CALL83] Mark A. Callahan, A 3-D display headset for personalized computing, M.S. thesis, Dept. of Architecture, Massachusetts Institute of Technology, 110 pp. (1983)

[CHUN89] James C. Chung, Mark R. Harris, Fredrick P. Brooks, Henry Fuchs, Micheal T. Kelley, John Hughes, Ming Ouh-Young, Clement Cheung, Richard L. Holloway, Micheal Pique, "Exploring Virtual Worlds with Head-Mounted Displays," *Non-Holographic True 3-Dimensional Display Technologies*, SPIE Proc. Vol. 1083 (1989)

[CLAR76] James H. Clark, , "Designing surfaces in 3-D," *Communications of the ACM*, 19(8), 454-460 (1976)

[EYLE88] John Eyles, John Austin, Henry Fuchs, Trey Greer, John 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)

[FISH86] Scott S. Fisher, Micheal McGreevy, James Humphries, Warren Robinett, "Virtual Environment Display System," *Proc. 1986 Workshop on Interactive 3D Graphics*, 77-87, (1986)

[GLIN86] C. V. Glines, "Brain buckets," *Air Force Magazine*, 69(8), 86-90 (1986)

[HECH74] E.Hecht and A. Zajac, Optics, Addison-Wesley Publishing Company (1974).

[HEND89] Breck W. Henderson, "Simulators Play Key Role In LHX Contractor Selection," Aviation Week magazine, Nov. 27, 1989

[HOLL87] Richard L. Holloway, "Head-Mounted Display Technical Report," Technical Report #TR87-015, Dept. of Computer Science, University of North Carolina at Chapel Hill

[HOWL83] Eric Howlett, , "Wide Angle Color Photography Method and System," U.S. Patent Number 4,406,532, Sept. 27, 1983

[KOCI86] Dean Kocian, paper on custom optics in HMDs at Wright-Patterson AFB

[LEEP90] LEEP Systems / Pop-Optix Labs, Waltham, Massachusetts, *Cyberface II Applications Note*, Aug. 1, 1990

[LONG73] R.S. Longhurst, Geometrical and physical optics, Longman, New york (1973).

[ORA90] Optical Research Associates, Code V software for optical systems analysis

[SAUN68] Bernard G.Saunders, , "Stereoscopic drawing by computer -- Is It Orthoscopic?", Applied Optics, Vol. 7, No. 8, pp1499-1504 (1968)

[SUTH65] Ivan E. Sutherland, , "The ultimate display," 1968 *Proceedings of the IFIPS Congress*, **2**, 506-508 (1965)

[SUTH68] Ivan E. Sutherland, , "A head-mounted three-dimensional display," 1968 *Fall Joint Computer Conference, AFIPS Conference Proceedings*, 33, 757-764 (1968)

[TIME89] London Times newspaper, "Mini screen is a real eye opener," May 28, 1989

[VICK74] D. L.Vickers, , *Sorceror's Apprentice: head mounted display and wand*, Ph.D. dissertation, Dept. of Computer Science, Univ. of Utah, Salt Lake City (1974)

[VPL89] VPL Research, Redwood City, California, "VPL EyePhone Operations Manual," p. B-4, June 27, 1989

THE UNIVERSITY OF NORTH CAROLINA

AT

CHAPEL HILL

Department of Computer Science

Campus Box 3175 - Sitterson Hall
Chapel Hill, NC 27599-3175

January 29, 1991

Steve E. Tice
Virtual Reality Applications Gallery Chair
SIGGRAPH '91
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

Dear Steve:

Enclosed are seven proposals from UNC for the Virtual Reality Applications Gallery. Since all of these applications from UNC require the same hardware, this is described in a separate, enclosed two-page document.

Thank you for the extension on submitting the videotape; we plan to get it to you by Friday. It has a segment for each of the seven proposals and then a segment on the hardware.

We hope you will accept several or all of these proposals. Feel free to contact any of us by phone (Brooks: 919/ 962-1931; Fuchs: 919/962-1911; Robinett: 919/962-1798).

Sincerely,

Frederick P. Brooks, Jr.
Kenan Professor and
Principal Investigator, HMD Project

Henry Fuchs
Federico Gil Professor and
Principal Investigator, HMD Project

Warren Robinett
HMD Project Manager

3-D X-Ray Synthetic Fluoroscopy With a Head-Mounted Display

Timothy Cullip
Department of Radiation Oncology and
Department of Computer Science
University of North Carolina
Chapel Hill, NC 27599-3175, USA
(919) 966-1101
(919) 962-1799 fax

System Category: Player Systems

System Application: Radiation Oncology, Radiology

Abstract

Computer-assisted tomography provides volume density data of objects such as the human body. Powerful image processing engines such as the UNC-built Pixel-Planes 5 enable the real-time computation of arbitrary projections that look like ordinary X-rays. A head-mounted, head-tracked display enables intuitive navigation from view to view, or synthetic fluoroscopy.

Wearing a head-mounted display, the visitor studies the relationships between tumor, anatomy, and planned treatment dose by viewing computed radiographs calculated on Pixel-Planes 5. Real-time computation of new radiographs corresponding to new viewing angles, coupled with stereoscopic images, yields three-dimensional information unavailable from static two-dimensional radiographs. In this exhibit the visitor moves his/her head to fully study the dose distribution produced by a prospective treatment plan in relation to the tumor within the patient's anatomy.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

The user views the stereo computed radiographs through the head-mounted display. By walking around the virtual patient, the user can view the patient from different angles, seeing through im by dynamic computed fluoroscopy.

Virtual Reality Applications Gallery Submission Form

Contact Person System Developer Name(s) Timothy Cullip
Organization University of North Carolina, Computer Science Department & Radiation Oncology Dept
Street Address
City Chapel Hill State NC
Zip/Postal Code 27599-3175 Country USA
Daytime Telephone (919) 966-1101 (Rad.Onc.Dept.) Fax (919) 962-1799 (Comp. Science Dept.)
System Title 3-D X-ray Synthetic Fluoroscopy With a Head-Mounted Display
System Category Player Systems
System Application(s): Radiology, Radiation Oncology

Abstract / System Description # (200 words suitable for the conference program)
See attached.

System Contents
See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, ~~regardless of how many are accepted.~~

Floorplan attached? [X] yes [ ] no Square footage required
Electrical power required See attached.
Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.
Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

- [X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.
- [X] My videotape-document(s) to be considered for broadcast.
- [X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.
- [X] Excerpts from the video-document or photos of my system to be used for the above purposes.
- [ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.

Signature Tim Cullip (...) Date 1/29/91

Please send completed submittals to:
Steve E. Tice
SIGGRAPH '91 Virtual Reality Applications Gallery Chair
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

Radiation Therapy Treatment Planning With a Head-Mounted Display

James C. Chung
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1889
(919) 962-1799 fax

System Category: Authoring Systems, Player Systems

System Application: Radiation Oncology

Abstract
To design an effective radiation therapy treatment plans, radiotherapists must possess a firm understanding of the three-dimensional arrangement of the tumor within the surrounding healthy tissue. This is not always possible with the static two-dimensional x-ray films used in conventional treatment planning. This exhibit samples ongoing research at UNC-CH into the possible advantages gained by using real-time, three-dimensional interaction to study patient anatomy and its spatial relation to tumor and dose distribution.

Wearing a head-mounted display, the visitor is presented with a polygonal surface model of a cancer patient's anatomy. Using a 6D mouse, the visitor grabs a treatment beam in its storage "rack", and drags it and positions it to pass through the tumor and as little surrounding healthy tissue as possible. This process can be repeated with several beams to ensure complete tumor coverage while minimizing the dosage received by healthy tissue.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

The user controls a cursor in the virtual world with a 6D mouse. The virtual treatment beams can be grabbed by placing the cursor into the desired beam and depressing the mouse button. As long as the button is depressed, the beam can be dragged around the virtual world. Releasing the mouse button releases the beam from mouse control and it will remain in its position at time of button release. The user is free to walk about and around the anatomical model within the constraints of the tracking system.

Virtual Reality Applications Gallery Submission Form

Contact Person System Developer Name(s)	James C. Chung
Organization	University of North Carolina, Computer Science Department
Street Address	
City	Chapel Hill
State	NC
Zip/Postal Code	27599-3175
Country	USA
Daytime Telephone	(919) 962-1700
Fax	(919) 962-1799
System Title	Radiation Therapy Treatment Planning With a Head-Mounted Display
System Category	Authoring Systems, Player Systems
System Application(s):	Radiation Oncology

Abstract / System Description # (200 words suitable for the conference program)

See attached.

System Contents

See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

Floorplan attached? [X] yes [ ] no Square footage required

Electrical power required See attached.

Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.

Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

- [X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.
- [X] My videotape-document(s) to be considered for broadcast.
- [X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.
- [X] Excerpts from the video-document or photos of my system to be used for the above purposes.
- [ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.

Signature James Chung (by H Fuchs) Date 1/29/91

Please send completed submittals to:
Steve E. Tice
SIGGRAPH '91 Virtual Reality Applications Gallery Chair
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

An Interactive Building Walkthrough System

Frederick P. Brooks, Jr.
John Alspaugh
Randall Brown
Curtis Hill
Amitabh Varshney
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1932
(919) 962-1799 fax

System Category: Virtual Worlds Navigation and Interaction (Single User)

System Application: Environment exploration (Architecture)

Abstract

Walkthrough is a system to help architects and their clients to explore a building design prior to its construction. Using a treadmill and a head-mounted display, the user can navigate through a three-dimensional virtual building and evaluate the design. This interface is more appropriate to the task than most others, for it both immerses the user in the virtual environment and allows him to explore it in a natural manner: by walking.

The building model is created using AutoCAD. Realism is enhanced by calculating a radiosity solution for the lighting model. This has an added feature of independent, interactively-switched lights. Thus, not only can the physical dimensions be tried out during design, but also such things as lighting, color combinations and decor. Due to the computationally intensive nature of the radiosity solution, modeling changes cannot be made on-line.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

The user wears a head-mounted display and walks through the virtual building using a treadmill. Translation speed through the model is governed by how fast the user walks and the length of user's stride, just as in the real world. Head movement and translation are independent. The user can orient his or her head to look in any direction, while walking and steering with a set of handlebars on the treadmill.

The operator starts the walkthrough program and then puts the head-mounted display on the user's head. The operator controls the light settings in the model.

Possible enhancements

We expect to add an optical tracker which would permit the user to walk around in a 20 x 20 foot virtual world. Another possible enhancement is the addition of procedural textures in the model.

Virtual Reality Applications Gallery Submission Form

Contact Person System Developer Name(s) Frederick P. Brooks, Jr., John Alspaugh, Randall Brown, Curtis Hill, Amitabh Varshn
Organization University of North Carolina, Computer Science Department
Street Address
City Chapel Hill State NC
Zip/Postal Code 27599-3175 Country USA
Daytime Telephone (919) 962-1700 Fax (919) 962-1799
System Title An Interactive Building Walkthrough System
System Category Virtual Worlds Navigation and Interaction (Single User)
System Application(s): Environment Exploration (Architecture)

Abstract / System Description # (200 words suitable for the conference program)
See attached.

System Contents
See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, ~~regardless of how many are accepted.~~

Floorplan attached? [X] yes [ ] no Square footage required
Electrical power required See attached.
Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.
Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

- [X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.
- [X] My videotape-document(s) to be considered for broadcast.
- [X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.
- [X] Excerpts from the video-document or photos of my system to be used for the above purposes.
- [ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.
Signature ________ Date 1/29/91

Please send completed submittals to:
Steve E. Tice
SIGGRAPH '91 Virtual Reality Applications Gallery Chair
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

Flying Through a Protein Molecule

Warren Robinett
Richard Holloway
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1700
(919) 962-1799 fax

System Category: Virtual Worlds Navigation, Control and Interaction (Single User)

System Application: Scientific Visualization (Biochemistry)

Abstract
Using a head-mounted display and hand-held input device, the user can fly through a three-dimensional graphical model of a protein molecule, getting inside of it and seeing from different angles. The atoms are represented as spheres, color-coded for atom type, with each sphere appearing to be about a foot in diameter. The protein molecule is dihydrofolate reductase (DHFR), and positioned nearby is a second molecule, methotrexate, which is a drug used in cancer treatment. Methotrexate is known to dock with DHFR, and the two molecules are positioned slightly disengaged from the docking position so that the user can fly in and see the key-and-lock spatial match at the docking site.

In addition to the touching-colored-spheres graphical model of the molecules, some other representations are available. In the ball-and-stick model, atoms are represented as smaller spheres and the bonds beween atoms are represented as narrow sticks. In the ribbon model, a curving ribbon-like shape shows the polypeptide backbone of the protein. These representations can be spatially superimposed.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

The user pushes the button on the hand-held input device and translates through the virtual world in the direction pointed by the input device. Flying speed is constant. The user must turn his or her head and body to look and fly in different directions. The scale, orientation and relative positions of the molecules are kept fixed.

The operator procedure is to run the molecule simulation program and then put the Head-Mounted Display onto the user's head and put the manual input device into his or her hand. Then the operator shows the user how to press a button and point the input device to fly through the virtual world. We limit the user to just flying to make it simple and quick for the user to get the idea and start flying around.

The operator can type a command at the keyboard to restore the program's starting configuration. The operator can also type commands to switch between the different molecular model types which the user is seeing. Since the system can display 490K spheres/second, quite complex molecules, such as the poliovirus, can be flown around.

Virtual Reality Applications Gallery Submission Form

Contact Person System Developer Name(s) Warren Robinett, Richard Holloway

Organization University of North Carolina, Computer Science Department

Street Address

City Chapel Hill State NC

Zip/Postal Code 27599-3175 Country USA

Daytime Telephone (919) 962-1700 Fax (919) 962-1799

System Title Flying Through a Protein Molecule

System Category Virtual Worlds Navigation, Control and Interaction (Single User)

System Application(s): Scientific Visualization (Biochemistry)

Abstract / System Description # **(200 words suitable for the conference program)**

See attached.

System Contents

See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, ~~regardless of how many are accepted.~~

Floorplan attached? [X] yes [ ] no Square footage required

Electrical power required See attached.

Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.

Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

- [X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.
- [X] My videotape-document(s) to be considered for broadcast.
- [X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.
- [X] Excerpts from the video-document or photos of my system to be used for the above purposes.
- [ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.

Signature [illegible] Date 1/29/91

Please send completed submittals to:
Steve E. Tice
SIGGRAPH '91 Virtual Reality Applications Gallery Chair
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

A Multi-Player Three-Dimensional Adventure Game

Warren Robinett
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1700
(919) 962-1799 fax

System Category: Multiple-User Virtual World Interaction

System Application: Entertainment

Abstract
This three-dimensional adventure game consists of a network of interconnected rooms, through which the user can move, with objects located in these rooms, which the user can pick up and manipulate. Each room or object is represented graphically. The user has two basic capabilities -- flying through the world and grabbing objects.

An object has a position and orientation in the virtual world. In addition, specific types of objects have behaviors and interactions with one another. Objects which move around and initiate actions are called *creatures*. Objects which have an interior within which other objects can be stored or transported are called *containers*. Objects which transport the user from place to place are called *vehicles*. Object which the user can manipulate to obtain desired results are called *tools*. Objects interact with one another, usually triggered by proximity or relative position, to cause events within the simulated world. Events are often marked by sounds or by animation of the graphics of one of the interacting objects. By picking up and moving an object near another object, the user can cause events to occur. This adventure game contains a several creatures, containers, vehicles and tools.

Two users simultaneously inhabit the network of rooms, and if they enter into the same part of the virtual world, they can see and interact with each other.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

The hand-held input device has two buttons -- one is for flying and the other is for picking up objects. When the fly-button is pressed, the user translates through the virtual world in the direction pointed by the input device. When the grab-button is pressed, if an object is near the hand, a rigid connection between the hand and object is formed, and when the button is released, this connection dissolves. The user can fly while holding an object, which is done by holding down the grab-button while using the fly-button to fly. The user can cause things to happen in the simulated world by moving one object near another one. For example, moving the throttle object of the train object will cause the train to start moving.

The operator procedure is to run the adventure game program and then put the head-mounted displays onto the users' heads and put the manual input device into their hands. Then the operator shows the users how to press a button and point the input device to fly through the virtual world, and how to pick up an object by pressing the other button. The adventure game world can be left running while switching users, so that only one new user needs to be instructed at a time.

The operator can type a command at the keyboard to restore the program's starting configuration.

Possible enhancements

We will have the system working with at least two users in the same virtual world, each with a stereoscopic head-mounted display, and a manual input device. If possible, we will configure the hardware to allow three or four users simultaneously interacting. We have enough head-mounted displays, trackers, and input devices for four users, and Pixel-Planes 5 can generate the 8 video signals needed. However, there are some problems that have to be solved to get all of these components to work together.

Virtual Reality Applications Gallery Submission Form

Contact Person	System Developer Name(s) Warren Robinett		
Organization	University of North Carolina, Computer Science Department		
Street Address			
City	Chapel Hill	State	NC
Zip/Postal Code	27599-3175	Country	USA
Daytime Telephone	(919) 962-1700	Fax	(919) 962-1799
System Title	A Multi-Player Three-Dimensional Adventure Game		
System Category	Multiple-User Virtual World Interaction		
System Application(s):	Entertainment		

Abstract / System Description # (200 words suitable for the conference program)

See attached.

System Contents

See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, ~~regardless of how many are accepted.~~

Floorplan attached? [X] yes [ ] no Square footage required

Electrical power required See attached.

Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.

Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

- [X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.
- [X] My videotape-document(s) to be considered for broadcast.
- [X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.
- [X] Excerpts from the video-document or photos of my system to be used for the above purposes.
- [ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.

Signature Warren Robinett Date 1/29/91

Please send completed submittals to:
Steve E. Tice
SIGGRAPH '91 Virtual Reality Applications Gallery Chair
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

A True-Scale Model of the Solar System and Nearby Stars

Warren Robinett
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1700
(919) 962-1799 fax

System Category: Virtual Worlds Navigation, Control and Interaction (Single User)

System Application: Education

Abstract
This three-dimensional graphical model includes the sun, 9 planets, all their moons and rings, several dozen major asteroids and comets, the Oort comet cloud, and also the several dozen nearest stars. Various graphical visualization aids can be selectively enabled, such as orbits, orbital planes, rotational axes, rotational planes, gravitational spheres of influence, planetary surfaces, surface features, atmospheres and clouds, interior features, and markers to help locate tiny bodies in vast spaces.

The user can fly through the model, tilt the model to any orientation, and make the model grow or shrink. Expanding and shrinking the model is necessary to deal with the vast scale differences -- from the diameter of Mercury to the diameter of Pluto's orbit is a factor of a million, and out to the nearest star is another factor of 4000. Shrinking down the model and then expanding it around a continuously-steered center of scaling is a powerful technique for moving great distances through the model.

The orbital motions and rotations of the bodies are modeled, and time can be run forward and backward, or stopped. There are also great temporal scale differences in the solar system, for example, between the rotational speed of Mercury and the orbital speed of Pluto, so the rate of time passage can be scaled, too.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

The user can push a button on the hand-held input device to change between control modes. Each mode has its own 3D icon-shape superimposed onto the user's hand, giving the user feedback about the current command mode. A second button initiates commands in the current command mode. In **fly** mode, the user translates through the virtual world in the direction pointed by the manual input device. In **tilt** mode, the user grabs the fabric of space and re-orients the model. In **shrink** mode, the user shrinks the model around the hand as the center of scaling. In **expand** mode, the model expands around the hand as the user continuously steers into where he or she wants to end up. In **rendevous** mode, the user locks onto the frame of reference of the body nearest the hand. In **time** mode, the motion of the hand controls the rate and direction of time passage. There are several modes which allow turning on and off the various visualization graphics such as orbits, surface features, etc.

The operator procedure is to run the solar system simulation program and then put the Head-Mounted Display onto the user's head and put the manual input device into his or her hand. Then the operator shows the user how to use one button to change command mode and the other button to activate commands. One of the commands restores the solar system model to the starting configuration. The operator can also do this by typing a command on the workstation keyboard.

Possible enhancements

If we have time, we might add the ability to follow a canned orbital trajectory through the solar system, such as the path taken by Voyager past the large planets, or Apollo 11 from Earth to the Moon. An orbital trajectory is an interesting four-dimensional object. As the user moved along a trajectory, time would be advancing and so the bodies would be moving in their own orbits.

Virtual Reality Applications Gallery Submission Form

Contact Person System Developer Name(s) Warren Robinett

Organization University of North Carolina, Computer Science Department

Street Address

City Chapel Hill State NC

Zip/Postal Code 27599-3175 Country USA

Daytime Telephone (919) 962-1700 Fax (919) 962-1799

System Title A True-Scale Model of the Solar System and Nearby Stars

System Category Virtual Worlds Navigation, Control and Interaction (Single User)

System Application(s): Education

Abstract / System Description # (200 words suitable for the conference program)

See attached.

System Contents

See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, ~~regardless of how many are accepted.~~

Floorplan attached? [X] yes [ ] no Square footage required

Electrical power required See attached.

Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.

Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

- [X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.
- [X] My videotape-document(s) to be considered for broadcast.
- [X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.
- [X] Excerpts from the video-document or photos of my system to be used for the above purposes.
- [ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.

Signature Warren Robinett Date 1/29/91

Please send completed submittals to:

Steve E. Tice

SIGGRAPH '91 Virtual Reality Applications Gallery Chair

SimGraphics Engineering Corporation

1137 Huntington Drive

South Pasadena, CA 91030-4563

A Mountain Bike with Force Feedback for Indoor Exercise

Ryutarou Ohbuchi
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1700
(919) 962-1799 fax

System Category: Force feedback mountain bike simulator (Single User)

System Application: Physical fitness

Abstract

This is an exercise bike with force feedback to the pedaling, in addition to 3D visual feedback, that rides around a synthetic mountainous terrain. Primary purpose of this system is to provide indoor cardiovascular exercise that is not too boring. A multiple speed bike fixed on a stand is outfitted with a computer controlled eddy current resistance device to give rider a variable resistance to the pedaling motion. As the rider rides the bike around in a synthetic mountainous terrain, which is displayed as a shaded 3D model on the screen, the pedalling resistance changes depending mainly on the grade of the terrain; you have to pedal harder if it is an uphill ride. Factors other than the grade, such as rider's weight as well as wind strength and direction affects the pedal resistance. Variable workload makes exercise bike ride physically more stimulating, improving the quality of exercise.

Heading of the bike is controlled by the handlebar, as in normal bike, whose orientation is sensed by the rotary encoder. Rotation of the wheel is sensed by another rotary encoder to control the bike's movement in the terrain. Braking motion is not sensed, but it actually decelerates the rotation of rear wheel. Since the resistance is provided by an eddy current device, there is a limitation; negative resistance can not be provided. That is, there will be no acceleration of bike's wheel even if the rider points the bike downhill, even though the bike in the virtual terrain will start accelerating downward. No force feedback other than the pedaling resistance, such as jolting, banking, tumbling over, etc. are simulated in the force feedback.

To make the ride a bit more stimulating, the terrain includes obstacles, such as trees and boulders, to be avoided. If the bike collides with the obstacle, there is a collision detection which let rider know of the collision. Collision is signaled by tinting the screen with red. The visual effects includes a fog, which makes collision avoidance more interesting, depending on its thickness.

Notes to Jury

All the Virtual Reality Application Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, regardless of how many are accepted.

User interaction, Operator procedures

Interaction by the user with the system is based on a very popular metaphor, a bicycle, which most of the people is familiar with. Thus the interaction requires very little explanation and/or training. It is safer since there will be no actual collision or falling off from the bike, not to mention collision with a motor vehicle. The rider of the bike, interact with the system by riding the bike. The rider's steering motion and pedaling motion is sensed to provide appropriate image and resistance change.

The operator procedure is to run the virtual exercise mountain bike is to run the program and let the user mount the bike, and ride. Most of the user will understand the user interface without any explanation. Before start riding, the user will be asked of desired strength of exercise as well as his/her weight to calculate proper resistance. The operator can turn on and off the fog upon user's request.

Possible enhancements

If we have time, we might add 'adventure game' like character to the terrain, to let the user explore the terrain which hides a few secrets. Also, it was pointed out by the past user that the perspiration of the rider can be a problem. Since the bike is not moving, perspiration will not dry out as in real bike ride that inevitably gets wind. It might also mechanically degrades the bike itself over time. A computer controlled fan to generate wind of variable strength, possibly of variable direction, may be added.

Virtual Reality Applications Gallery Submission Form

Contact Person System Developer Name(s) Ryutarou Ohbuchi

Organization University of North Carolina, Computer Science Department

Street Address

City Chapel Hill State NC

Zip/Postal Code 27599-3175 Country USA

Daytime Telephone (919) 962-1700 Fax (919) 962-1799

System Title A Mountain Bike with Force Feedback for Indoor Exercise

System Category Force Feedback Mountain Bike Simulator (Single User)

System Application(s): Physical Fitness

Abstract / System Description # (200 words suitable for the conference program)

See attached.

System Contents

See attached.

Notes to Jury All the Virtual Reality Gallery submissions from the University of North Carolina require the same hardware, so if more than one is accepted, only one of them can be run at a time. The space requirement (one 26 x 20 foot area) and shipping costs are the same, ~~regardless of how many are accepted.~~

Floorplan attached? [X] yes [ ] no Square footage required

Electrical power required See attached.

Maximum system dimensions: length 26 ft. width 20 ft. depth 15 ft.

Estimated set-up/take-down time 18 hrs./8 hrs. Estimated shipping cost $5000

Release I grant SIGGRAPH my permission for (check all that apply):

[X] My videotape-document(s) to be considered for the SIGGRAPH '91 virtual reality applications gallery video documentary or the art show video catalog and/or video review.

[X] My videotape-document(s) to be considered for broadcast.

[X] My videotape-document(s) to be reproduced for virtual reality applications gallery publicity purposes.

[X] Excerpts from the video-document or photos of my system to be used for the above purposes.

[ ] My work to be considered by other SIGGRAPH '91 events (interactive systems only)

All of the above information is true and complete.

Signature R. Ohbuchi (Ryutarou Ohbuchi) Date 1/29/91

Please send completed submittals to:
Steve E. Tice
SIGGRAPH '91 Virtual Reality Applications Gallery Chair
SimGraphics Engineering Corporation
1137 Huntington Drive
South Pasadena, CA 91030-4563

UNC Hardware for Virtual Reality Application Gallery

contact person: Henry Fuchs
University of North Carolina, Computer Science Department
Chapel Hill, NC 27599-3175, USA
(919) 962-1911
(919) 962-1799 fax

System Contents (hardware and software)

Two unusual hardware capabilities will be featured:

- **Pixel-Planes 5**, a new graphics engine designed and built at UNC by Henry Fuchs, John Poulton, and team from UNC's Microelectronic Systems Laboratory.

- A new **outward-looking optical tracker** prototype being developed at UNC. It has a very large working volume (20 x 20 feet) which is achieved by erecting over the working area a special ceiling in which are embedded infrared LED beacons. The optical tracker system has a triad of photodiodes mounted on the head, each one of which measures a 2D position of an LED in its field of view, and the appropriate LEDs are turned on as the user moves around. The user's position and orientation are computed from the known geometry of the LEDs, photodiodes and optical elements. This large working volume will permit the user to walk around through a 20 x 20 foot virtual world. We propose to demonstrate this by the illusion of walking around part of a house.

We will therefore require an unusually large area, shown on the attached floor plan. In addition to these specialized items, we plan to use Polhemus magnetic trackers, VPL EyePhones, hand-held input devices, and a Macintosh dedicated to playing sounds. The software for each submission is a stand-alone simulation program that makes use of the Pixel-Planes and Virtual Worlds software libraries (PPHIGS and VLIB) written at UNC. A large monitor will allow spectators to see and hear what the user sees and hears. All required hardware will be brought from UNC.

The proposed demonstration of a 3D adventure game will put multiple users into the same virtual world. Pixel-Planes 5 can generate 8 independent viewpoints as 8 separate video signals, so we plan to put at least two, and possibly three or four, users into the same virtual world. Each user will have a dedicated stereoscopic head-mounted display, head and hand tracker, and manual input device.

Demo date at UNC: Friday, April 26, 1991

Electrical Power Requirements

2 circuits of 208 volt 3-phase Y @ 30 Amps per phase
1 circuit of 120 volt 1-phase @ 30 Amps
3 circuits of 120 volt 1-phase @ 20 Amps

Estimated Set-up and Take-down time

We estimate that it will take 12 to 18 hours to set up everything, and 8 hours to break it down.

Estimated Shipping Cost

We estimate an expense of $5000 for transporting the hardware from North Carolina to Las Vegas. We plan to drive a truck containing the equipment.

We would hope that our shipping expenses were fully covered by SIGGRAPH, but if our work is accepted to the Virtual Reality Applications Gallery and this cost is a problem, we can negotiate partial support.

Floor Plan

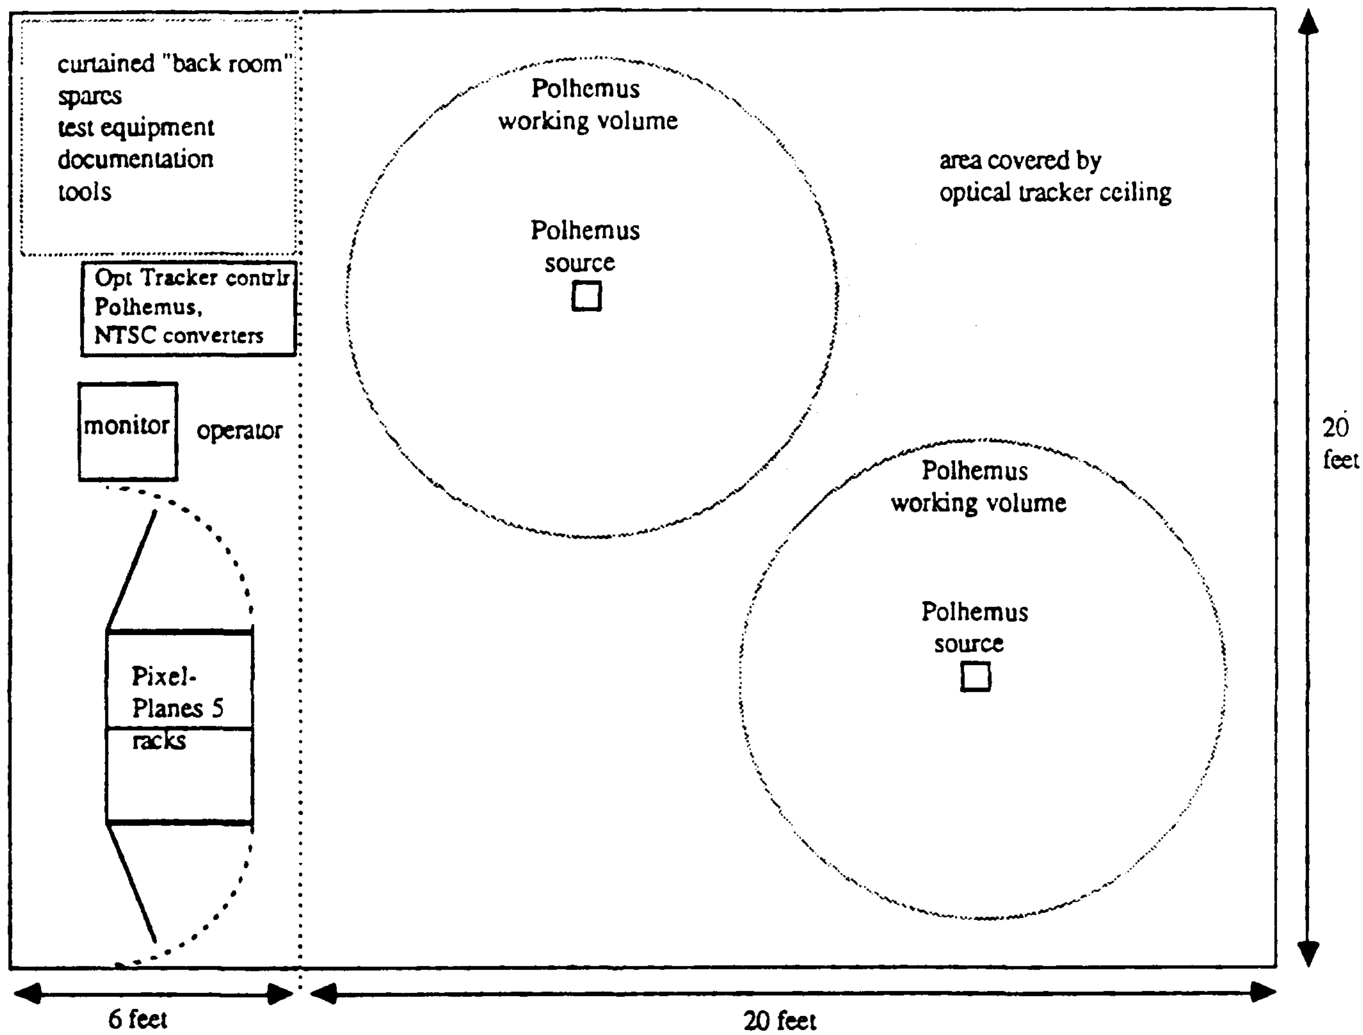

Floor plan for UNC

26 x 20 = 520 sq. ft.

Note that this space can be shared by all accepted submissions from UNC

Note that the 20 x 20 foot stippled area in the floor plan will be covered by a special ceiling brought from UNC. The support structure for the ceiling panels will be supported by posts at the four corners of the 20 x 20 foot area. (Polhemus tracking systems will be used for the multiple user applications because the optical tracker prototype is expected to support, by Siggraph'91, only a single user.)

Possible application of portable personal visualization systems in the near future: Ultrasound Examination

Physician wearing a portable personal visualizer observes a fetus by ultrasound imaging, allowing him to correlate the 3D visual information with that obtained by tactile examination. The hand-held transducer is tracked stereoptically by miniature cameras mounted on the user's visor. Two upward-looking cameras on the head-gear track an array of IR LED beacons on the ceiling as the user moves about the room. Conventional ultrasound equipment can be seen in the background.

University of North Carolina at Chapel Hill
Department of Computer Science
Head Mounted Display Project