Virtual Reality - A General Overview

Michael E. Bauer
Sterling Software, Palo Alto, CA, U.S.A.

Abstract

During the past few years, Virtual Reality (VR) has transitioned from a subject for basic research to a concept that is being employed in many scientific and engineering disciplines (e.g., architecture, computational fluid dynamics, planetary exploration, medical research, etc.). A major reason for this transition is that VR technology has been teamed with advanced computer graphics technology to create new scientific visualization tools, which give scientists and engineers the opportunity to immerse themselves in their data and afford them a new and effective perspective of this data.

This paper gives an overview of VR systems by focusing on VR hardware technologies as well as design approaches of successful virtual worlds. The concepts in this paper were obtained from the 1993 SIGGRAPH course on Applied Virtual Reality, "Virtual Reality Overview", held in Anahaim, California, and taught by Carolina Cruz-Neira of the Electronic Visualization Laboratory, University of Illinois at Chicago.

Introduction

What is virtual reality? There are many different factors that characterize the definition of virtual reality. The following characteristics are generally accepted throughout the industry:

- Virtual reality has to perform its function in real time; the computations associated with a changing viewpoint as well as the changes in the model have to take place in real time. Any significant lag time will have a detrimental effect on the VR sensation.
- Virtual reality space exists in its own 3D world which may be either abstract or concrete. The scenes themselves are limited only by the bounds of one's imagination.
- Virtual reality should provide the illusion of immersion. A participant should feel the sensation of being completely surrounded by the virtual scene.
- Virtual reality should provide and encourage interaction. A participant should be able to interact with the virtual world via some input devices (e.g., gesture recognition, voice recognition, applying forces, etc.). Objects that are manipulated by these input devices should respond with some sort of reaction.
- Virtual reality should provide feedback to the participant. This may be accomplished by either providing visual, sound, or force feedbacks.

To produce an effective virtual reality system, many of the above mentioned characteristics should be realized. The following discussions on VR hardware technologies and VR design considerations should assist a user to accomplish this task.

VR Hardware Technologies

1) A virtual reality system requires many different hardware components, which control the participant's sensory inputs (e.g., visual, audio, tactile, etc.), track the participant's movements, and receives and interprets the participant's input into the system. The purpose as well as advantages and disadvantages of each class of device is addressed in this section.

1.1) Visual Displays

The visual display of the VR system is the most important component from the participant's perspective. It provides the visual stimuli, which determines the degree of immersion that can be achieved. Currently, there are many devices that can function in this capacity; however, they can be grouped into four categories: monitors, head mounted display systems (HMDs), booms, and projection systems.

1.1.1) Most monitors employed as the visual display include some form of stereographics support. Typically, they are the least expensive option for a visual display because they come standard on many graphics workstations or they are an inexpensive option. There are two different stereo viewing systems available for monitors (reference #1). One system consists of stereo eye glasses which contain liquid crystal electronic shutter lenses which, in turn, select the appropriate image for the appropriate eye. The light weight glasses are linked to an infrared emitter which controls the shutter frequency. The other system uses a screen-sized liquid crystal electronic shutter that presents different polarized states for left and right eye images from the CRT . The shutter is controlled directly through the monitor. The user is required to wear only light weight polarized glasses. Due to the larger liquid crystals, this system is usually more expensive.

The advantage of using monitors is that the stereoscopic scene is displayed with a full workstation resolution requiring only the space of a desktop. It is a minimal-invasive system, i.e., light weight glasses need to be worn. In addition, small groups of people may experience the virtual scene simultaneously from the same monitor. Simple interaction with external physical interfaces is maintained. The primary weakness of a monitor system is that it provides the least degree of immersion. This is due to the small angle of view of the scene, and a non-surrounding view as well. There is no superimposition of the real world and the virtual space; essentially the virtual scene is inside of a box.

1.1.2) Head Mounted Display Systems (HMDs)

HMDs use a pair of liquid crystal displays (LCDs) as their screens which are placed directly in front of the users' eyes. In addition the HMDs use special optics to attain a wide field of view. The combination of the wide field optics and the screen's close proximity to the eyes, produces a large angle of view as well as a panoramic view of the virtual scene. Since they are very light weight and flat they can be easily mounted on a users head. This provides for many degrees of freedom in moving about in the real world allowing for good navigation through large databases . The degree of immersion is currently higher than monitors. The potential for increased immersion, once some of the disadvantages listed below are solved, is one of the main advantages of HMDs.

The disadvantages of HMDs can be summarized as follows. The LCDs have a very low resolution compared to a monitor, especially when they are in full color. The LCDs display quality is weak due to problems with contrast and brightness. Since they are worn on a user's head they can be highly invasive. The HMDs cannot be shared at the same time unless the same video signal is sent to multiple HMDs. Without half-silvered mirrors as part of the screen, HMDs provide no superimposition of the real and virtual world, i.e., a user is visually separated from the real world. For this reason physical interfaces as well as hand and body parts need to be reconstructed in the virtual environment. Nevertheless, HMDs are one of the most widely used visual displays in virtual reality systems.

1.1.3) Booms

Booms primarily use monitors (CRTs) as the display technology. The booms' display technology differs from the monitors display technology, described above, in that booms use two small monitors, one for each eye, rather than one larger stand alone monitor. CRTs provide a higher resolution than LCDs and have a better display quality. The presence of high-power, high- frequency electromagnetic fields and a heavy weight make CRTs unsuited for mounting on a user's head. For this reason they are mounted on a boom comprised of several arms connected via linkages. As with LCDs, the CRTs incorporate wide field optics to increase the field of view. Since booms are not physically mounted to the head, they have the advantage of being only moderately invasive. A user can easily place or push away the CRT from the eyes. This allows for tolerable interaction with physical interfaces. Due to the simplistic design of the boom and the many degrees of freedom it provides, navigation through large databases is good.

The boom's disadvantages are similar to those of the HMDs. Only one person can realize the virtual experience at a time, unless the video signal is shared with other devices. Once the CRT is in front of the eyes, physical interfaces and body parts need to be reproduced in the virtual world, as with HMDs, there is no superimposition of the real and virtual worlds.

1.1.4) Projection Systems

With projection systems the virtual scene is projected on large screens which completely surround the view. An example of such a projection system is the CAVE (Audio Visual Experience Automatic Virtual Environment) developed at the Electronic Visualization Laboratory at the University of Illinois (reference #2). The projection systems have several advantages. A user may have a 360-degree top and bottom view of the virtual scene. Since the stereoscopic images are projected on surrounding screens occupying a large space, many people can share the experience simply by wearing polarized glasses. The projection systems are the least invasive of all the display systems since only a tracker is connected to the body. The high video resolution, panoramic view provides a strong sense of immersion. This makes good navigation of databases possible. The use of physical interfaces is very easy since there is full superimposition of the virtual and real spaces. In addition there is no need to regenerate the hand and body parts.

The projection systems do have some disadvantages. Projection systems are required to be placed in a large room due to the large surrounding screens. Supporting the synchronization and generation of the stereoscopic images for all the screens is very difficult and expensive. Since there is full superimposition of the virtual and real spaces, stereo violations may occur. As an example, the user's hands may have the same spatial location of a computer generated solid object, hence appearing to be inside of the object.

1.2) Tracking Systems

Current tracking systems are used to measure the user's head position and orientation as well as the position of other body parts such as the hands and fingers. Tracking the head position and orientation allows for correct computation of the view of the world from the users point of view. This is critical to obtaining a good illusion of immersion. Tracking the other body parts allows for interaction and control of the virtual world. Several methods employed in tracking systems (electromagnetic, mechanical, acoustical, optical, and inertial systems) that will be discussed next.

1.2.1) Electromagnetic Tracking Systems

This type of tracking system uses a transmitter or source which emits electromagnetic fields along three orthogonal axes. The position and orientation of one or more sensors with respect to the source are reported. Examples of such tracking systems are the Polhemus Isotrack as well as the Ascension Flock of Birds. Electromagnetic tracking systems have some problems. One problem is the interference from other electromagnetic fields. Trying to track the head position using a HMD, which may produce its own electromagnetic field, could result in erroneous information if care is not taken.

The primary problem with electromagnetic tracking systems is the latency and accuracy issues. Latency is associated with the elapsed time between change in position of a sensor and the time required for the change to be reported. To have effective tracking, the latency should be less than 0.1 seconds. Accuracy is extremely sensitive to the presence of metal and other electromagnetic fields as mentioned above. The advantage of the electromagnetic tracking systems is that the sensors may be moved around freely. In addition, object obstructions (objects placed between the source and sensors) do not interfere with the electromagnetic signal. Electromagnetic Tracking Systems are fairly inexpensive compared to other tracking systems.

1.2.2) Mechanical Tracking Systems

These tracking systems are usually associated with mechanical linkages. A good example of such a device is the boom discussed above. The position and orientation of the user's view point are computed from joint angles. This technique produces very small latency and very high accuracy. The problems associated with the boom are the range restriction due to the mechanical linkages. The price associated with accurate mechanical tracking systems can be very high.

1.2.3) Optical Tracking Systems

This tracking system, which is still in an experimental stage, uses light emitting diodes (LEDs), video cameras, and image processing techniques to determine the position and orientation of a user. Fixed cameras may either track the LEDs mounted on the user or head mounted cameras track fixed LEDs in the surrounding space. The positions of the LEDs on the video screens can be established using signal processing techniques. These positions are then used to determine the positions of the LEDs in the actual 3D space. This technique requires a clear line of sight between the LEDs and the cameras, hence object obstructions are critical. Optical tracking systems which are still in the experimental stages, are rather expensive due to the electronic equipment required.

1.2.4) Acoustic Tracking Systems

Acoustic tracking systems use ultrasonic sound which is picked up by several microphones arranged in a triangular fashion. The position and orientation are calculated by the difference in elapsed time that is required for the sound pulses to reach the microphones. This type of tracking system has the same latency problem as magnetic trackers; the acoustic tracking systems, however, are not affected by metal interference. Since sound waves are affected by the media through which they travel, object obstructions will affect the sound pulses. Examples of acoustic tracking systems are Logitech Mouse and Mattell Powerglove, which are less expensive tracking systems.

1.2.5) Inertial Tracking Systems

There are two types of inertial tracking systems: gyroscopes and accelerometers. Both of these systems are in experimental stages. Gyroscopes measure the angles of pitch, yaw, and roll. This system is based on the principle of conservation of momentum. Accelerometers detect the acceleration of a tracked object, and, through integration, the position can be determined. The advantage of inertial tracking systems is they do not require a separate source, so their range is limited only by the length of their cord.

1.3) Computation Systems

The computer hardware is the "heart" of a virtual reality system. The computer controls all input and output devices, determines the state of the environment as well as the stereoscopic viewpoints associated with the left and right views of the virtual scene. Last, but not least, the computer also renders all of the objects of the virtual world. All of these functions have to be synchronized on a frame by frame basis at a video rate no less than ten frames per second (i.e., the graphics update rate, which includes the computation time, time to read the input devices, and time to produce the stereo rendering, is greater than ten frames per second). To create an effective virtual reality system, a multi-processor computer is often used, where each processor is assigned one of the following tasks: handling the operating system, processing the computations, performing the rendering, and controlling the input and output devices. Some VR systems employ a dual headed graphics system (two independent graphics pipelines) or a pair of graphics workstations working in tandem to produce the stereographic image pairs. If the virtual scene requires complex computations, these can be performed on a separate supercomputer. The computed data as well as the state of the virtual environment can be passed to and from the workstation and supercomputer via a fast network. Suitable networks may be either the UltraNet or HIPPI networks.

1.4) Haptic Devices

Haptic devices allow a user to manually explore the virtual environment. By measuring the position, orientation, and forces of the user's body parts, virtual objects may be easily manipulated. Force and torque feedback as well as tactile devices provide a user with the sensation that the objects are real. In addition, devices that can produce the stimulus of temperature highlight the illusion of reality.

1.5) Audio Systems

Audio systems can provide excellent interaction with the virtual environment. Voice recognition for input commands and voice as well as sound output synthesizers create feedback about the virtual world and are powerful interaction tools. The output synthesizers may provide two different sound cues: monaural, and spatially localized sounds. Monaural sounds can be used to indicate events or user feedback on actions such as glove gesture recognition, when objects have been grabbed, or display states of the environment. Spatially localized sounds are used to give individual sounds a spatial location. Assigning these sounds to objects in the virtual environment can simplify the task of localizing objects.

Virtual Environment Design Considerations

2.0) There are many issues to consider when designing an effective virtual reality system. They include immersion, presence, combining real world objects with virtual objects, degree of intrusiveness of the interface, physical and audio feedback, navigation and control, as well as successive refinement or level of detail, which will be discussed below.

2.1) Immersion

An important aspect of a VR environment is that the participant feels completely surrounded by or immersed in the virtual world. This virtual world may be generated dynamically by a computer graphics system or scripted from pre computed images on a video disc player. The key to immersion is in the quality and quantity of visual cues provided to the participant. Currently many of the human factors issues involved with creating a good sense of immersion are still unknown. This is a subject of ongoing research at the NASA Ames Research Center. From experience, researchers have learned that obtaining a good sense of immersion depends on the performance of the virtual reality system. As a rule of thumb, and as mentioned earlier, the graphics update rate should be greater than ten frames per second. Furthermore, the latency between the participant's actions and the instances at which these actions are presented in the scene should be less than one tenth of a second.

2.2) Presence

Presence is the sensation that objects in the VR scene are "really out there in space". Two properties dominate in the field of VR to provide computer graphics scenes with a sense of presence. These are depth cues in the virtual scene as well as behavior and interactivity of the virtual objects. There are several parameters that produce stimulating depth cues. Perspective rendering simulates the view as we see it in the real world. Objects that are farther away appear to be smaller, two parallel lines converge to a point at the horizon. Stereoscopy creates a strong sense of depth perception. Two stereoscopic images of the virtual scene are created by rendering the images from two different perspective points of view. The disparity between the two images is interpreted by the brain to create the sensation of depth. It should be noted that stereo is most effective when objects are within a ten meter proximity. Objects that are farther away cease to provide stereo depth cues since the disparity between the two objects is significantly diminished.

Head motion parallax produces strong depth cues as well. As the viewpoint of the scene changes, the brain receives images from different viewpoints at infinitesimal time intervals, thus inducing perceived disparity. In addition, objects that are closer hide the view of objects that are farther away. Changing the location of light sources to highlight features of objects provides additional depth cues. Depth cues can also be obtained by introducing shadows into the scene. Real shadows are very expensive to render; their approximations can, however, be rendered at lesser cost while still providing good depth information. Atmospheric blurring is used to create the impression that objects are far off in the distance.

Providing interaction with an object, such as touching an object and having it respond to the person's actions, certainly creates a true sense of presence. Assigning a behavior to the objects, particularly in the response to a participant's interactions, as well as objects possessing their own behavior and the capability of self-generated motion also increases the sense of presence.

2.3) Combining Real World Objects with Virtual Objects

When designing a virtual reality system, it should be determined in advance whether a user needs to operate specialized devices in the real world. It may be necessary to see the devices as well. In this case there are several options. The display device could provide complete superimposition of the real and virtual worlds, such as the projection system. As an alternative, half-silvered mirrors could be used in a head mounted display to see the real and virtual worlds simultaneously. Using a desktop monitor allows a user to have full interaction of the real world while sacrificing the sense of full immersion. To maintain the full sense of immersion, the operating devices could be reproduced in the virtual world to facilitate interaction with virtual operating devices. In this case, the user's body parts would need to be reproduced as well.

2.4) Degree of Intrusiveness of the Interface

The user's interaction that is required with the virtual world affects the selection of the user interface. Some user interfaces require a participant to wear encumbering devices and attachments which create unacceptable motion restrictions. It may be necessary to minimize intrusion by sacrificing other features such as full immersion.

2.5) Physical Feedback

Currently physical feedback is created by using force/torque controllers based on sensor technology used to measure the forces and torques to control robotics. This technique significantly increases the level of immersion Unfortunately the current understanding of physical feedback is still very limited. There are two major areas of research in physical feedback. These can be described as force and tactile feedback.

Force feedback relates to the forces acting on the muscles and joints of the participant. Caution must be exercised in providing reasonable force feedback to guarantee the participant's safety.

Tactile feedback can provide the sensation of touch, i.e., forces acting on the participant's skin as well as the sensation of temperature.

2.6) Audio Feedback

There are several methodologies to incorporate audio feedback into a virtual environment. As discussed above in the section on audio systems, there are monaural sounds and spatially localized sounds. Audio feedback of spatially localized sounds may be used as a navigational tool by using audio as a clue to orient users in the virtual world. It may also be used as a localization tool where sounds are associated with objects or a particular location in the virtual space. The monaural sounds can be used as a sonification tool, thus transforming numerical information into sound. In addition monaural sounds are useful as an interaction tool, using audio for input and output such as voice recognition and speech synthesis, or providing sound feedback for actions. Voice commands used for environment control can free the user's hands for more essential tasks such as object manipulation.

2.7) Navigation and Control

Navigation within the virtual world is an important aspect of the virtual experience and it has an affect on the participant's immersion. Of the many different navigation techniques implemented for virtual environments, the most common is to "point and fly" through the virtual space. With this method a participant sets the desired direction of travel via an input device such as a mouse or the DataGlove (produced by VPL Research). Additional interaction with the input device activates the "flying" motion. By scaling the size of the virtual world a user can easily travel through a macroscopic landscape or the microscopic world of atoms. Tele-transport is also a very useful navigation method. By stating the location (e.g., next room) via voice control for instance, a user can automatically be tele-transported to the specified virtual location. A graphical interface may be used by displaying two or three dimensional maps in the virtual world to provide the capability of selecting a location on the map and instantly placing the participant at the respective location in the virtual world. Other exotic navigation devices such as bicycling on a stationary bicycle, or walking on a steerable treadmill, where both actions occur in the real world, have been successfully employed. In any case, if a participant is allowed to walk around freely in the real world without a clue of object obstructions (e.g., bumping into walls, tripping over cables), collision control must be provided.

Manipulating the parameters which control the state of the environment, such as changing the light settings, or manipulating an object's position and orientation in the environment, require other user interfaces. Pull down menus or 3D widgets in the virtual space have proven useful.

2.8) Successive Refinement or Level of Image Detail

Scenes in a virtual reality environment may be very complex and require much time to render. Such complex images may prevent the virtual environment from running within the minimal graphics update rate of ten frames per second. In this case, the use of successive refinement (images are rendered in greater detail only when there is no motion), is recommended. As soon as the participant moves his viewpoint, a much lower resolution image is displayed to guarantee real-time update rates.

Minimizing the complexity of the image to assure high update rates can be accomplished in several ways as discussed below.

2.8.1) Visibility Culling

A virtual reality scene may be comprised of a large virtual space containing many virtual objects. At any point in time a user may only see a fraction of this entire space in the viewing frustum (3D scene displayed to the user). There is no need to process the data (pass any of this data through the graphics pipeline), that is outside of the viewing frustum. This selective processing is called visibility culling. Visibility culling may also be extended to objects that are in the viewing frustum and are hidden behind other objects, e.g., objects hidden behind a hill or building. There are several papers published that cover visibility culling extensively (reference #3-#7).

2.8.2) Texture Mapping

Texture mapping can provide rich scenery with complex textures on simple geometry. With texture mapping any complex 2D digitized image, e.g., a tree, may be mapped onto a simple geometrical object such as a polygon. To render a tree with polygons would require much computational time. However, rendering the 2D image of the tree mapped on simple polygons can be accomplished in a fraction of the time using hardware texture mapping available in most medium- and high-end graphics workstations.

Texture mapping can also be used to animate a sequence in a virtual world, e.g., displaying an explosion or flames. Presenting this type of animation in a virtual environment using standard rendering techniques (e.g., particle animation and polygonal rendering) would be very difficult to accomplish within the desired frame rate. A good solution to this problem is to display a series of preprocessed animated textures of the animation on a simple polygon.

2.8.3) Level of Detail (LOD) Modeling

The concept of level of detail modeling suggests that objects only be modeled at their necessary resolution. For instance, objects that are closer to the viewer should be rendered with more detail than objects that are further away. There are many different schemes to provide LOD modeling, some of which are very specific to the virtual scene. A more general scheme is described as follows. LOD models of an object may be switched based on the displayed object's size. The closer a particular object is to the viewer, the larger it appears to be. This size appearance can be used to switch to different LOD models. This technique has serious problems in that it introduces visual "popping", i.e., when the LOD changes a visual artifact is produced. One method to reduce popping is to blend both models together for a period of time. This however, does not altogether eliminate visual popping. A more appropriate method to produce the best visual effect, is to incorporate a morphing technique (reference #8, #9, #10). This technique takes the vertices of one LOD model and displaces them over time to coincide in position with corresponding vertices on the next higher LOD model, with color and attributes similarly interpolated. A specific LOD management tool (reference #11) was used at NASA Ames Research Center to accommodate the large databases associated with planetary terrain data.

2.8.4) Load Management

The objective with load management is to maintain an even load balance on the VR rendering system. In designing a virtual scene, one should reduce heterogeneously distributed scenery. Frequently this is very difficult to do, due to the nature of the scene. To maintain the desired frame rate, one could set high and low threshold values as a percentage of this frame rate. If the frame rate exceeds the high threshold value, the system is stressed and a lower LOD model of the scene would be displayed. Once the frame rate drops below the low threshold value, a higher LOD model could replace the lower LOD model. Such a technique is discussed in reference #11.

2.9) Shared Virtual Environments

Shared virtual environments can play a significant role in virtual environments since multiple users may share and control a VR environment; scientist may exchange knowledge and ideas using shared data. Generally a shared virtual environment has a host machine which is a single computational platform. It maintains the state of the environment as well as providing a large storage space. Graphics workstations which handle all of the graphics and communicate with the host machine via a high speed network (UltraNet or HIPPI). Additional information on shared virtual environments is available in the upcoming paper written by Steve Bryson.

Summary

Virtual Reality is evolving to provide a new paradigm of how we view and interact with data. It is a serious tool that can now be applied to many disciplines such as scientific and information visualization, architecture, engineering, design and medical applications, remote presence, learning and training, leisure, sports, health, telecommunications, entertainment, and art and advertising. Several books (reference #12-#14) are available on VR to obtain detailed information on VR systems. Currently the state of the art of VR is: "it is almost here"; it suffers from slow frame rates and lag time, poor display technology, poor registration with real world, tedious model-building, and various problems with the interface technology. With many of the technological developments in the hardware being fueled by the entertainment industry in VR, many of the limitations associated with current VR systems will be surmounted within the next few years and VR will become an overwhelming experience as well as a powerful tool to visualize and interact with information.

Acknowledgments

The author thanks Carolina Cruz-Neira for organizing and teaching (amongst other speakers) the course on Applied Virtual Reality taught at SIGGRAPH '93, Anaheim, California. The material presented at the course provided the foundation for this overview paper.

The author is very thankful to Randy Kaemmerer, Glenn Deardorff, Margaret Bradshaw and Eric Hibbard for commenting on the drafts, which resulted in an improvement of the paper in many aspects.

References

1) E. Hibbard, M. Bauer, M. Bradshaw, G. Deardorff, K. Hu, and D. Whitney, "On t he Theory and Application of Stereographics in Scientific Visualization," *Eurographics '91, State of the Art Reports*, Vienna, Austria, September, 1991.
2) C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, "Surround-Screen Projected-Based Virtual Reality: The Design and Implementation of the Cave," *Computer Graphics*, pp. 135-42, August 1993. Proceeding of SIGGRAPH '93, held in Anaheim, CA, 1-6 August, 1993
3) N. Greene, M. Kass, G. Miller, "Hierarchical Z-Buffer Visibility," *Computer Graphics*, pp. 231-38, August 1993. Proceeding of SIGGRAPH '93, held in Anaheim, CA, 1-6 August, 1993
4) S. J. Teller and Pat Hanrahan, "Global Visibility Algorithms for Illumination Computations," *Computer Graphics*, pp. 239-46, August 1993. Proceeding of SIGGRAPH '93, held in Anaheim, CA, 1-6 August, 1993
5) S. J. Teller and C. H. Sequin, "Visibility Processing For Interactive Walkthroughs," *Computer Graphics*, vol. 25, no. 4, pp. 461-69, July 1991. Proceeding of SIGGRAPH '91, held in Las Vegas, NV, 28 July - 2 August, 1991
6) T. A. Funkhouser, C. H. Sequin, and S. J. Teller, "Management of Large Amounts of Data in Interactive Building Walkthroughs," *1992 Symposium on Interactive 3D Graphics, special issue of Computer Graphics*, pp. 11-20, ACM SIGGRAPH, Cambridge, MA, March, 1992
7) H. Fuchs, Z. M. Kedem, and B. F. Naylor, "On Visible Surface Generation by A Priori Tree Structures," *Computer Graphics*, vol. 14, no. 3, pp. 124-133, July 1980. Proceedings of SIGGRAPH '80.
8) T. Beier, S. Neely, "Feature-Based Image Metamorphosis," *Computer Graphics*, pp. 35-42, July 1992. Proceeding of SIGGRAPH '92, held in Chicago, IL, 26-31 July, 1992

9) J. F. Hughes, "Scheduled Fourier Volume Morphing," *Computer Graphics*, pp. 43-45, July 1992. Proceeding of SIGGRAPH '92, held in Chicago, IL, 26-31 July, 1992
10) J. R. Kent, W. E. Carlson, and R. E. Parent, "Shape Transformation for Polyhedral Objects," *Computer Graphics*, pp. 47-54, July 1992. Proceeding of SIGGRAPH '92, held in Chicago, IL, 26-31 July, 1992
11) L. E. Hitchner, and M. W. McGreevy, "Methods for user-based reduction of model complexity for Virtual Planetary Exploration", IS&T/SPIE's Symposium on Electronic Imaging, San Jose, Feb. 1993
12) Rheingold, H., *Virtual Reality*, Summit, New York, 1991
13) Krueger, M., *Artificial Reality II,* Addison and Wesley, 1991
14) Pimentel, K, Teixeira, K., *Virtual Reality: Through the new looking glass*, McGraw-Hill, 1993