

23. VPL-Research, *DataGlove Model 2 Users Manual*, Inc., 1987.
24. Ware, C., and Osborne, S., *Exploration and Virtual Camera Control in Virtual Three Dimensional Environments*, ACM SIGGRAPH 1990 Symposium on Interactive 3D Graphics, March, 1990.
25. Weimer, D., and Ganapathy, S., *A Synthetic Visual Environment with Hand Gesturing and Voice Input*, Proceedings of the ACM SIGCHI Human Factors in Computer Systems Conference, April, 1989, 235-240.
26. Zachary, G., and Gentile, C., *Personal Communication* (letter), VPL Research, Inc., July 18, 1990. VPL/AGE Power Glove Support Program, VPL Research, Inc., 656 Bair Island Road, Suite 304, Redwood City, CA 94063, (415) 361-1710.
27. Zimmerman, T., Lanier, J., Blanchard, C., Bryson, S., and Harvill, Y., *A Hand Gesture Interface Device*, Graphics Interface '87, May, 1987, 189-192.

many researchers' hands as possible. The virtual reality systems previously described in the literature cost more than most researchers can afford. We have shown that for less than \$5,000, or five dollars per day over three years, researchers can use a head-mounted display with glove and voice input. Our system has a higher spatial resolution than any previous system, and is significantly lighter than previous systems [4,7]. For glove input, the Power Glove has provided excellent spatial accuracy and usable finger bend data. Based on experience with our system, we have found that interaction latency is significantly more important than display resolution or stereoscopy, and that the user can greatly benefit from the display of reference objects, such as a ground plane and a virtual vehicle.

ACKNOWLEDGMENTS

This work could not have proceeded without the help we received from Chris Gentile of AGE. Novak of Mattel, Inc. also provided assistance with an early draft of the paper. We would also like to thank Ronald Williams, Pramod Dwivedi, Larry Ferber, Rich Gossweiler, and Chris Long at the University of Virginia for their help.

REFERENCES

1. Becker, A., *Design Case Study: Private Eye*, Information Display, March, 1990.
2. Blanchard, C., Burgess, S., Harvill, Y., Lanier, J., and Lasko, A., *Reality Built for Two: A Virtual Reality Tool*, ACM SIGGRAPH 1990 Symposium on Interactive 3D Graphics, March, 1990.
3. Brett, C., Pieper, S., and Zeltzer, D., *Putting It All Together: An Integrated Package for Viewing and Editing 3D Microworlds*, Proceedings of the 4th Usenix Computer Graphics Workshop, October, 1987.
4. Brooks, F., *Walkthrough - A Dynamic Graphics System for Simulating Virtual Buildings*, Proceedings of the 1986 ACM Workshop on Interactive Graphics, October, 1986, 9-21.
5. Brooks, F., *Grasping Reality Through Illusion: Interactive Graphics Serving Science*, Proceedings of the ACM SIGCHI Human Factors in Computer Systems Conference, Washington, D.C., May 17, 1988, 1-11.
6. Eglowstein, H., *Reach Out and Touch Your Data*, Byte, July 1990, 283-290.
7. Fisher, S., McGreevy, M., Humphries, J., and Robinett, M., *Virtual Environment Display System*, Proceedings of the 1986 ACM Workshop on Interactive Graphics, October, 1986, 77-87.
8. Fisher, S., *The AMES Virtual Environment Workstation (VIEW)*, SIGGRAPH '89 Course #29 Notes, August, 1989. (included a videotape).
9. Fisher, S., *Personal Communication* (electronic mail), Crystal River, Inc., September 28, 1990.
10. Foley, J., van Dam, A., Feiner, S., and Hughes, J., *Computer Graphics, Principles and Practices*, Addison-Wesley Publishing Co., 1990.
11. Kaufman, A., Yagel, R. and Bakalash, R., *Direct Interaction with a 3D Volumetric Environment*, ACM SIGGRAPH 1990 Symposium on Interactive 3D Graphics, March, 1990.
12. Kelley, A. and Pohl, I., *A Book on C*, second Edition, Benjamin/Cummings Publishing Company, Inc., 1990.
13. Lanier, J., *Plenary Address on Virtual Reality*, Proceedings of UIST: the Annual ACM SIGGRAPH Symposium on User Interface Software and Technology, November, 1989.
14. Ming, O., Pique, M., Hughes, J., and Brooks, F., *Force Display Performs Better than Visual Display in a Simple 6-D Docking Task*, IEEE Robotics and Automation Conference, May, 1989.
15. Novak, *Personal Communication* (telephone call), January 3, 1991.
16. Ouh-young, M., Pique, M., Hughes, J., Srinivasan, N., and Brooks, F., *Using a Manipulator For Force Display in Molecular Docking*, IEEE Robotics and Automation Conference 3 (April, 1988), 1824-1829.
17. Pausch, R., *A Tutorial for SUIT, the Simple User Interface Toolkit*, Technical Report Tech. Rep.-90-29, University of Virginia Computer Science Department, September 1, 1990.
18. Pausch, R., and Williams, R., *Tailor: Creating Custom User Interfaces Based on Gesture*, Proceedings of UIST: the Annual ACM SIGGRAPH Symposium on User Interface Software and Technology, October, 1990, 123-134.
19. Pausch, R., and Gossweiler, R., "UserVerse: Application-Independent Object Selection Using Inaccurate Multi-Modal Input," in *Multimedia and Multimodal User Interface Design*, edited M. Blattner and R. Dannenberg, Addison-Wesley, 1991 (to appear).
20. Rabb, F., Blood, E., Steiner, R., and Jones, H., *Magnetic Position and Orientation Tracking System*, IEEE Transaction on Aerospace and Electronic Systems, 15, 5 (September, 1979), 709-718.
21. Scully, J., *Personal Communication* (letter), Ascension Technology, Inc., PO Box 527, Burlington, VT 05402 (802) 655-7879, June 27, 1990.
22. Sturman, D., Pieper, S., and Zeltzer, D., *Hands-on Interaction With Virtual Environments*, Proceedings of UIST: the Annual ACM SIGGRAPH Symposium on User Interface Software and Technology, November, 1989.

low display resolution, but notice lag in the 200 milli-second range.

Stereoscopy is not essential

Users of bi-ocular and monocular (one eye covered with a patch) versions of our system could maneuver and interact with objects in the environment. Since a straightforward implementation of stereo viewing slows down graphics by a factor of two or doubles the hardware cost, it is not always an appropriate use of resources.

A ground plane is extremely useful

Non-head-mounted virtual worlds sometimes introduce a ground plane to provide orientation [3,22]. In expensive head-mounted systems, the floor is usually implicitly included as a shaded polygon. We found the need in our system to include an artificial ground plane for reference, drawn as a rectangular grid of either lines or dots.




Figure 5. Displaying the Vehicle to the User

Display the limits of the “vehicle” to the user

In virtual reality, a user's movement is always constrained by the physical world. In most systems this manifests with the user straining an umbilical cord. Even in systems with no umbilical and infinite range trackers, this problem will still exist. Unless the user is in the middle of a large, open space the real world will

limit the user's motions. In the VIEW system [7,8] a waist-level hexagon displays the range of the tracker, but is part of the world scene and does not move as the user flies. We treat the user as always residing in a “vehicle” [24]. The vehicle for a Polhemus is roughly a ten foot hemisphere. If the user wishes to view an object within the range of the vehicle, he may walk over to it, thereby changing his own location within the vehicle. If, however, the user wishes to grab an object not currently in the vehicle, he must first fly the vehicle until the desired object is within the vehicle, as shown in Figure 5. Note that the user may be simultaneously moving within the vehicle and changing the vehicle's position in the virtual world, although in practice our users do not combine these operations. For small vehicles it is probably appropriate to always display their bounds but for larger vehicles it may be better to show their bounds only when users are near the edges.

FUTURE WORK

Adding voice input will allow us to experiment with a model we have developed to support object selection via simultaneous voice and gesture input. We have already built a prototype of this selection model using a display screen in combination with voice and gesture input and will attempt to repeat those results using a head-mounted display[19].

We also will be addressing the registration problem, or the correct matching of real and synthetic objects. Until force-feedback technology improves from its current state[14,16], glove-based systems will have to use real-world objects as tactile and force feedback to the user for some tasks. For example, one could perform a virtual version of the popular magic trick “cups and balls” by moving real cups on a real table, but having arbitrary virtual objects appear under the cups. The graphics for the cups, which can be grasped and moved, must closely correspond to the real world cups. By attaching trackers to real world objects, we will study how closely the visual image must match reality to avoid user dissatisfaction. A second approach to this problem is to use the Private Eye as a heads up display, wearing it over only one eye and allowing the user to correlate the real world and synthetic graphics.

We are currently pursuing support to create a laboratory with between ten and twenty low cost virtual reality stations. By providing reasonable access to an entire graduate or undergraduate class, we suspect we may quickly develop a large number of new interaction techniques. Jaron Lanier has commented that in virtual reality, “creativity is the only thing of value” [13]. A good way to spark creative breakthroughs is to increase the number of people actively using the technology. We are also exploring the possibility of creating a self-contained, portable system based on a laptop machine.

CONCLUSIONS

The field of virtual reality research is in its infancy, and will benefit greatly from putting the technology into as

away from the receivers, the signals degrade. Although some signal is received up to a 90 degree angle, Mattel claims the glove is only usable at up to roughly 45 degrees. When the glove is within five to six feet of the receivers, its (x, y, z) coordinate information is accurate to within 0.25 inches [15]. In addition to position information, the Power Glove provides roll information, where roll is the angle made by pivoting the hand around the axis of the forearm. Roll is reported in one of twelve possible positions.

Finger bend is determined from the varying resistance through materials running the length of the finger. The user's thumb, index, middle, and ring finger bend are each reported as a two-bit integer. This four-position granularity is significantly less than the resolution provided by the VPL DataGlove, but most of the gestures used in previously published virtual reality systems can be supported with only two bits per finger [2,8,11,25].

The only hardware we plan to add to our system is for voice input. Several small vocabulary, speaker-dependent input devices exist for the PC, all costing several hundred dollars. Once this is added, many of the commands currently given by hand gesture will be replaced by voice input.

All software for our system is locally developed in ANSI-standard C [12]. We have a simple version of PHIGS [10] and are using a locally developed user interface toolkit [17]. Our low-level graphics and input handling packages have been widely ported, and allow our students to develop applications on SunsTM, Macin-

toshesTM, or PCs before running them on the machine equipped with the head-mounted display. We are currently developing a three-dimensional glove-based object editor.

Although fast enough to be used, the limiting factor of our system's performance is the speed of line scan conversion. We draw monochrome wire frame objects, but are limited by the hardware's ability to draw lines. The hardware can render 500 vectors per second (of random orientation and length) but our CPU can execute the floating point viewing transformations for 3,500 vectors per second. In practice, we tend to use scenes with roughly 50 lines and we sustain a rate of 7 frames per second. High-performance scan-conversion boards currently exist which would substantially improve our rendering capabilities, and we expect their price to drop substantially in the coming year.

The major limitation of our system's usability is the lag of the Polhemus Isotrak. Other researchers using the Isotrak have also reported this problem; no one has precisely documented its duration, but it is within 150 and 250 milliseconds[9]. Ascension Technology, Inc. recently announced the BirdTM, a \$5,000 competitor to the Polhemus Isotrak with a lag of only 24 milliseconds[21].

The existing system, when augmented with voice, will still cost less than \$5,000 in hardware (\$750 for each eye, \$3,000 for the head tracker, \$80 for the Power Glove, and ~\$400 for the voice input). For less than the cost of a high resolution color monitor, we have added the I/O devices to support a complete virtual reality system.

RESEARCH OBSERVATIONS

Fred Brooks [5] has commented that:

A major issue perplexes and bedevils the computer-human interface community -- the tension between narrow truths proved convincingly by statistically sound experiments, and broad 'truths,' generally applicable, but supported only by possibly unrepresentative observations.

Brooks distinguishes between findings, observations, and rules-of-thumb, and states that we should provide results in all three categories, as appropriate. Most research presented to date in virtual reality are either what Brooks calls observations or rules-of-thumb, and we continue in this vein, stating our experience:

The quality of the graphics is not as important as the interaction latency

If we had to choose between them, we would prefer to decrease our tracking lag than increase our graphics capabilities. Although we have much greater spatial resolution than other head-mounted displays, this does not seem to significantly improve the quality of our system. Our experience confirms what has been discovered at VPL Research and NASA AMES research center: if the display is driven by user head motion, users can tolerate



Figure 4.Attaching to a Baseball Cap




Figure 2.Private Eye Design: Top View

280 vertical red pixels against a black background. Optics between the user's eye and the display surface make the image appear to be one to three feet wide, "floating" several feet away.

The Private Eye is implemented with a vertical column of 280 red LEDs, manufactured as a unit to pack them as densely as possible. To fill the entire visual display area, the LEDs are switched on and off rapidly as a vibrating mirror rotates through the 720 different vertical columns of the display, as shown in Figure 2. The Private Eye can "shadow" a standard CGA display with resolution of either 640 by 200 or 320 by 200 pixels, or it can be accessed a library which supports a spatial resolution of 720 by 280 resolution. The library allows the painting of text and bitmaps, but does not support graphics primitives such as lines; therefore, we use the device by shadowing a CGA display.

Reflection Technologies is marketing the Private Eye primarily as a "hands-busy" display; Figure 3 shows how the company expects most users to wear the device. The user can look down into the display without obstructing normal vision. Figure 4 shows how we mount two Private Eyes underneath a baseball cap. We have also used sunglasses with leather sides to shield the user from peripheral distractions. Our head-mounted display can either be stereoscopic or bi-ocular (each eye receives the same picture).

We use a Polhemus 3Space Isotrak[20] to track the position and orientation of the user's head. The Isotrak senses changes in a magnetic field and reports three spatial (x, y, z) and three angular (yaw, pitch, roll) coordinates 60 times each second. Our system uses the Mattel Power Glove as an input device for position and gesture information. The glove is manufactured by Mattel, Inc., under licence from Abrams-Gentile Entertainment, Inc. (AGE). The Power Glove is provided to retail stores at a wholesale cost of 62 dollars and is sold at a retail cost ranging between 70 and 100 dollars. Although Mattel does not release unit sales figures, they report that in 1989 the Power Glove generated over 40 million dollars in revenue, implying that over half a million gloves were sold that year.

Early glove research was conducted at VPL Research, Inc., the manufacturers of the DataGlove™[23,27]. The DataGlove uses fiber optics to determine finger bend and a Polhemus tracker to determine hand position. Neither of these technologies could be mass produced easily, so the Power Glove uses variable resistance material for finger bend, and ultrasonics for hand position.

The Power Glove is marketed as a peripheral for the Nintendo Entertainment System™. To thwart rival toy manufacturers, the data stream between the Power Glove and the main Nintendo unit is encrypted. When the Power Glove was originally introduced, it was rumored that dozens of research groups across the country began working on decrypting this data stream, and that several groups actually broke the code. An article appeared in Byte magazine describing how to attach the glove as a serial device, but only allowed the glove to emulate a joystick-type input device[6]. Rather than engaging in cryptography, we phoned Chris Gentile at AGE and described our research goals. He allowed us to sign a non-disclosure agreement and within days sent us a decrypting device that allows us to use the glove as a serial device communicating over an RS232 line. AGE and VPL Research have recently announced the VPL/AGE Power Glove Education Support Program[26] and plan to provide a low cost glove with 5 degrees of freedom for between 150 and 200 dollars.

The Power Glove uses two ultrasonic transmitters on the back of the user's hand and three wall-mounted receivers configured in an L-shape. The glove communicates successfully within ten to fifteen feet of the receivers when it is oriented towards them. As the glove turns




Figure 3. Expected Private Eye Use

Virtual Reality on Five Dollars a Day

Randy Pausch

Computer Science Department
University of Virginia
Thornton Hall
Charlottesville, VA 22903
Pausch@Virginia.edu

ABSTRACT

Virtual reality systems using head-mounted displays and glove input are gaining popularity but their cost prohibits widespread use. We have developed a system using an 80386 IBM-PC™, a Polhemus 3Space Isotrak™, two Reflection Technology Private Eye™ displays, and a Mattel Power Glove™. For less than \$5,000, we have created an effective vehicle for developing interaction techniques in virtual reality. Our system displays monochrome wire frames of objects with a spatial resolution of 720 by 280, the highest resolution head-mounted system published to date. We have confirmed findings by other researchers that low-latency interaction is significantly more important than high-quality graphics or stereoscopy. We have also found it useful to display reference objects to our user, specifically a ground plane for reference and a vehicle containing the user.

KEYWORDS: Virtual reality, head-mounted display, glove input, computer graphics, teleoperation, speech recognition, hand gesturing, three-dimensional interaction.

INTRODUCTION

Virtual reality systems are currently gaining popularity but the cost of the underlying hardware has limited research in the field. With any new technology, there is an early period where informal observations are made and large breakthroughs are possible. We believe that the best way to speed up this process with head-mounted display/glove input systems is to provide low cost versions of the technology so larger numbers of researchers may use it. We have developed a complete virtual reality system for less than \$5,000, or less than five dollars per day if amortized over a three-year period. We built the system because we had an immediate need and also to show that virtual reality research can be done without expensive hardware.

Our immediate interest in virtual reality interaction comes from the Tailor project[18], whose goal is to allow severely disabled children to control devices via gesture input. The Tailor system adjusts to each child's possible range of motion and converts motion in that range into analog con-

trol signals that drive software applications. To specify motion mappings, therapists with no technical background must specify one dimensional curves and two dimensional surfaces in three dimensional space. Using our low cost system, we will allow therapists to interactively manipulate a wire frame mesh by using the glove to grasp control points on the mesh.




Figure1. The Private Eye

Our system provides 720 by 280 spatial resolution and weighs 6 ounces, making it higher resolution and lower weight than head-mounted displays previously reported in the literature. In this paper, we present several design observations made after experience with our system. Our first observation is that increasing spatial resolution does not greatly improve the quality of the system. We typically decrease our resolution to increase our rendering speed. We also observe that stereoscopy is not critical, and that reference objects such as a ground plane and a virtual vehicle are extremely helpful to the user.

SYSTEM DESCRIPTION

The main processor for our system is a 2.5 MIP, 20 Mhz 386-based IBM-PC™ compatible with 640K of RAM, a 80387 floating point co-processor, and MS-DOS™. Our head-mounted display uses a combination of two Private Eye displays manufactured by Reflection Technology, Inc. [1]. Figure 1 shows a Private Eye, a 1.2 by 1.3 by 3.5 inch device weighing 2.5 ounces. The 1 inch square monochrome display surface has a resolution of 720 horizontal by

*Proceedings of the ACM SIGCHI Human Factors
in Computer Systems Conference, April, 1991,
New Orleans*

Virtual Reality on Five Dollars a Day

Randy Pausch

Computer Science Report No. TR-91-19
October 7, 1991

This work was supported in part by the National Science Foundation, the Science Applications International Corporation, the Virginia Engineering Foundation, the Virginia Center for Innovative Technology, and the United Cerebral Palsy Foundation.