


Virtual Reality

STATUTORILY EXEMPT



Editor's Note: This paper was awarded Second Prize in the Computer Hardware and Software Systems category of the 1992 CISI Essay Contest.

1. INTRODUCTION

Virtual reality represents both the hope and the hype of computer science. With its promise of providing highly interactive and responsive interfaces to computer applications, virtual reality suffers from a media blitz that threatens its success before it becomes reality. The term *virtual reality* conjures up notions of science fiction in many people's minds. Despite an abundance of hype, virtual reality is slowly emerging as an important new way of interacting with computer applications and simulations. Virtual reality is an attempt to provide an interactive, immersive interface to a synthetic, computer-generated world. It is the logical extension and coupling of recent progress in the fields of computer graphics and human-computer interaction (HCI). In the following discussions, we will attempt to separate the *fiction* of virtual reality from the *science*. In addition to providing a general introduction to virtual reality, we will focus on current technology and techniques as well as future possibilities. Discussions include history, display devices, interaction devices, design tools, and applications. Before delving into the details and technology of virtual reality, let us first consider what we mean by the term *virtual reality*.

1.1. What Is Virtual Reality?

Virtual reality is the effort to create an encompassing, simulated environment that provides a natural and intuitive way for humans to interact with computers. Such immersive virtual environments may perhaps best be described as "the integration of computer graphics and various input and display technologies to create the illusion of immersion in a computer generated reality" [BRY92a, 1.1]. Several terms – virtual environments, virtual worlds, VR, synthetic worlds, artificial reality, cyberspace, reality engines, and telepresence – are commonly used interchangeably to denote virtual reality. Although people sometimes create their own definitions for these terms, the above definition is generally accepted as an adequate description of virtual reality. We will use many of these terms throughout these discussions but will most often refer to virtual reality as simply *VR*.

Two primary characteristics of virtual environments are the associated display devices and human-computer interaction methods. The displays encompass the user in a scene viewed in three-dimensional space. As the user's head turns, the scene changes accordingly as in real life. VR interfaces imitate the way humans normally interact with the physical world. Aside from viewing in three dimensions, interactions typically include reaching out, grabbing, pointing, and touching. Some systems even include listening, speaking, and smelling. The virtual world, in turn, responds interactively and naturally to the user.

Some current applications of virtual reality include simulation of surgery for instructional purposes, a modeling system to manipulate molecules easily in 3D space, a VR-controlled robot arm, a building designed in virtual reality before it was built in physical reality, and various arcade games. Researchers are exploring and developing other uses of VR in such areas as medical imaging, data management, modeling, scientific hypothesis testing, and prototype design. Virtual reality is changing the way people work with computers to an intuitive, more direct method. However, despite all the media attention and hype, VR development is still in its early stages. Major hardware deficiencies exist, and integration of the various components is no trivial task.

1.2. A Brief History

The ideas behind virtual reality can be traced back to early flight simulators. Some might consider Morton Heilig's "Sensorama" as one of the first VR machines and many of the current trends in VR are derived from Myron Krueger's responsive environments. However, most researchers in the computer science field consider Ivan Sutherland's work at the University of Utah and MIT as the basis of virtual reality.

The responsive, interactive nature of virtual reality technology resembles various systems developed in the mid-1900s. A flight simulator displayed different airplane "out-the-cockpit" screens according to the actions of its "pilot" user. VR systems also respond to user actions. However, instead of being displayed on flat, two-dimensional screens as in flight simulators, VR images project into three-dimensional, virtual worlds. The Sensorama Simulator patented by Morton Heilig in 1962 might be considered one of the first VR prototypes. The Sensorama user sat in a wooden booth, peered into a large viewer, and gripped a set of handlebars. Sensorama took its "rider" on a motorcycle trip through Brooklyn, a California helicopter ride, a drive in a convertible with a young blonde companion, and to a personal belly dancer performance. Besides creating a visual journey, Sensorama provided appropriate smells, stereoscopic sounds, a lurching seat, vibrating handlebars, and "wind" blowing at its rider. Although the ideas behind Sensorama were essentially ignored when they were first introduced by Heilig, today's virtual reality systems closely resemble his multisensory environment. In the early 1970s, Myron Krueger began exploring Artificial Reality, a term he coined to describe "a new kind of environment, created by human perceptions triggered or mediated by video and computer technologies" [RHEI91, 122]. Although Krueger first explored many of the VR-related

research ideas currently being pursued, and some of today's VR systems resemble his responsive virtual room, virtual reality as commonly defined today developed primarily from the widespread enthusiasm for Ivan Sutherland's work in the late 1960s and early 1970s.

Ivan Sutherland and his colleagues at MIT's Lincoln Laboratory created the first head-mounted display (HMD) known as the "Sword of Damocles" [SUTH65]. This device was a bulky, metal display suspended from the ceiling and locked at the user's head. The Sword of Damocles displayed appropriate views according to the way its users moved their heads. Later, Sutherland, Daniel Vickers, and other researchers at the University of Utah developed the "first fully functional HMD system," which consisted of six interconnected subsystems: a clipping divider, a matrix multiplier, a vector generator, a headset, a head position sensor, and a general-purpose computer. This work was a natural extension of Sutherland's previous research called Sketchpad [SUTH63]. (Sketchpad is the seminal work on which all modern computer graphics systems are based.) Eventually, a wand called the "Sorcerer's Apprentice" was added to the system that allowed an HMD user to "reach out and 'touch' synthetic objects." The addition of this interactive device greatly increased the user's illusion of immersion and provided a method for system input.

Sutherland's innovative work provided a foundation for subsequent work in computer graphics and human-computer interaction. From the early 1960s to 1990, much research and development was carried out in these related fields. However, most of this work can be characterized as incremental advances. Computer graphics capabilities progressed from wireframe drawings to shaded displays to photorealistic rendering. Likewise, display technology progressed from vector-based storage tubes to high-resolution, color raster displays [FOLE90]. The increase in the capabilities of display devices and graphics systems, in turn, led to changes in the way we interacted with computers. New interaction methods and user interfaces evolved which took advantage of windowing systems, interactive pointing devices (i.e., mouse, trackball, joystick), and graphical user interfaces (GUI). Throughout this evolution, system costs continued to spiral downward which resulted in widespread use of graphics and GUIs.

While the central ideas of VR have existed since the 1960s, it was not until recent years that graphics hardware technology and interaction devices advanced sufficiently to allow the creation of immersive interfaces. Perhaps even more important was the overall increase in performance and storage capabilities of computer systems. Only recently have critical thresholds of technology and demand come together to provide the beginnings and feasibility of virtual reality. (For a complete history of the virtual reality movement, see Howard Rheingold's *Virtual Reality* [RHEI91]. Another very good introduction to related techniques and technology is [BRAN87].)

2. HUMAN-COMPUTER INTERACTION

Human-computer interaction (HCI) is a very complex and well-studied topic. Even with its detail and wealth of knowledge, there remain many unknown factors and many open areas of research. At its core, virtual reality is a high-bandwidth, highly interactive means for humans to interact with computer applications. Lessons learned from past HCI knowledge should be consulted when exploring new endeavors in VR. (Good introductions to HCI topics may be found in [SHNE87] and [RUBI84].)

Interaction methods may be broken down into three main styles: menus, command languages, and direct manipulation [SHNE87]. A menu presents a user with a list of choices. A command language allows complex communication based on a grammar or a set of language rules. Direct manipulation allows users to treat the interface as a collection of objects that may be manipulated, transformed, or controlled. Each of these interaction styles has tradeoffs such as ease-of-use, speed, error rates, and complexity; no single style is best in all situations. However, direct manipulation is currently a very popular method for human-computer interaction. Users like the feeling of control that direct manipulation interfaces provide. In addition, direct manipulation interfaces allow users to take advantage of their real-world experiences in object manipulation and visual feedback.

In considering this traditional taxonomy of HCI, virtual reality may be classified as a direct manipulation interaction method. In fact, virtual reality is the *ultimate* form of direct manipulation. The user is given the ability to directly manipulate all objects in the synthetic world, taking full advantage of his senses and perception skills. It is interesting to note that other interaction styles may be incorporated within VR. For example, some VR applications use virtual menus within their system. Virtual reality provides a high-bandwidth, interactive form of direct manipulation for communicating with computers.

3. VIRTUAL ENVIRONMENT SYSTEM COMPONENTS

The basic components of virtual reality systems are display devices, interaction devices, and a computational environment. (Here, we assume the computational environment incorporates an integrated graphics subsystem.) The display devices feature wide field-of-view and stereoscopic images that move with the user. Current VR displays include head-mounted displays, virtual "rooms," and binocular omni-oriented monitors. VR interaction devices allow natural, human movements and gestures as input into the system. Such devices may track actions of the user's hand, head, eye, and body. An overlap exists between VR display devices and VR interaction devices. Display devices inherently involve interaction, specifically system input through head movements. The following discussion on displays describes devices that present *visual* images to the user, while the section on interaction devices covers all forms of VR input and output devices except those previously discussed. The computational environment integrates the various

system components, calculates the current state of the system (scene), and draws the resulting display. VR systems typically use high-end graphics workstations for integrating this input and output. Currently, a Virtual Environment Operating System (VEOS) is also under development for supporting VR applications. The following sections describe the more popular devices on the market and those being developed at research institutions. [AIEX92] has a more complete list of commercial VR products.

3.1. Displays

VR display devices have the primary responsibility of creating a feeling of immersion in a simulated world. Various techniques used to create this illusion include stereoscopic images, wide field-of-view, and viewer-centered perspective. Stereoscopic images give the illusion of three dimensionality by displaying the correct perspective of a scene to each eye individually. The field-of-view measures the angle one can see without rotating one's head. A larger field-of-view creates a greater suspension of disbelief. For instance, pictures on a Sony Watchman (a hand-held television) dwarf in comparison to the realistic images created in large screen IMAX theaters. Viewer-centered perspective is the ability to display a scene from the appropriate point of view of the user as he rotates and moves. The smoothness and continuity of the changing images play an important part in the illusion of immersion and correlate directly with the speed that images are generated and displayed. Lag time between head movement and image transformations remains a primary problem with VR displays. Currently, head-mounted displays and virtual reality rooms are popular viewing devices. Other displays include binocular omni-oriented monitors and the visual retina display. Most of today's VR display devices consider only head movements, but some attempts have been made to incorporate eye movements as well.

3.1.1. Head Mounted Displays

Head-mounted displays (HMDs) are the most common visual interface to virtual environments. Figure 1 shows a head-mounted display, a helmet-like device with a pair of screens mounted in front of the user's eyes. The screens show separate images with slightly different perspectives for each eye in order to produce a three-dimensional effect. Most head-mounted displays currently use liquid crystal displays (LCDs); however, some systems use cathode ray tubes (CRTs). LCDs are flat and lightweight, making them suitable for head-mounted displays, but they have poor resolution and relatively low light output. CRTs are considerably more expensive and quite heavy, but they offer higher image quality, high light output, and small display areas. The fields-of-view for HMDs range from 100 to 140 degrees. (For comparison, a normal, nineteen-inch diagonal, flat CRT display screen offers a forty-five-degree field-of-view for a user located eighteen

inches away.) Jerky image displays and poor image resolution remain problems in HMD systems.




Fig. 1. Head-mounted display

Research laboratories and VR companies have produced various head-mounted displays. VPL Research produces the EyePhone, perhaps the leading commercial HMD. NASA and the Air Force Institute of Technology have produced LCD and inexpensive CRT prototypes. The Visually Coupled Airborne Systems Simulator (VCASS) at Armstrong Aerospace Medical Research Lab provides a quality, high-resolution CRT head-mounted display system. Honeywell has produced HMDs for the Apache attack helicopter, while GEC manufactures displays for British researchers and the F-16 night attack system "Falcon Eye." LEEP Systems Inc. markets Cybernetic 2, a LCD system optically designed to enhance image resolution. A see-through HMD that superimposes images onto real world objects is being developed at the University of North Carolina, Chapel Hill. Figure 2 shows a sketch of the UNC see-through HMD. Head-mounted displays are available from various vendors for use in custom VR systems. [BRY92a] offers a more extensive discussion on commercially available headmount systems and a technical description of display optics and resolution.




Fig. 2. See-through head-mounted display

3.1.2. Virtual Reality Rooms

Virtual reality "rooms" have large projection screens as the walls, ceiling, and floor. Figure 3 shows a model of the Audio Visual Experience Automatic Virtual Environment (CAVE), a VR room developed at the Electronic Visualization Laboratory at the University of Illinois at Chicago. The CAVE tracks the user's location and projects the appropriate images on the screens as the user moves throughout the virtual environment. The CAVE provides stereoscopic display, complete field-of-view, and viewer-centered perspective. However, computational lags cause noticeable display delays similar to those of head-mounted displays. To experience a three-dimensional effect, the CAVE user wears a pair of Stereographics Corporation's CrystalEyes, lightweight, electronic glasses that shutter asynchronously. The CAVE environment has the advantage of allowing multiple users simultaneously although only the perspective of one viewer may be displayed at a time. Developers are modifying the infrared signals controlling the CrystalEyes shutters in an effort to produce the correct perspective to all CAVE users. The CAVE and other VR rooms provide less restrictive, but more costly and slightly less immersive environments than head-mounted displays.




Fig. 3. Virtual reality room (CAVE)

3.1.3. Other Displays

Other types of visual displays include the binocular omni-oriented monitor (BOOM) and the visual retina display. Figure 4 shows a BOOM system, which was originally designed at NASA Ames Research Center and is currently produced by Fake Space Labs. The BOOM places small screens in front of the user's eyes similar to those in head-mounted displays; however, a BOOM is attached to a mechanical, arm-like structure. The user holds and moves the BOOM manually to change the views. The BOOM has a slightly smaller field-of-view (90 to 120 degrees), but the friction of moving the BOOM arm

prevents the lag problems common to HMDS and virtual rooms. The Human Interface Technology Lab (HIT Lab) at the University of Washington is experimenting with a visual retina display, where a filtered laser actually draws images on a user's retinas. (This obviously has important biological concerns and implications!) While the system is still in its earliest stages, such a laser microscanner would solve the current problems of poor image quality in virtual displays. The goal of virtual reality display devices is to create an easy-to-wear, easy-to-use visual interface.




Fig. 4. BOOM display

3.2. *Interaction Devices*

Interaction devices control the input and output to the virtual reality system. Most system input is tied to tracking devices that determine the position and location of the hand and head. Aside from visual output devices, some systems also include auditory, olfactory, and tactile sensory devices.

3.2.1. *Gloves*

Today's typical VR system consists of a head-mounted display connected to a glove input system. Gloves react according to the location, orientation, and sometimes gesture of a user's hand. Instead of typing commands at a keyboard or using a mouse to imitate two-dimensional motions, gloves allow users to communicate to a computer using more natural three-dimensional movements.

Glove systems have two responsibilities: computing hand position and orientation, and determining finger orientation (gesture). Two technologies used to track hand locations are orthogonal electromagnetic fields and ultrasonic signals. To determine finger angle and position, fiberoptic cables, Hall-effect sensors, and plastic strain gauges have been used to line gloves. The following discussion describes three popular glove systems--DataGlove, Dexterous Hand Master, and Power Glove -- each using a different combination of these technologies. Figure 5 shows photographs of the three gloves.



Fig. 5. Interaction gloves

VPL Research manufactures the DataGlove, perhaps the most popular VR glove system. The DataGlove uses orthogonal electromagnetic fields for hand tracking and fiberoptic cables for finger tracking. The Polhemus 3Space Tracker comes packaged with the DataGlove for determining hand location. The Tracker has a transmitter and a receiver, each with three orthogonal wire coils. A control unit pulses each of the coils in the transmitter with a radio frequency electromagnetic signal. The receiver measures the strength of the signal from each transmitter coil in each direction, a total of nine measurements, to determine the position and orientation of the receiver. The 3Space transmitter sits on a stand near the DataGlove user, while the receiver is mounted to the wrist of the DataGlove. The DataGlove consists of fiberoptic cables that run in loops along each finger and the thumb to determine finger position. Figure 6 shows a sketch of the cabling. The ends of each loop are connected to an interface board by the wrist. An LED sends light down one end of each cable, while a phototransistor converts the light into an electrical signal indicating the amount of light that passes through. As a finger bends, the fiberoptic cables along it allow light to escape. The light intensity along that cable indicates to the computer how much the finger is bent. The DataGlove provides a comfortable, fairly accurate input device to virtual reality systems.

EXOS produces the Dexterous Hand Master. The Hand Master is a lightweight, aluminum exoskeleton of sensors. The sensors are connected by hinged joints and held in place along each finger segment with pads and Velcro straps. Each sensor has a Hall-effect magnetic pickup and a magnet to measure the bending angle of the finger joint. The magnet moves towards or away from the Hall-effect sensor according to the bending and lateral motions of the finger. The sensor creates a voltage according to the strength of the magnetic field created by the magnet. PC software computes the finger position using the voltages of the various sensors. To determine hand position and orientation in space, the Hand Master must be attached to a electromagnetic tracking system, such as the




Fig. 6. DataGlove cabling

Polhemus 3Space Tracker (used in the DataGlove) or the Ascension Bird. Supposedly, the Hand Master is more comfortable than it appears [EGLO90]. Compared to the DataGlove, the Hand Master offers more precision in finger position measurements, but it also has a more complex structure and higher cost. One problem with magnetic trackers (as used in the Hand Master) is that they have limited range of precision in environments with metal objects. The metal in the environment interferes with the magnetic tracking process. Unfortunately, most computer labs (and hence environments of VR systems) have lots of metal objects and are prone to these problems.

Power Glove has an ultrasound system for hand-tracking and plastic strain gauges for finger measurements. Mattel produces the Power Glove for use with the Nintendo TV game system. Two transducers on the back of the glove send clicks to three receivers located around the monitor. An 8-bit processor determines the hand position and orientation from the data sent and received by the ultrasound system. Plastic strain gauges attached to each finger change in resistance according to the amount of bend. The processor determines the approximate position of each finger. [EGLO90] describes how the Power Glove may be directly connected to a PC compatible for experimenting with simple hand-tracking tasks. The Power Glove offers more durability but less precision than the DataGlove and the Hand Master. The cost of the Power Glove, however, is considerably less (about \$100, compared to \$8,800 and \$15,000, respectively).

3.2.2. Navigation

Numerous methods of movement in virtual reality systems have been implemented. For most applications, the user "flies" through the virtual environment using specific hand signals or pressing particular buttons to indicate speed and direction. Simple hand gestures such as one, two, or three-finger pointing may be mapped to specific navigation functions by the developer. One common approach is to use the index finger for indicating the direction of movement and the up and down motion of the thumb as a throttle/brake (see fig. 7). Some researchers have suggested mapping movement in VR to paradigms of vehicle control. The user uses a virtual vehicle to move within the virtual world. By leveraging off of people's current knowledge and experience with vehicle control (i.e., driving a car), vehicle control in VR may be a natural method of navigation. Initial prototypes of virtual vehicles used a simple flying platform. However, more sophisticated vehicles are quite possible.

When exploring complex virtual environments, users may be given a mechanism to enter another world or "room" in the scene. Such mechanisms that allow users to freely move to other environments, rooms, or scenes are called "portals." Portals function much like the transporter from the popular "Star Trek" television series by allowing users to instantly pass through a portal to another world of their choice. In implementation, portals may be represented by familiar real-world objects such as doors or windows or may be represented by some more exotic means. Since the developers of a VR system establish the navigation rules and are not restricted to real-world navigation rules, portals could be



Fig. 7. Navigation by gesturing

implemented in very creative ways. Portals provide a means of random-access movement much like that of a hypertext system.

Another method of movement requires the user to walk on a steerable treadmill. Researchers have also considered movement controlled by simulating bicycling, rowing, and driving. Note that users often need to be confined to stationary platforms or teathered so that they don't bump into objects in the "real world" while exploring the "virtual world"! An attempt to create a body glove similar to VR hand gloves has had limited success because of the extensive wire cabling. However, VPL Research does offer such a device called the Body Suit. The Spaceball by Spatial Data Systems offers six degrees-of-freedom force and torque control that resembles a joystick (see fig. 8). This device has proven useful for navigating in VR environments, but movements are not as natural as glove or other movement systems.

3.2.3. Touch, Hear, Speak, Smell

Other areas that are being explored for potential integration into VR systems are tactile and force feedback, spatially localized sound, and voice input. The development of some of these technologies has been motivated by applications not directly related to virtual reality. None are presently considered part of today's "typical" virtual reality system.




Fig. 8. Spaceball

Some research has been performed to incorporate tactile and force feedback into virtual reality glove systems. VR users would feel pressure on their fingertips when picking up virtual objects and resistance when trying to unite two repelling bodies. Three approaches have been used to provide tactile feedback to VR gloves [FOLE87]. One method derives from a feedback system for the blind which uses solenoids to control the amount of pressure of blunt wires against the skin. Another approach uses vibrating piezoelectric crystals activated by electric currents. Some developers are also trying to push "memory metals" against the skin by heating them with electric current and causing them to change shape. The addition of tactile and force feedback in VR systems can provide very convincing physiological cues that add much to the illusion of immersion and

the suspension of disbelief [BROO90] [IWAT90]. Figure 9 shows an experimental force feedback system. Methods of simulating texture and temperature have yet to be fully developed although early experiments are being performed. It appears that technology of tactile and force feedback for VR systems is still in its early stages.




Fig. 9. Experimental force feedback system

Spatially localized sound has a dramatic effect on the sense of immersion in virtual environments. However, creating the illusion of spatial location involves much more than just stereo headphones. In humans, the pinna (outer ear) modifies sounds based on their direction and amplitude. The brain determines the three-dimensional, spatial location of the sound by using the modifications performed by the pinnae and the differences in the signals received by each ear. The Concolvotron produced by Crystal River Engineering imitates the transformations of the pinnae to give sound a sense of three-dimensional location. A microphone placed in the ear canal measures the results of a test sound, creating a model for computing the actions of an individual's pinnae on sounds. The Convolvotron presents sounds to the user through earphones and may have up to eight binaural sound sources. In addition to this device, which tries to imitate the three-dimensional localizing capabilities of the human auditory system, there is also current research in automatically generating sounds from virtual environments and underlying simulations. A technique called "sound rendering" has been developed at The George Washington University for the automatic generation of sounds in animation systems and virtual worlds [TAKA92]. Sound is considered a very important aspect of VR since users

often get "lost" while moving in the virtual world. Many feel that sound will provide a good stimulus and feedback mechanism for navigation cues.

VR systems may eventually incorporate verbal input devices as yet another method of natural human interaction. Voice recognition systems are already being developed independently for use in word processing and military control input applications. At IBM's Thomas J. Watson Research Center, a system with a 20,000-word vocabulary (approximately 98 percent of the English speaking vocabulary) is being developed. In the late 1970s, the MIT Media Lab produced a voice-gesture command system called "Put That There" [HERO80]. The user viewed a wall-sized display of an ocean, pointing to ships and locations to relocate objects on the screen. The user could also scale, shape, manipulate, and move objects by speaking and pointing. Research in voice and speech recognition is a continuing effort. VR is one useful application of these developments.

The inclusion of smell in VR is often speculated. However, simulation of the sense of smell has several difficulties: how to generate various smells and how to eliminate remnants of previous scents. It is possible that smell may be integrated in VR in the future, but currently this remains unexplored.

3.3. Computational Environments

Virtual reality places great computational demands on the underlying platforms. High performance computers are needed to run simulations, read input devices, access and update scene databases, render scenes into images, and finally update displays. In addition, all of these tasks need to be performed at "interactive" speeds in order to provide a highly interactive and responsive system to the user. Unfortunately, very few (if any) of today's computers can satisfy such stringent requirements for complex VR applications. Often, design tradeoffs are made to compensate for machine performance deficiencies. For example, one might choose to sacrifice scene detail, advanced rendering techniques, or display resolution in order to provide tolerable display updates or consistent frame rates. It is generally accepted that 30 Hz refresh rate is needed for flicker-free animation and interaction [FOLE90]. However, current developers are often elated to "squeeze out" 10 Hz of interaction. Given less demanding or simplified applications, some machines can adequately support VR interaction. High-performance graphics and computation demands usually dictate the use of high-end workstations. Currently, Silicon Graphics, Inc. (SGI) is the machine of choice for developing virtual environments. SGI IRIS workstations boast very fast rendering rates for Gouraud and Phong-shaded polygons. The newest SGI platforms are capable of real-time texture mapping. Note that the heated battles of the workstation market cause companies to "leap-frog" one another with technology from month-to-month. Certainly, other vendors such as IBM and HP offer reasonable VR platforms as well. However, today the vast majority of VR research and development is being done using SGI workstations.

Because of the high performance demands placed on VR systems, single monolithic computers are often inadequate to address these problems. A combination of specialized hardware, supercomputers, parallel processing, distributed computing, and networks of heterogeneous computers is often needed to successfully implement virtual reality. Even in the most simple cases, two workstations are often used to provide a stereoscopic view. Each workstation computes a separate view for each eye to provide stereoscopy. Distributed systems seem like a good approach to VR implementation. However, issues include how to update a global scene database, how to replicate changes and object collisions across networked nodes, and how to insure simulation consistency throughout the network. All of these issues are very much open areas of research. Some VR applications use networked supercomputers to run demanding simulations with the graphical display, interaction, and integration work being performed on workstations. Such complex, heterogeneous computational environments add to the complexity and burden of VR development. To address such problems, there is ongoing research in developing specialized operating systems for supporting virtual reality in networked environments. The Virtual Environment Operating System (VEOS) is currently being developed at the Human Interface Technology Laboratory (HIT Lab) at the University of Washington. This operating system is designed to integrate the various components of a virtual reality system and to provide resource management. The philosophy of the HIT Lab is that others will successfully solve the problems with display and interaction devices; they want to solve VR development, management, and implementation problems that go beyond device inadequacies.

VEOS has been described as "software for the construction of, maintenance of, and interaction with arbitrary virtual environments" [BRIC92]. The VEOS project is a joint research effort, whose goal is to demonstrate potential methods of coordinating all the components of a VR system rather than to produce a marketable product. VEOS consists of three subsystems: the kernel, the entity interface system, and the interaction tools system. The kernel handles communication tasks, process management, and memory management. The entity interface system maintains the virtual environment's database using an object-oriented approach to define objects in the scene and the scene itself. The interaction tools system manages the user's input and output in the virtual environment. VEOS attempts to provide transparent low-level database, process, and communications management for the integration of virtual reality system components. While VEOS is a research effort, it is hoped that lessons learned from this project will eventually be incorporated into systems that better support virtual reality. This, in turn, will make VR applications easier to develop and to use.

4. DESIGN AND DEVELOPMENT TOOLS

VR systems and applications are very much in the early stages of development. Therefore, today's VR developers are faced with a very difficult task of integrating disparate hardware and software components. The current vendors of VR hardware

components usually provide no software for integrating these devices into a working system. Typically, VR interaction devices return bitstreams of data to be interpreted and utilized by custom programs written by the system developer. Although few companies currently offer extensive software development tools, several vendors have developed toolkits to aid in the development of applications. These toolkits usually consist of libraries of subroutines that may be called by a developer's custom code. In addition, some tools are packaged as complete, end-user virtual reality systems (including necessary hardware and software components). However, these "packaged" tools tend to be very limiting for custom application development. They are designed for end-users versus application developers. It should be noted that VR developers often spend a year or more of effort establishing a minimum suite of software and a minimal hardware environment before they even *begin* application development. The task of integrating VR components is a difficult and time-consuming process. A rich supply of VR development tools is needed to aid in these efforts. In this section, we will briefly discuss some current design and development tools. These tools are the WorldToolKit by Sense8 Corporation, Minimal Reality Toolkit by University of Alberta, Visual Simulation toolkit (Performer) by Silicon Graphics, Inc., Reality Built for Two (RB2) by VPL Research, and RS/UniVRS by IBM. Note that most of these development tools are very new and unproven; details are sketchy, and some have just recently become available. With these caveats noted, we provide a brief "snapshot" of current VR development tools.

4.1. WorldToolKit

The WorldToolKit produced by Sense8 Corporation offers an object-oriented library of over 230 "C" functions for developing three-dimensional graphics and virtual reality applications. WorldToolKit provides a flexible and easy method of integrating objects, lighting, animation sequences, input sensors, and graphics display devices through the use of a simulation manager, an object manager, and a real-time rendering pipeline. Some highlights of the WorldToolKit include real-time texturing, portal creation, terrain generation, interactive polygon color editing, data import from various formats, and device drivers for many popular interaction devices. WorldToolKit can best be classified as a programmer's development library.

4.2. Minimal Reality Toolkit

The Minimal Reality (MR) Toolkit, a project at the University of Alberta, is a library of subroutines for building virtual reality and three-dimensional user interfaces. The MR Toolkit supports several popular interaction devices including the Polhemus Isotrak and VPL's DataGlove and EyePhone. Features include the support of data and user interface distribution over multiple workstations, automatic mapping of actual world to virtual world geometry, and real-time performance analysis tool. The University of Alberta offers

the MR Toolkit as free software to academic and research institutions. The MR toolkit is also a programmer's development library.

4.3. Visual Simulation Toolkit (Performer)

Silicon Graphics Inc. (SGI) is developing a visual simulation toolkit for creating high-performance, IRIS-based simulation applications [JONE92]. The toolkit (known both as Visual Simulation Toolkit and Performer) comes packaged with a software library, sample applications, sample databases, and documentation. The software library is divided into three subgroups: the rapid rendering library, the image generation library, and the sample simulation applications. While developing the toolkit, SGI has focused on high performance, rapid development, rich functionality, and enhanced portability. Long term goals of the toolkit project include optimized performance, an easy-to-use development environment, and integrated support for high-function image-generation features. This software supports only SGI IRIS workstations.

4.4. Reality Built For Two

VPL Research Inc. markets Reality Built For Two (RB2), a VR development platform [BLAN90]. RB2 allows the developer to interactively change behavior constraints and interactions while the system is running. However, [BERG92] claims that RB2 is geared more towards nonprogrammers, and some developers may soon find themselves limited by its capabilities. The RB2 system consists of RB2 Swivel, Body Electric, a DataGlove 2 system, a Polhemus Isotrak, the EyePhone head-mounted display, and Isaac. RB2 runs on a Mac II workstation and two Silicon Graphics IRIS POWER Series workstations. Objects are designed with the solid modeling application called RB2 Swivel, which then provides the data to Body Electric and Isaac. Body Electric, a data flow/real-time animation control package, collects and interprets data from the various interaction devices and translates the data into graphics commands for Isaac. Isaac renders the images onto the monitors in the EyePhone. RB2 is an example of a "packaged" system geared towards end-users rather than custom development.

4.5. RS/UniVRS

IBM and Division Ltd. have developed RS/UniVRS, an IBM RISC System/6000-based development kit for creating virtual reality applications. The product is already on the market in United Kingdom and is expected in the United States soon. The RS/UniVRS package includes a workstation running AIX and the dVS/UniVRS development software. A primary advantage is that RS/UniVRS can act as a front end to any computer-aided design package for the RS/6000s. Little information is currently available on this forth-

coming product. This IBM RS6000-specific software is anticipated to be a programmer's library.

5. APPLICATIONS OF VIRTUAL REALITY

Virtual reality offers the potential for highly interactive interfaces to any computer-based simulation. Therefore, VR interfaces may be applied to a limitless range of applications. Existing applications of virtual reality involve such diverse fields as medicine, architecture, space science, visualization, and entertainment. From assisting in radiation therapy to walking through imaginary buildings, from playing with virtual molecules to allowing viewers to take part in movies, the potential applications of virtual reality are endless. However, before concluding the VR is the ultimate solution, one should realize that the current state of VR technology suffers from a number of problems. These problems include poor image quality, limited precision, computation delays, and limited display capabilities. In addition, cost and complexity can be overriding issues. (Section 6.0 offers more details on the obstacles to overcome in virtual reality research.) Despite these problems, VR is being used today. In the following discussions, we offer a brief synopsis of current VR applications. This list is by no means definitive. Rather, it is meant to show the diverse nature and applicability of VR.

5.1. *Medical Applications*

Researchers at the University of North Carolina at Chapel Hill are exploring the use of virtual reality in radiation therapy treatment and ultrasound imaging. Conventional methods of planning radiation therapy treatment involve using static, two-dimensional x-ray images. The UNC head-mounted display system allows radiotherapists to easily visualize the three-dimensional location of tumors within healthy tissue and design more effective treatment plans. Within the virtual world, the user grabs radiation beams from a storage "rack" and places them in the most beneficial positions to maximize the dosage to the tumor while minimizing the radiation to the healthy surrounding tissue [LEVO90a]. The Ultrasound Vision application transforms real-time ultrasound data into an image physically superimposed on the patient. Using a specially designed "see-through" head-mounted display system, the doctor can virtually "see" inside the patient [BAJU92]. Elsewhere, researchers are exploring the possibility of "Nintendo surgeons," doctors that learn surgical techniques by practicing on computer-generated versions of patients.

5.2. *Architectural Walkthroughs*

Architectural applications allow clients and architects to "walk through" a virtual model of a planned building. The three-dimensional VR model allows the architect to better visualize such aspects as lighting, acoustics, and physical stress, as well as to facilitate the understanding between the architect and client. At the University of North

Carolina at Chapel Hill, members of the Computer Science Department used their architectural walkthrough application to convince architects to remove a hallway partition that impeded traffic flow before the building was actually built. Matsushita Electric Industrial Company offers walkthroughs of remodeled kitchens to shoppers at a Tokyo department store. The coupling of VR interfaces with advanced real-time rendering techniques such as progressive radiosity make such virtual "walkthroughs" possible. Other VR explorers are experimenting with a system that allows architects to test the handicap access of their buildings designs from a virtual wheelchair.

5.3. *Space Sciences*

Virtual reality is of great interest to the space sciences. It is anticipated that VR interfaces to robotic systems will allow scientists to perform maintenance and repairs of spacecraft and conduct scientific experiments in space without leaving the confines of earth. This capability, known as "telepresence," will allow users to directly manipulate remote objects through robotics and virtual simulations. Figure 10 shows the telerobotic system being developed at NASA Ames Research Center. A VR robot located outside a space station simulates the actions of a system user inside the station. The robot's camera "eye" follows the user's head movements, displaying the appropriate images of the outside of the space station in the HMD. The user's gloved hand controls the movements and actions of a robot "arm." Other researchers are considering the possibility of using VR to explore Mars and other planets by remote control.

5.4. *Visualization*

One of the greatest strengths of virtual reality is its ability to model abstract and intangible entities in a concrete, easily manipulable, visual form. Scientific and mathematical problems commonly involve large amounts of complex data. Virtual environments allow users to intuitively manipulate and transform various graphical representations of data. An application developed by Sense8 Corporation presents a dynamic, real-time, three-dimensional model of the stock market [POTT92]. Small chips represent each stock, while the stock price and activity determine chip color and location. Blue squares are winners, while red ones represent falling stocks. Within a virtual world, a stock trader can instantly grasp the performance of various stocks by selecting chips with a wand and exploring stocks of interest by "flying" towards them using a Spaceball. The ability to buy and sell shares more quickly than other traders can lead to large financial profits. Similarly, intangible objects such as molecules, solar systems, and abstract data may also be examined in virtual reality. The molecular docking system being developed at UNC-Chapel Hill assists chemists in discovering new drugs by letting them handle virtual molecules that correctly fit into receptor proteins [IWAT90]. The system uses the Argonne Remote Manipulator (ARM), a six degrees-of-freedom, force-reflective feedback mechanism to control the movement and simulate the repulsive forces of the molecules (see fig. 11). Another UNC-Chapel Hill project allows users to explore a




Fig. 10. Space science application (telepresence)

true-to-scale version of the solar system, changing the scale, orientation, amount of detail, and time passage with a push of a button. A three-dimensional surface modeling application called *3 dm* allows users to intuitively build, modify, and understand their models [BUTT92].

At the Navy Research Lab (NRL) in Washington, D.C., virtual reality is being used to explore war gaming and electronic warfare simulations. The goal is to provide an immersive environment in which a general can experience various battle scenarios, get immediate feedback on command decisions, and visualize electronic emanations of various platforms (i.e., ships, planes, submarines, radars) [BERG92]. Electronic signals, radar networks, and vehicle movements are simulated within a virtual world. Using conventional VR gear (glove, HMD), a user enters the simulated battle and can interactively direct various vehicle movements and electronic warfare techniques. This simulation allows multiple participants, and decisions may be made by either a computer application or another person. In theory, users "inside" the environment do not know if they are engaging other people or simulations. Immersion within the virtual environment adds much realism and credence to traditional simulation interfaces.

Perhaps the most famous VR application in the scientific world is the Virtual Wind Tunnel. Conventional physical methods for visualizing fluid flow have limited capabilities and involve complicated techniques. Computer simulations of fluid flow require fast computation, high-speed rendering, high-resolution graphics, and massive data storage and retrieval capabilities--making virtual reality a suitable environment for fluid flow visualization. The Virtual Wind Tunnel places its user in a virtual flow field




Fig. 11. Molecular modeling in virtual reality

using a BOOM display system, and thus, allowing the user to easily examine the flow from any direction (fig 12). The user controls the speed, time, and scale of the flow using a glove input system. The Virtual Wind Tunnel is a good example of applying a virtual reality interface to visualization and understanding of complex physical phenomena such as fluid dynamics [BRY92b].

5.5 Entertainment

The entertainment industry has particular interest in virtual reality. Partially originating from science fiction, virtual reality has great potential in bringing new forms of entertainment to the general public. W Industries of Great Britain has already begun manufacturing various VR arcade games. Virtual World Entertainments has a Battle Tech Center in Chicago. Sixteen pilots divided into two teams control their individual "machine of destruction" attacking members of the opposing team. Other possibilities include interactively participating in VR-based "movies," taking a realistic preview of a vacation site, and creating new forms of art. The entertainment industry offers huge financial rewards for those willing to invest resources and capital in VR-based entertainment. Some preliminary glimpses of VR entertainment are offered by Nintendo game systems with higher visibility coming in the near future (watch at toy stores near

you!). The theme of a recent science fiction film called *The Lawn Mower Man* revolved around virtual reality. All too often, the entertainment application of VR receives much publicity with other (arguably) more useful and more realistic applications being overshadowed. Consequently, this hype masks and hinders subsequent progress and acceptance of VR. However, the entertainment application of VR is also viewed as very positive. It is the mass-market appeal of VR entertainment that will drive software and hardware components into low-cost affordable systems. But perhaps the most important factor is *it's just plain fun!*




Fig. 12. The virtual wind tunnel

5.6. CAVE Applications

Most of the applications mentioned above involve using head-mounted display systems; however, CAVE (VR rooms) applications also exist [CRUZ92]. Some applications developed for the CAVE include regional-scale weather display, brain surgery planning, the Visible Embryo, the Evolving Universe, the Snowstorm, Fractal Exploratorium, and Bio Modeling. The weather application creates a three-dimensional display of weather systems over a region. In planning brain surgery, the CAVE transfers MR data into three-dimensional anatomical models projected physically onto patients. CAVE users may take

a trip through a human fetus in the Visible Embryo or through the galaxies and stars of the Evolving Universe. The Snowstorm displays three-dimensional vector fields using interactive particle systems, while Bio Modeling generates interactive models of biological macromolecules. The Fractal Exploratorium allows users to explore fractals and chaotic attractors.

6. OBSTACLES TO OVERCOME

Creating a virtual world is not a simple endeavor. Although various VR system components may be purchased from commercial vendors, the integration of these devices involves a great programming effort. The development of useful applications has been plagued by performance/lag problems and poor image resolution.

Vendors offer a vast array of display devices, interaction devices, computational environments, and even application development software for virtual reality systems. Choosing among all the different devices can be a time-consuming process; integrating all the components is even more difficult. No step-by-step instruction manual exists for setting up a system, and there are no fixed guidelines for designing and developing applications. Further, the VR quest for complete immersion requires high-speed rendering, fast computation involving complex equations, and unique system constraints. Aside from the task of meeting such requirements, VR software usually must be written in low-level languages to avoid the unacceptably slow algorithms and cumbersome data structures of high-level languages. With the youth of the VR industry and the relatively small number of working applications, most efforts to develop applications should be viewed as experiments rather than solutions.

In addition to the complexities of development, VR applications are also limited by certain hardware and software characteristics of current VR devices. The two major obstacles in creating an illusion of immersion are the lag problem and poor image quality. The lag problem refers to the slight delay between the time when a user's head moves and the time when the rest of the virtual world moves to the user's new perspective. With each motion of the user, the system must quickly update the world model and produce the appropriate outputs to the appropriate interaction devices. The lag results from the delays inherent in tracking devices, the delay in processing the data produced by interaction devices, and the computations necessary to maintain the model of the virtual environment and display the appropriate perspective. Several approaches to the lag problem are being examined at UNC-Chapel Hill. One approach sacrifices detail for smooth, natural motion. As the user moves, simplified images transform smoothly to the correct perspective. When the user pauses, detail is reinserted into the images. Another form of this progressive refinement approach uses more detailed rendering in the specific gaze of the user. That is, a tracker measures the directed gaze of the user's eyes and more detail is shown in the center of the image (the gaze position) with less detail at the image edge [LEVO90b]. The Pixel-Planes project is an attempt to speed up computations by linking thousands of custom processing chips to function as a single computer [FUCH89]. Methods of

improving the performance of tracking devices are also being examined. The problem of poor image quality results from size and space limitations of current display devices. The clarity of displays has been measured to be that of a person with visual acuity between 20/85 for BOOM displays and 20/425 for head-mounted displays [CRUZ92]. (The visual acuity measurements are Snellen fractions, meaning the viewer can see at twenty feet what a viewer with average eyesight can see at the number of feet indicated by the second number.) Today's display devices exhibit such poor image quality that users are legally blind! Efforts to improve image quality include manipulating various optical properties, distorting rendered images, and developing fundamentally different display devices. These system deficiencies are primarily technological problems and should disappear with more research and improved devices.

VR applications are also limited by the computation speed of the underlying simulation as well as the rendering performance of the graphics system. To a certain extent, we will always want more performance than we currently have. However, within recent years substantial improvements in computation and rendering capabilities have made "interesting" VR applications possible. Classic rendering "tricks" such as texture mapping can be used to compensate for inadequate rendering speeds. Texture mapping allows applications to "fake" the level of detail in scenes by simply mapping 2D images or textures on 3D objects [HECK86]. Some high-end workstations (e.g., SGI Reality Engine) offer real-time texture mapping, which is very useful for developing convincing VR applications. Clearly, improvements in computation and rendering speed will add much to our ability to develop sophisticated virtual worlds.


Given these obstacles and problems, what types of applications are best suited for current virtual reality technology? Many researchers feel that any application involving three or more dimensions that can conform to current limitations will more likely benefit from an immersive interface compared to a conventional interface. Applications that do not require precise control, complex graphical display, or large data accesses are potentially good candidates. VR interfaces seem to offer more realistic depth cues, better use of human visual perception capabilities, and more intuitive interaction methods. However, much research and empirical evidence is needed to support and prove these conjectures. As the current obstacles to VR are reduced, better applications and uses of VR will emerge. Consequently, better understanding of VR interaction methods and suitability will be possible.

7. CONCLUSIONS

Virtual reality is an attempt to provide an interactive, immersive interface to a synthetic, computer-generated world. This new method of human-computer interaction provides a natural, intuitive way of interacting with computer applications. VR provides direct manipulation and control of synthetic objects as well as providing the user with the illusion of immersion. In the preceding discussions, we have seen that the notion of virtual reality is a relatively old concept that can be directly traced to Sutherland's historic

Sketchpad system. However, only in recent years have graphics and computation performance sufficiently improved and display and interaction devices matured to warrant enthusiastic pursuit of VR. Enthusiasm for VR has reached a critical level such that an overabundance of media hype threatens further advances. Laymen are beginning to expect *science fiction* capabilities and not *scientific* capabilities. Today's display devices suffer from poor resolution and poor image quality, and current interaction and tracking devices have significant lag problems. Even though these system components are plagued by problems, great improvements are being made to the point that useful applications are emerging. Successful VR applications include applications from medicine, architectural design, space sciences, and visualization. Probably the most touted VR application is the virtual wind tunnel that was implemented at NASA Ames.

In the upcoming years, research and development in VR will be directed at navigation issues, software development tools, improved display devices and interaction devices, operating systems that better support VR, and interaction methods. It is widely believed that virtual reality interfaces offer a better and more intuitive means of interacting with computers. However, much investigation needs to be done to fully understand and support such assumptions. Surely, some applications are not well suited for immersive interfaces. Further research is needed to find out what applications are and are not suitable for VR. Much interest in virtual reality is being expressed by industry, academia, and government. Many institutions that have a vested interest in computer science have begun preliminary investigations into VR and its applicability to their problems. Most major institutions have established virtual reality research labs to investigate these possibilities first-hand. Despite current limitations and hype, virtual reality is emerging as a very exciting and innovative tool for highly interactive computer interfaces. Perhaps in the near future, the promise and the fiction of virtual reality will become fact.



REFERENCES

- [AIEX92] "Virtual Reality Resource Guide," *AI Expert*, August 1992, 42-47.
- [BAJU92] Bajura, Michael, Henry Fuchs, and Ryutarou Ohbuchi. "Merging Virtual Objects with the Real World: Seeing Ultrasound Imagery within the Patient," *Computer Graphics (Proc. SIGGRAPH)*, Vol. 26, No. 2, July 1992, 203-10.
- [BERG92] Bergen, Daria, and Rudy Darken. "Getting Started in Virtual Reality," Navy Research Laboratory and The George Washington University, SGI User's Group Meeting Presentation, Washington, D.C., August 1992.
- [BLAN90] Blanchard, Charles, S. Burgess, Y. Harvill, J. Lanier, and A. Lasko. "Reality Build for Two: A Virtual Reality Tool," *1990 Symposium on Interactive 3D Graphics*, ACM SIGGRAPH, March 1990.
- [BRAN87] Brand, Stewart. *The Media Lab: Inventing the Future at M.I.T.*, New York: Penguin Books, 1987.
- [BRIC92] Bricken, William. "VEOS Design Goals," *Implementation of Immersive Virtual Environments*, course notes no. 9, ACM SIGGRAPH '92, July 1992.
- [BROO90] Brooks, Frederick P., Jr., Ming Ouh-Young, James J. Batter, and P. Jerome Kilpatrick. "Project GROPE - Haptic Displays for Scientific Visualization," *Computer Graphics (Proc. SIGGRAPH)*, Vol. 24, No. 4, August 1990, 177-85.
- [BRY92a] Bryson, Steve. "Survey of Virtual Environment Technologies and Techniques," *Implementation of Immersive Virtual Environments*, course notes no. 9, ACM SIGGRAPH '92, July 1992.
- [BRY92b] Bryson, Steve, and Creon Levit. "The Virtual Wind Tunnel," *IEEE Computer Graphics and Applications*, July 1992, 25-34.
- [BUTT92] Butterworth, Jeff; Andrew Davidson, Stephen Hench, and T. Marc Olano. "3DM: A Three Dimensional Modeler Using a Head-Mounted Display," *1992 Symposium on Interactive 3D Graphics*, ACM SIGGRAPH, 135-38.
- [CRUZ92] Cruz-Neira, Carolina, Daniel J. Sandin, Thomas A. DeFanti, Robert V. Kenyon, and John C. Hart. "The Cave Audio Visual Experience Automatic Virtual Environment," *Communications of the ACM*, Vol. 35, No. 6, June 1992, 65-72.
- [FOLE87] Foley, James D. "Interfaces for Advanced Computing," *Scientific American*, October 1987, 127-35.

- [FOLE90] Foley, James D., Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics, Principles and Practice*, Addison-Wesley Publishing Company, 1990.
- [FUCH89] Fuchs, H., J. Poulton, J. Eyles, T. Greer, J. Goldfeather, D. Ellsworth, S. Molnar, G. Turk, B. Tebbs, and L. Isreal. "Pixel-Panes 5: A Heterogeneous Multiprocessor Graphics System Using Processor-Enhanced Memories," *Computer Graphics (Proc. SIGGRAPH)*, Vol. 23, No. 3, July 1989, 79-88.
- [HECK86] Heckbert, Paul S. "Survey of Texture Mapping," *IEEE Computer Graphics and Applications*, November 1986, 56-67.
- [HERO80] Herot, C., R. Carling, M. Friedell, and D. Kramlich. "A Prototype Spatial Data Management System," *Computer Graphics (Proc. SIGGRAPH)*, Vol. 14, 1980.
- [IWAT90] Iwata, Hiroo. "Artificial Reality with Force-feedback: Development of Desktop Virtual Space with Compact Master Manipulator," *Computer Graphics (Proc. SIGGRAPH)*, Vol. 24, No. 4, August 1990, 165-70.
- [JONE92] Jones, Michael. "A High Performance Visual Simulation Toolkit," Silicon Graphics Inc. technical report/white paper, 23 March 1992.
- [LEVO90a] Levoy, Marc. "A Hybrid Ray Tracer for Rendering Polygon and Volume Data," *IEEE Computer Graphics and Applications*, March 1990, 33-40.
- [LEVO90b] Levoy, Marc, and Ross Whitaker. "Gaze-Directed Volume Rendering," *1990 Symposium on Interactive 3D Graphics*, ACM SIGGRAPH, March 1990.
- [POTT92] Potts, Mark. "Virtual Reality: Sci-Fi Technology on the Verge of a Billion-Dollar Boom," *The Washington Post*, 16 August 1992, H1.
- [RHEI91] Rheingold, Howard. *Virtual Reality*, New York: Summit Books, 1991.
- [RUBI84] Rubinstein, Richard, and Harry Hersh. *The Human Factor: Designing Computer Systems for People*, Digital Press, 1984.
- [SHNE87] Shneiderman, Ben. *Designing the User-Interface: Strategies for Effective Human-Computer Interaction*, Addison-Wesley Publishing Company, 1987.
- [SUTH63] Sutherland, Ivan E. "Sketchpad: A Man-Machine Graphical Communication System," *Proceedings of the AFIPS Spring Joint Computer Conference*, Vol. 23, 1963.
- [SUTH65] Sutherland, Ivan E. "The Ultimate Display," *Proceedings of the IFIPS Congress 2*, 1965.
- [TAKA92] Takala, Tapio, and Jame Hahn. "Sound Rendering," *Computer Graphics (Proc. SIGGRAPH)*, Vol. 26, No. 2, July 1992, 211-20.