Art-Related Virtual Reality Applications at the University of North Carolina at Chapel Hill

Richard L. Holloway

Department of Computer Science
University of North Carolina

Abstract

"Virtual Reality" is a field that has attracted much interest from people of many different disciplines. This is not surprising, since there are so many applications for virtual reality, from architectural building walk-throughs, medical applications and scientific visualization, to performance art.

The emphasis at the University of North Carolina at Chapel Hill (USA) is on improving virtual reality technology and using it to solve real problems. Some of the current applications of virtual reality at UNC are: architectural building walk-through, radiation treatment planning, molecular docking, and medical "X-Ray vision."

This paper gives an overview of virtual reality techniques and highlights work in progress and recent developments at Chapel Hill and discusses this work's relevance to art.

Introduction

For me to discuss how virtual reality relates to art is a little like a toolmaker who makes better chisels talking to sculptors about how to do statues. I will give my impressions of how virtual reality relates to art at the end of this paper, but for the most part I will stick to the explanation of how virtual reality is accomplished and of the many applications for which we are using it at the University of North Carolina at Chapel Hill, and leave the vivid imaginations of thousands of artists to decide how best to use it, and where it might go from here.

I. Overview of Virtual Reality Ideal

A. The Promise of Virtual Reality

The promise of virtual reality (VR) is that it offers us the ability to enter a "parallel universe" or "virtual world" in which all things are possible. This universe contains whatever we want or need it to at the moment- we can choose how it will act in order to satisfy our current need. If we are architects, this universe may consist of the building we are designing and the surrounding landscape; if we are radiologists, it may consist of a model of the current patient's anatomy, so that we can properly place a radiation beam to destroy a tumor; if we are biochemists, this universe may consist of a huge protein molecule, whose structure we come to understand by flying through it on a "magic carpet" and by manipulating its twistable bonds. In the parallel universe, we are not bound by physical laws which are not useful- what is imaginable is possible.

That's the *promise* of virtual reality; but what is the *reality* of VR?

The reality is that we do our best to *simulate* the parallel universe described above: we don headgear, handgear, and body suits; we walk on treadmills and "fly" by pointing a three-dimensional mouse and pressing a switch. The more convincing this simulation is, the closer we come to the promise of virtual reality.

B. Simulating Another World

In order to simulate reality, we need an understanding of the ways we sense and interact in the real world. Here in the real world, we use our *senses* to perceive the reality around us: there are sensory receptors scattered throughout our bodies which sense external stimuli and report to our brain. In order to simulate another world, we need to feed in different stimuli to at least some of these receptors. In addition, we need a way of detecting actions taken by the user so that we can reflect these actions in the virtual world in order to maintain the illusion over time.

One way to characterize this complex interaction is by imagining an interface between the user and the world (real or virtual), and categorizing sensory stimuli as his *inputs*, and anything that the user might do that must be reflected in the virtual world as his *outputs*.

Sensory inputs are described well by the literature of sensory-perceptual psychology; the list that follows is one possible categorization (of many) of the components of this interface. I have listed the outputs with the inputs, since every output can be associated with some input.

Type[1]	**Input / Output**
visual	what the user sees / his appearance
auditory	sound / voice, user-generated sounds
haptic[2]	touch, contact forces experienced / forces exerted
proprioceptic[3]	current body configuration
vestibular[4]	acceleration, orientation
olfactory	smell
gustatory	taste

With this model then, we associate inputs with real sensory receptors; the outputs are defined to be those actions that could serve as a stimulus to some other (real or virtual) user. For example, the auditory input is any sound that the user hears; his auditory output is any sound that he makes that another user in that virtual world could hear. The user's visual input is anything that he sees; his visual output is his appearance in the virtual world. Finally, the position and orientation of the user's body parts is his proprioceptic input *and* output: he can sense where his limbs, etc, are, and the system must also detect where they are in order to properly display his appearance.

In summary, in order to fool the user into believing that he's in another world, we must present alternative stimuli (sights, sounds, forces, etc), or "inputs" to his sensory systems. In order to keep the illusion up to date, we must detect any user "outputs": changes in his position or orientation, sounds generated, etc.

---

[1] This list contains only the most relevant to this field, and is a subset of those found in the psychology literature (see [5], [10]); I have omitted some categories for brevity and simplicity.

[2] haptics- "pertaining to sensations such as touch, temperature, pressure, etc. mediated by skin, muscle, tendon, or joint." [14]

[3] proprioception is "the sense of movement and bodily position," [5] and is one of the haptic senses.

[4] the vestibular system of the inner ear provides information about the linear and angular accelerations of the head. [10]

II. Virtual Reality Equipment at UNC Chapel Hill

To implement the model above, we need devices for providing the stimuli to the user, as well as sensing devices to detect user output. The VR equipment currently in use at UNC does not cover every category in the model above- there is simply too much to simulate and detect. What follows is a list of VR equipment now in use for our current applications.

A. Input Devices

1. Displays

Displays are arguably the most important feature of a VR system, since the visual system has such an impact on our perception (Christman [10] characterizes the human eye as "an extension of the brain."). The visual input in a VR system is usually computer-generated, and displayed on some type of screen (cathode ray tube (CRT), liquid crystal display (LCD), etc). Although the *head-mounted display* (HMD) is often used in VR systems, it is by no means the only way to give the user visual input. At UNC, we currently use four types of display:

a. *Regular CRT's:* due to their availability, these displays enjoy wide use, especially for the Walkthrough and molecular studies applications.

b. *Large, rear-projection screen:* This is a Sony Super Data® 12700 projection display, projected onto a 4' x 5' 4" rear-projection screen. It gives a wide-field-of-view color image, and is used often for the Walkthrough and molecular studies applications.

c. *Opaque Head-Mounted Displays:* We use both a commercially available VPL Eyephone® and a bicycle-helmet-mounted display, developed at the Air Force Institute of Technology in collaboration with UNC; these two are used for all applications except X-Ray Vision.

d. *See-Through Head-Mounted Displays:* We have two units (only one of which is currently operational) that were developed at UNC that use half-silvered mirrors to superimpose virtual objects on the real world [6] [9]; these are currently used only for X-Ray Vision. A new, more robust, high resolution model is under development.

The first two display types above can be fitted with stereo viewing plates (we use Tektronix alternating polarization plates) so that a stereo image is displayed when viewed with polarized glasses; the Eyephone® and see-through HMD's normally display in stereo; the AFIT is currently used in non-stereo mode.

2. Sound Equipment

Audio input is usually computer-controlled and may be recorded or computer-generated.

Currently, most of our sound generation is done by an Apple Macintosh®IIcx, and uses computer-generated and recorded sounds. Any sound can be recorded and played back under application control through the built-in speaker in the Macintosh®. Examples of sound in the featured applications are:

a) In a simple molecular application for the force feedback arm, collisions between two atoms is marked by a "knock" sound played by the Macintosh® (this is not yet feasible for more complicated molecules and so is not currently done).

b) In the Adventure Game application, a similar "knock" sound is generated when the user bumps into a virtual wall; a "whoosh" sound is played when the user fires the application's "vortex gun".

c) The Virtual Piano application's notes are generated by the Macintosh® as well.

d) The force feedback arm (described below) uses a simple bell sound to give user feedback when the arm is out of range.

3. Force Feedback / Motion Simulators

Haptic input can be given in all sorts of ways, since there are receptors all over the body for this type of input.

The most prominent haptic input device in use at UNC is the ARM (Argonne Remote Manipulator, a donation from Argonne National Laboratories), a six degree-of-freedom (3 forces, 3 torques) force feedback device. The ARM is currently used primarily for molecular docking research.

Another type of haptic input is given by the treadmill, which is used for the Walkthrough application. Although this form of feedback is not as direct as the ARM's forces, the treadmill nonetheless provides haptic input by giving the user the feeling that he is walking on a real surface.

B. Output Devices

1. Trackers

One of the most critical user outputs is the position and orientation of his body: head, torso, extremities, etc. If the user moves his head, we need to update his view of the world as shown on the display (assuming we are doing head tracking) to reflect his new viewpoint. In order to detect user motion, we have three types of tracking systems at our disposal:

a) *Polhemus Trackers:* We currently have three Polhemus Navigational Sciences 3Space® trackers. These trackers use a low-frequency magnetic field generated by one or two *sources* , and report the positions and orientations of (depending on the model) from one to four *sensors* at once. The working range of these trackers is a hemisphere less than five feet in radius, with an update rate equal to 60 Hz divided by the number of source-sensor pairs that are in use. The two main problems with the Polhemus are its *lag* (the time between a user's movement and the appearance of data reflecting that movement) and its limited working volume.

The Polhemus is currently used for all head-mounted display applications; one sensor is attached to each of the HMD's for head tracking, and another sensor has been placed inside of a hollowed out billiard ball, which serves as a 3-D mouse for hand tracking. Alternatively, we sometimes use the hand sensor on a guitar finger pick equipped with a microswitch.

b) *Optical Tracker:* An optical tracker is currently under development at UNC that should solve the problems we are having with the Polhemus. Its goal is to allow room-size working volumes with update rates in the neighborhood of 200 Hz. It is at present a working, limited-volume bench prototype [11] [12] [13].

c) *The ARM:* In addition to being a haptic input device, the ARM is also an effective hand tracking device, since the ARM's handgrip position can always be calculated from its joint angles.

In addition to these more general-purpose trackers, we also have a VPL DataGlove® for tracking the fingers of the user's hand (the DataGlove™ uses the Polhemus tracker for the hand position and orientation).

2. Audio Output

One of the most natural ways of communicating for humans is via voice; thus, devices for detecting auditory output in the form of speech are a natural part of a virtual reality system. We have recently acquired a DragonWriter® speech recognition system for our lab. This system is speaker-dependent, and thus has to be trained for each user. The system has not been used yet, but is promising for applications where the user will spend enough time with the system to justify the training time.

3. Motion Devices

If the user actually walks very far in the room where his VR equipment is connected, he may run into a wall that is not part of his virtual world, but is part of the real building, or he may run out of cable. For applications (such as Walkthrough) where the user will need to "walk through" a building or building-sized model, the user's translation output is detected by a specially modified treadmill, which is equipped with bicycle handle bars for steering and electronics for reporting changes in its virtual position and orientation.

For greater distances, a similarly modified bicycle has been used.

4. Other "Input Devices"

There is also a host of other devices that can be used in virtual reality systems which are traditionally called "input devices", since they provide input to a computer. We use them to detect user output in the form of button presses, knob twists, etc. Examples from the featured applications:

a) *Joysticks and Sliders:* The user's vertical position in the building model in Walkthrough is controlled by a slider. Also, when the user is not using the treadmill for moving through the floor plan, his movements can be controlled by a joystick; another joystick can be used for the head's orientation when the head-mounted display is not used.

b) *3-D Mouse Buttons:* The 3-D mouse mentioned in the Tracker section also has two microswitches for detecting user output. In several of the applications, a button press on the 3-D mouse allows the user to "grab" a virtual object, which stays attached to the user's hand for as long as the button is pressed. In the adventure game application, the user picks up the "vortex gun" with one sequence of button presses, and "fires" it with another. Flying through the virtual world is also initiated with a button press.

c) *ARM Dials and Switches:* There are dials on the ARM for controlling the level of force output and for general use, as well as a safety switch that enables any output forces at all.

5. Haptic Output

With the ARM system, the user outputs real forces in response to those produced by the ARM; these may move the ARM to a new position, which it must detect and act accordingly. In the other systems where the user grabs virtual object, the forces are virtual, but still may be classified as haptic output.

III. Virtual Reality Applications at the University of North Carolina at Chapel Hill

The department of Computer Science at the University of North Carolina at Chapel Hill has been active in virtual reality research for over two decades. The emphasis at UNC is on improving virtual reality technology by tackling real-world problems and letting these problems drive the research.

Descriptions of some of the most relevant projects follow. The projects are grouped into *research projects*, which are collaborative efforts which may take years, and *personal projects*, which are usually done by one or two people for exploring various aspects of virtual reality, or for class projects. The more mature projects are listed first.

(The majority of the VR research projects have been supported by the DARPA project "Advanced Technology for Portable Personal Visualization", Frederick P. Brooks, Jr. and Henry Fuchs, principal investigators (contract #DAEA18-90-C-0044), and by NIH Grant #5-R24-RR 02170-07, ONR Grant # N00014-86-K-0680, and NSF Grant # CCR-8609588.)

A. Research Projects

1. Molecular Studies

Tool-building for biochemists doing research on protein molecules has been going on at UNC for over twenty years, dating back to the GROPE I project in 1967 [4].

More recently, the GROPE III project has been used to show that force feedback improves performance in a molecular "docking" task (finding a minimum energy position and orientation for a drug molecule in the active site of a protein molecule) [4]. The user holds the handgrip of the ARM (described in II) with a virtual drug molecule attached to it. As he moves it around the large protein molecule, he experiences a simulation of the forces that would be exerted on a real drug molecule in that configuration by the protein molecule.

Viewing can be done either with the head-mounted display or with a PS-330 color vector display with a Tektronix alternating polarization plate and polarized glasses (the latter display is preferred for large models).

A newer molecular application is the "fly-through" of a molecule. The user can load in a protein molecule modeled as a cluster of spheres, put on the head-mounted display, and enter a world where angstroms are as long as meters, and atoms are as big as beach balls. The user can scale the model up and down, and fly through it by pointing the three-dimensional mouse and pressing a button. As long as the user presses the button, he will "fly" in the direction he's pointing. When he stops, he will be suspended in space at that point, free to walk around in that area and explore any interesting local structure.

Credits: ARM: Ming Ouh-Young, James J. Batter, P. Jerome Kilpatrick, William V. Wright, Russ Taylor, and Frederick P. Brooks, Jr. (P.I.)

Flythrough: Warren Robinett, Jim Chung, Bill Brown, David C. Richardson

Grants: ARM: NIH Grant #5-R24-RR 02170-07

Flythrough: DARPA #DAEA18-90-C-0044, NIH Grant #5-R24-RR 02170-07 and ONR Grant # N00014-86-K-0680

2. The Walkthrough Project

The Walkthrough project started in 1985 with the basic goal of "a virtual building environment, a system which simulates human experience with a building, without physically constructing the building." [1] Using this system, the user can "walk through" in real time a building that may not yet exist. As the user moves through the three-dimensional building model (by whichever means), perspective views are generated that make it seem as if the user is really inside the building.

Viewing can be done with a regular CRT, the large rear-projection screen, or a head-mounted display. With the HMD, the user's viewpoint is controlled naturally through his head movement: if he wants to see off to the left, he simply turns his head to the left. If the other displays are used, the head orientation and position are controlled by joysticks or the three-dimensional mouse.

Movement through the building can be controlled by walking on a treadmill, by the 3-D mouse or by another joystick. If the treadmill is used, the steps taken on it are translated one-to-one to steps in the virtual building. Turning while walking is accomplished by turning bicycle handlebars that are attached to the front of the treadmill.

When the joysticks are used for moving through the model, a technique called *adaptive refinement* [2] is applied to the image whenever the user stops (this can't be done with the HMD because there are always small head movements). This technique takes advantage of the reduced demand for interactivity and uses the compute time to improve the image realism and anti-aliasing.

This system was used to walk through the UNC computer science building before it was built and was actually used to make some design modifications before the construction began.

Credits: John Airey, John Rohlf, Randy Brown, Curtis Hill, John Alspaugh, and Amitabh Varshney, and Frederick P. Brooks, Jr (P.I.)
Grants: NSF Grant # CCR-8609588 and ONR Grant # N00014-86-K-0680.

3. Radiation Treatment Beam Placement

James Chung, a graduate student at UNC, is working on a CAD tool using a head-mounted display for designing radiotherapy treatment beam configurations in the hopes that it will aid radiotherapists in making better treatment plans [7].

In this VR system, a model of the patient's anatomy is explored by a doctor who is preparing a radiation therapy treatment plan in order to deliver a lethal dose of radiation to a tumor, while minimizing the exposure of healthy tissue to the radiation. The doctor puts on a head-mounted display and enters a virtual world containing a model of the patient's anatomy (as obtained by standard computed tomography (CT) methods and rendered as a set of polygons) and some number of polygonally defined radiation beams. While in this virtual world, the doctor experiments with different beam placements by moving the virtual beams by "grabbing" them with the 3-D mouse in order to find the optimal placement of the beams as described above.

The problem that this system addresses is that currently, radiation treatment planners have to look at the patient's anatomy on a two-dimensional screen, which makes it difficult to understand views other than the "cardinal" (orthogonal) views, so treatment geometries involving odd angles are not often used, even though they might result in a better overall treatment. The projected advantage of this system is that in the virtual world, the doctor is free to examine all angles of beam placement in a very natural manner, and should thus allow better treatment plans in less time.

Credits: James Chung, Julian Rosenman, Henry Fuchs and Stephen Pizer
Grants: DARPA #DAEA18-90-C-0044, NIH Grant #5-R24-RR-02170 and
ONR Grant # N00014-86-K-0680

4. X-Ray Vision

The X-Ray Vision project is a planned head-mounted display application where, instead of blocking out the view of the real world in favor of the computer-generated world, the computer-generated world is superimposed on the real world, using the aforementioned see-through head-mounted display.

There is no limit on what type of image can be superimposed, but I have chosen to focus on the problem of cranio-facial reconstruction (CFR) planning as my driving problem. In this application, computer-generated images of a patient's bony tissue (and possibly soft tissue) would be superimposed *on the real patient*. This would allow the surgeon planning reconstructive surgery to see the real soft tissue, yet have a three-dimensional image of the underlying bone at the same time. According to Davis [8], there is often a need in CFR planning to superimpose CT images of bony tissue onto CT images of soft tissue, and a tool that accomplished this in 3-D would be quite useful for planning.

In addition to the passive visualization, the system would have a marking capability, so that the surgeon could mark points in the computer image that were important in the planning process while with the patient.

Credits: Richard Holloway, Jefferson Davis
Grants: DARPA #DAEA18-90-C-0044, NIH Grant #5-R24-RR 02170-07 and
ONR Grant # N00014-86-K-0680

B. Personal Projects

1. Adventure Game

Another application developed at UNC to investigate navigation and action in virtual reality is an adventure game created by Warren Robinett.

This world consists of a network of rooms connected by portals. The spatial relationship between rooms is not necessarily one that is possible in normal three-space: ie., the room that you access through a portal to your right may also be accessible by going through the portal on the left and then up one level.

Movement in this world is accomplished by flying, as in the molecule fly-through described previously.

Each room has its own "thrills and chills": one room has an elevator that will take you up or down when you fly into it. Another room contains a mirror, which mimics your every movement. Still another contains a giant, hungry bird that chases you around the room and tries to eat you; however, if you're quick enough, you can "pick up" (again, with the 3-D mouse) a "vortex gun" and blast the bird with a swirling polygonal vortex.

This virtual world comes complete with sound effects: if the user bumps into a wall while trying to fly through a portal, a "thump" sound is generated by the Macintosh®. When the vortex gun is fired, it makes a "whoosh" sound, and if you hit the bird with the vortex, a buzzing sound confirms the hit.

2. A Sampling of Student Projects

Other personal projects have been done by students, both on their own and for class projects. The "Exploring Virtual Worlds" class at UNC has spawned many interesting virtual worlds. A few interesting examples of independent and class projects follow.

a) Virtual Piano

Bill Brown created a virtual piano, consisting of a model of a virtual eleven-key keyboard floating in the middle of the room. When the user puts on the VPL Eyephone® and the VPL DataGlove®, a large, animated white glove comes to life, and shadows his real hand. As he presses on a virtual key with his virtual hand, the key descends and a note comes floating up from the Macintosh®, corresponding to the key that was pressed.

b) Fly-through of a city

Ron Azuma and Ulrich Neumann created a virtual world with a city, tunnels, and a lake. The user can fly through the virtual world, experiencing visual, audio and force feedback. In the city, the user flies between buildings that range from small to skyscraper size through targets that give directions on where to go next. Collisions are detected between the user and objects in the city, and audio and force feedback (if the ARM is used for flying) make the user aware of the collision. In the lake, the user swims with fish that are animated in such a way that they avoid collisions with each other and objects in the lake.

c) Virtual Golf

Virtual miniature golf is the objective of Curtis Hill's virtual world. Here, the user is armed with a special putter (whose movements are tracked by the system) in order to sink a virtual golf ball into a virtual cup. When the user hits the golf ball with the putter, a putting sound is emitted to give audio confirmation of the hit. When the ball stops rolling, the user is transported to the point where the ball is so that he can continue play. When the user sinks the putt, the flag sticking out of the hole turns red and the sound of a ball falling into a cup is produced by the Macintosh.

d) Virtual Mountain Bike

Ryutarou Ohbuchi created a virtual mountain bike system which featured kinetic feedback of changing terrain through computer-generated resistance on the pedals, as well as real-time visual feedback. The large, rear-projection screen displays terrain data while the user pedals through the scene, while a kinetic resistance device changes the load on the pedals to reflect the slope of the path. Collision detection with virtual trees is also detected.

IV. Relevance To Art

While many of the aforementioned applications are scientific, their relationship to art does not require too great a leap of imagination. If this medium is ever to be useful for artists, then surely some advances made in the technology in pursuit of scientific goals will also benefit the artistic uses of it. For example, advances made in image generation facilitate better, faster images; improvements in tracking make the system more responsive and more usable; better auditory and haptic systems also increase the realism and the possibilities for interaction.

Also, when you look at it, the very nature of virtual reality is artistic: before we start, there is no virtual world- someone has to *create* it. The very nature of the medium is therefore creative, and, I assert, *artistic*. Much of the work in creating a VR application is in modeling the virtual world: choosing the colors, shapes and sizes of the objects in the virtual world, which seems to me to be an inherently artistic task. Moreover, since the creator is unbound by physical laws, his choices are freer (and thus harder!) than ever.

Finally, a note of realism. As great as the *promise* of virtual reality technology is, the reality is that it is an immature field and that many VR systems are still quite hard to use and prohibitively expensive. Except for the top-of-the-line flight simulators, most VR systems fall far short of actually fooling the user into believing he's in a virtual world. Of course, some will not be daunted by these failings and will use VR technology as is; but with continued hard work VR systems should continue to improve and be usable by more and more artists.

References

[1] Airey, J.M., Rohlf, J.H., Brooks, F.P. Jr. Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments. UNC Technical Report. Chapel Hill, NC. January, 1990.

[2] Bergman, L., Fuchs, H., Grant, E., Spach, S. Image rendering by Adaptive Refinement. Proceeding of SIGGRAPH '86. Appeared in *Computer Graphics*, August 1986.

[3] Brooks, Frederick P. Jr. Grasping Reality Through Illusion. Proceedings of CHI '88 (Washington DC, May 1-11, 1988).

[4] Brooks, F.P. Jr., Ouh-Young, M., Batter, J.J., Kilpatrick, P.J. Project GROPE- Haptic Displays for Scientific Visualization. Proceedings of SIGGRAPH '90 (Dallas TX, August 6-10, 1990). In Computer Graphics 24,4, August, 1990. 177-185.

[5] Christman, R.J. *Sensory Experience.* Harper & Row, New York, 1979.

[6] Chung, J.C., Harris, M.R., Brooks, F.P. Jr., Fuchs, H., Kelley, M.T., Hughes, J., Ouh-Young, M., Cheung, C., Holloway, R.L., Pique, M. Exploring Virtual Worlds with Head-Mounted Displays. Non-Holographic True 3-Dimensional Display Technologies, SPIE Proceedings, Vol. 1083, Los Angeles, CA, January 15-20, 1989.

[7] Chung, James C. Application of Head-Mounted Display to Radiotherapy Treatment Planning. Dissertation proposal draft. UNC, November 1990.

[8] Davis, Jefferson. Personal communication. May, 1990.

[9] Holloway, R.L. Head-Mounted Display Technical Report. UNC Technical Report. 1987.

[10] Ludel, Jacqueline. *Introduction to Sensory Processes.* W.H. Freeman and Company, San Francisco, 1970.

[11] Wang, J., Azuma, R., Bishop, G., Chi, V., Eyles, J., Fuchs, H. Tracking a Head-Mounted Display in a Room-Sized Environment with Head-Mounted Cameras. Proceedings: SPIE '90 Technical Symposium on Optical Engineering & Photonics in Aerospace Sensing, Orlando, Florida, April 1990.

[12] Wang, J., Chi, V., Fuchs, H. A Real-time Optical 3D Tracker for Head-Mounted Display Systems. Computer Graphics: 1990 Symposium on Interactive 3D Graphics, Vol. 24, No. 2, March 1990, pp. 205-215.

[13] Wang, J. A Real-time Optical 3D Tracker for Head-mounted Display Systems. PhD dissertation, UNC-Chapel Hill, March 1990.

[14] *Webster's New International Dictionary*, 2nd ed. Merriam: Springfield, MA.