VIRTUAL REALITY

Virtual Reality (VR) is the use of computer graphics and other technologies to create a simulated environment in which the user interacts. While computer graphics supplies the visual component of the experience via a normal screen, a stereoscopic display, or head-mounted technology; a variety of other devices can be used to allow the user to experience the interaction. For example, a head tracking device can allow the scene to change when the user moves their head, a force feedback joystick can be used when a pilot is flying a VR simulator, or data gloves can be worn to touch and move objects in the virtual world (see Figure 1). The key is that the user has a sense of being immersed in a virtual world in which they are engaged.

Figure 1. Virtual reality devices.

History

It is hard to pinpoint a starting point for VR, since it is not always clear when an application should be considered a virtual experience. For example, Ed Link, an inventor who was enthralled with aviation, saw a need for training pilots in a safe environment. He patented the mechanical “pilot trainer” simulator in 1929 to provide future pilots with the feel of how to control an airplane. Link Trainers have been used in the military since 1934 when the U.S. Army Air Corps purchased six trainers to produce instrument capable pilots. While this original device did not contain a visual component, it did create a virtual sense of flying an aircraft. In 1939/1940, A. E. Travis added a visual component using film loops that reacted to the pilot’s control of the instruments. By 1972, General Electric had created computerized flight simulators with multiple large screens surrounding the cockpit. It was not until the late 1970’s that Tom Furness began experimenting with putting virtual displays in cockpit simulators through the Visually Coupled Airborne Systems Simulator (VCASS), a specially designed pilot’s helmet with displays.

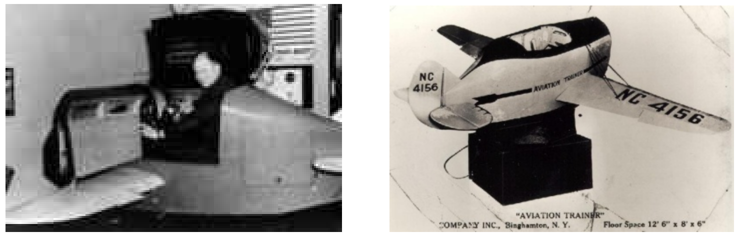

Figure 2. Ed Link's flight simulator, 1929.

Another early mechanical device that attempted to create a full sensory immersive experience was the *Sensorama* prototyped and patented in 1962 by Morton Hielig (see Figure 3). The Sensorama was able to display stereoscopic 3D images in a wide-angle view, provide body tilting, supply stereo sound, and triggered wind and aromas during a film. While the device was heralded as the "Cinema of the Future," it never received sufficient financial backing to become commercialized and remained a technical curiosity.

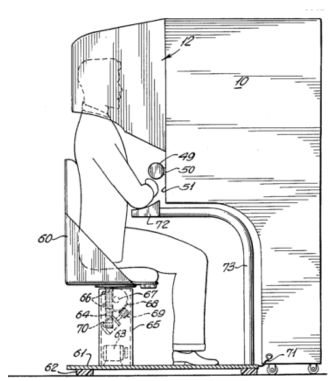

Figure 3. Morton Heilig's Sensorama*.*

Ivan Sutherland began experimenting with the first head-mounted 3D graphics display in 1966 while at MIT's Lincoln Laboratory (Figure 4). While bulky and cumbersome, the device included the capability for tracking head movement and allowed the user to investigate a 3D cube floating in space. By 1970, Sutherland created the first fully functional head-mounted display. Also in the 1970's, Fred Brooks began experimenting with tactile feedback devices at the University of North Carolina.

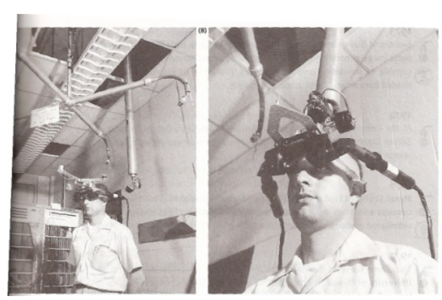

Figure 4. Sutherland's head-mounted display.

In recent years, tremendous increases in computing power, vastly improved displays, and huge storage capacities have increased the advancement of virtual techniques and reduced the cost of VR-type devices such as head-mounted displays and data gloves. Real-time interaction with realistic displays is now possible. Looking forward, there is currently research looking at effective use of other sensory inputs such as smell. Others are investigating better VR technology, including a rather remarkable claim by Sony who, in 2005, filed for and received a patent for the idea of the non-invasive beaming of different frequencies and patterns of ultrasonic waves directly into the brain to recreate all five senses.

Virtual Reality Applications

Trainers and simulators, as described earlier, were some of the first applications for VR. Besides the obvious safety advantage of learning a skill in a non-threatening environment, simulators can create situations that would be difficult, if not impossible, to experience otherwise. In addition, the cost of training can be significantly reduced in a virtual environment. VR also enhances learning by allowing students to replay certain scenarios, slow down action, back up, and experiment with cause-effect relationships. VR simulators have been developed for everything from operating cars, planes, and spacecraft to virtual surgery allowing doctors to practice complex operative tasks prior to entering the operating room (Figure 5).

Figure 5. Performing virtual surgery.

Another well-known application of VR is in the area of computer games. By the 1990's, VR gaming platforms were developed and being used for a series of commercial games. Racing, aerial dog fighting, sports, rock band accompanying, and first-person shooter games can all be conducted in a virtual environment. An interesting aspect of certain games is the social context they create. Games, such as *Second Life,* create a 3D virtual world in which players construct *avatars* for themselves and interact with other players playing different roles in world. The concept has been extended to the point that some universities are offering classes in the virtual world for players to attend via their avatar. Jumping on the VR entertainment bandwagon, and not known for doing things small, Disney has created *DisneyQuest,* a virtual reality interactive theme park with five floors of virtual games, 3D encounters, and VR experiences based on Disney characters, animations, and movies.

Scientific data visualization has been another application making use of virtual reality. Complex 3D data from scientific simulations and measurements can be difficult to effectively view and understand on traditional displays. By allowing the user to become immersed in the data for exploration, patterns and relationships can become more apparent. One of the pioneering technologies in this field was the Cave (Cave Automated Virtual Environment [recursive definition]) developed at the University of Illinois at Chicago in the early 1990's. The Cave uses 3-6 rear projectors in a room composed of screens that surround the viewer(s) (Figure 6). Images are projected on the walls of the Cave so that the viewer is immersed in the virtual world created from the images. Some type of interactive device is used to allow the user to navigate through the space. A 3D stereo Cave projects stereoscopic images and the user wears 3D goggles that allow objects to appear suspended in space such that the viewer can walk around them. Caves are currently being used at several universities and research centers for visualization and immersion research. The biggest downsides of a Cave are the size (basically a room within a room) and the cost of high quality projectors and screens.

Figure 6. CAVE virtual reality display.

Another application of VR technology has been the field of architecture. Traditional methods for designing buildings and spaces involve the use of 2D plans which may provide an accurate representation of a space, but are hard for people to visually translate into what the final spatial relationships will look like. VR allows a space to be designed in a virtual environment that people can physically walk through and see what the space will look like in actual 3D once constructed. The ability to visually experience a set of rooms in 3D can provide a much more realistic understanding of the look and feel of the overall space.

Many other VR applications are being developed as the technology improves and becomes more affordable. We are becoming closer to creating some of the futuristic artificial environments depicted by movies such as *Brainstorm, Lawnmower Man,* and *Avatar*.

Virtual Reality Technology

The difference between a virtual reality experience and a normal interactive computer graphics application is that VR creates a sensation that the user is a part of the virtual environment. A high quality display that immerses the viewer in the environment can help create this sense of reality. The Cave, described earlier, is one means of creating this immersion. A lower cost, and significantly more mobile, approach is to use some type of *head mounted display* (HMD), an example of which, the Sony Glasstron, is shown in Figure 7. The basic concept of an HMD is to use small displays in front of each eye. The advantage of this is that it can block out most or all of the surrounding visual distractions. A HMD parameter to consider when purchasing one is the Field of View (FOV). Humans have around 160°

FOV horizontally and 135° vertically, although the area of focus is much smaller. A small display screen in an HMD may offer only 25-40° FOV; 40° is equivalent to viewing a 105 inch movie screen from 12 feet away. Another factor to consider is whether the device supports stereo vision. Because each eye has a separate display, it is possible to display stereoscopic images to enhance the sense of 3D.

Figure 7. Head mounted display.

An interesting variation of the HMD is the Virtual Retinal Display, or VRD. VRD, also known as Retinal Scan Display (RSD) or retinal projector, is similar to HMD in appearance and can be worn on the head (Figure 8) or viewed from a desktop or handheld device. The head-mounted version, rather than having small displays for each eye, uses lasers or focused LEDs that are projected into each eye painting the picture in a raster format directly on the retina. This is not to be confused with the VISOR (Visual Instrument and Sensory Organ Replacement) worn by Lieutenant Commander Geordi La Forge on *Star Trek: The Next Generation* (the fictional VISOR interfaced directly with the optic nerve). VRD technology has been under development since the late 1980's. The image appears suspended in space in front of the viewer and provides a high resolution, full-color, screen-less display. The level of light entering the eye is small enough, both with laser and LED versions, to be well within established safety limits for human vision. Several applications have been proposed for VRD technology besides fully immersive virtual reality. For example, a mobile device, such as a phone could use VRD to display a full-size monitor when viewed by the user. Another possible application is to show the image into only one eye to provide auxiliary information for the real scene being viewed with the other eye, such as a repair manual or schematic. This type of application is sometimes referred to as *augmented reality* in which the computer is used to provide additional information and benefit to the real world versus a virtual environment.

Figure 8. Virtual Retinal Display.

In addition to seeing the virtual environment, it is necessary for the viewer to be able to change the view and navigate around for exploring the world. Several approaches can be used to provide this new viewing position and direction including handheld devices such as a mouse, wand, or dataglove. A more realistic approach for viewing the scene from a given position is for the scene to react to the viewer's head movement. That is, when the viewer turns their head, the scene changes accordingly. One means of accomplishing this is to use some type of viewer head tracking sensor so that the image being displayed is in sync with the location, position, and viewing direction of the user. Experiments have shown that head tracking is significantly more effective for accomplishing virtual tasks than other navigation methods. Different technologies are used to perform head tracking. The viewer may wear a device, perhaps as part of an HMD, that signals its location and orientation. Another approach uses advanced video processing techniques that capture the user's face from a camera mounted on the display device. The video is processed in real-time to locate the user's face and direction of viewing. Eye tracking devices can also be used to monitor the direction that a viewer is looking. Once the viewer's position and orientation are known, the software has to update the displayed view quickly. Even a small time lag between movement and view change significantly decreases the effect of realism.

The dataglove, or wired glove, is a device that allows for interaction with the virtual environment being displayed (Figure 9). Recently, the wired glove was popularized by the film *Minority Report.* Sensors in the glove use motion tracking to determine the glove's position and orientation in space. Additional sensors can track other gestures such as the bending of the fingers. This can be used by the software for interactions such as grasping an object. Expensive high-end wired gloves can include *haptic feedback* capability which simulates the sense of touch when the user interacts with a virtual solid object. An alternative to a wired glove for gesture interaction is to use a camera and vision processing to locate the user's hand and its orientation. While sacrificing the ability for tactile feedback, it is a low cost alternative that can be used to capture hand gestures.

Figure 9. Wired glove for interaction.

Haptic, or force feedback, devices allow users to interact with the virtual environment via the sense of touch. In addition to the wired gloves described above, simple haptic devices such as game joysticks and steering wheels can provide the user feedback based on the current simulated situation. For example, a steering wheel becomes harder to move as the user accelerates through a turn. Mobile devices, such as phones, use haptic feedback for touch screen interaction. Styluses, or hand held devices, can also provide haptic feedback. For example, in virtual surgery, the surgeon can wield a force feedback scalpel that provides different levels of resistance based on whether skin, muscle, or different organs are being incised.

Virtual Reality and the Web

Many vendors claim to offer a virtual experience (in its broadest possible definition) on the web. This can mean anything from displaying and interacting with 3D images to interactions with large numbers of other participants in virtual worlds. Some web applications use virtual devices such as HMD's and wired gloves. Many applications are limited to non-immersive desktop VR allowing users to create or participate in virtual worlds with limited interactions. Some of the limitations of using the web for virtual applications are bandwidth, efficient representation of virtual objects, and standardization. Bandwidth and efficient representation is critical for realistically rendered images to be manipulated and rendered in real time. Standardization applies to model representations, browser display capability, and commonly defined interactions. Different standards, such as VRML and X3D, have been proposed to create a common 3D representation for browsers to work with. Currently, most web-based VR applications require special plug-ins for modern browsers to handle.

Despite the limitations, many pundits project VR will become widespread on the web with many applications. Virtual marketing, in which a user can review and experience 3D products, has huge commercial potential. Virtual real estate (already partially available) would allow buyers to walk through a house listed on the web without having to physically visit it. Collaborative environments will allow geographically separated individuals to interact together in virtual 3D worlds for designing new products, conducting research, and jointly producing creative works. And, of course, social networking and gaming in a true 3D interactive VR world will have tremendous impacts.

Mixed and Augmented Reality

While virtual reality deals with simulated environments, mixed and augmented reality (AR) focuses on combining computer-aided interaction and information with the real world. A common and simple example is the yellow first-down line seen in televised football games. The line is automatically drawn by computers based on a 3D model of the field and camera positions and superimposed on the live action. Military applications include the *heads up display* that can combine targeting and instrument information with reality on an aircraft windshield or helmet display. AR applications typically require some type of camera input of the current scene being viewed along with tracking sensors, so the viewer's position and viewing direction is known. Large computing power is required to perform real-time image processing on the camera stream along with processing whatever application is desired. Possible future applications of AR include virtual screens providing information on whatever object is being viewed, virtual conferences (think "holodeck"), location-based navigation, and virtual road signs.

LAGNIAPPE – Avatars

It is common in VR applications for individuals to create an artificial representation of themselves for interacting in the virtual world, known as an *avatar.* The word comes from the Sanskrit word *avatāra* which, in Hinduism, refers to the appearance of a deity in physical form, having descended from heaven to earth. In computing, an avatar can refer to a 3D virtual character, a 2D icon used for applications such as chat rooms, or even a screen name. The name was first coined in 1985 by the developers of Lucasfilm's online role-playing game, *Habitat.* The use of avatar to mean online virtual bodies was popularized by Neal Stephenson in his novel *Snow Crash* (1992). In his book*,* the term was used to describe the virtual simulation of the human form in a fictional virtual-reality application on the Internet. Today, many games allow users to design their own avatars to reflect their own image, or characteristics they want to project. Avatars may be humanoid, creature-based, or mythical. There is a growing second industry devoted to the creation of products and items for avatars in games such as *Second Life.*

In 1994, the concept of avatars became one of the motivations for a proposed movie by James Cameron who was completing *Titanic* at the time. The proposed movie, *Avatar,* was to use virtual actors and, according to Cameron, was inspired by "every single science fiction book I read as a kid." The proposed movie was estimated to cost $100 million to produce. Although a script was created, filming and production did not start and many felt the movie would never be made. In 2006, Cameron announced that the original film was not produced because he was waiting for technology to improve, and that he was now creating it for a summer release in 2008. A new script was written and new technologies, including an innovative camera system, were created to produce the final film in 3D. Estimates put the film's production cost at $300 million. Some critics claimed the cost was excessive and predicted the film would be a financial bust. The film was released in December 2009 and grossed $232,180,000 after three days despite a major snowstorm on the east coast. Cameron says if the film is successful, he will produce two follow-on movies.