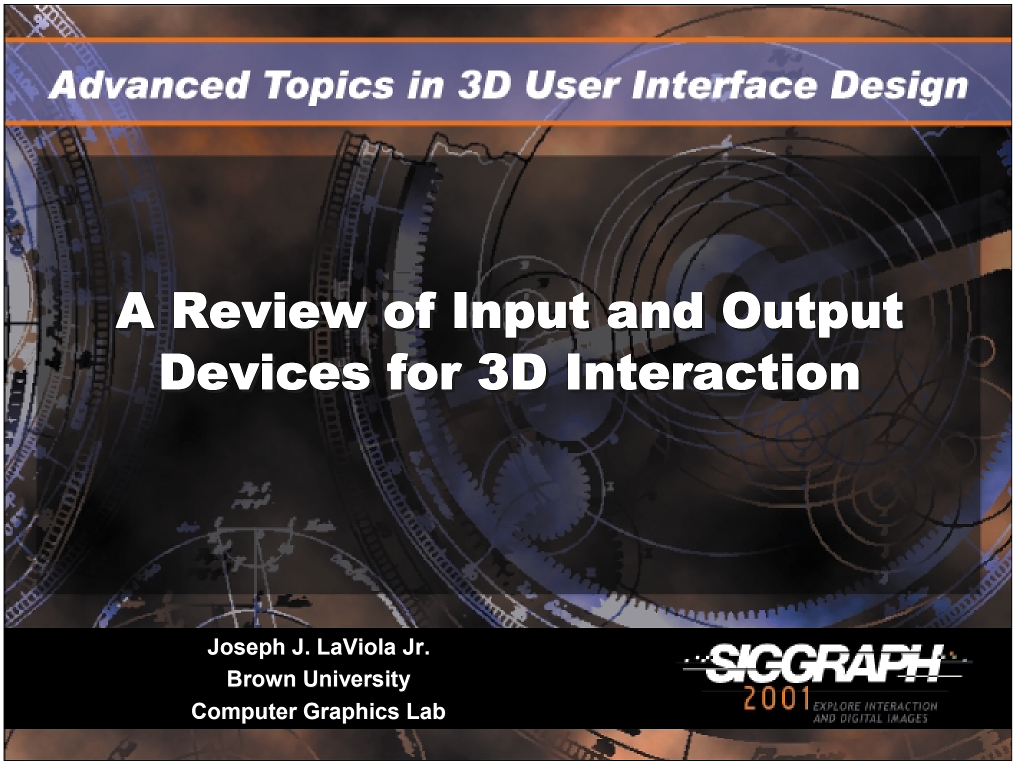

A Review of Input and Output Devices for 3D Interaction

Joseph J LaViola Jr
Ph.D Candidate
Brown University,
Department of Computer Science
Providence, Rhode Island

Lead Consultant and Founder
JJL Interface Consultants, Inc.

http: http://www.cs.brown.edu/people/jjl
email: jjl@cs.brown.edu

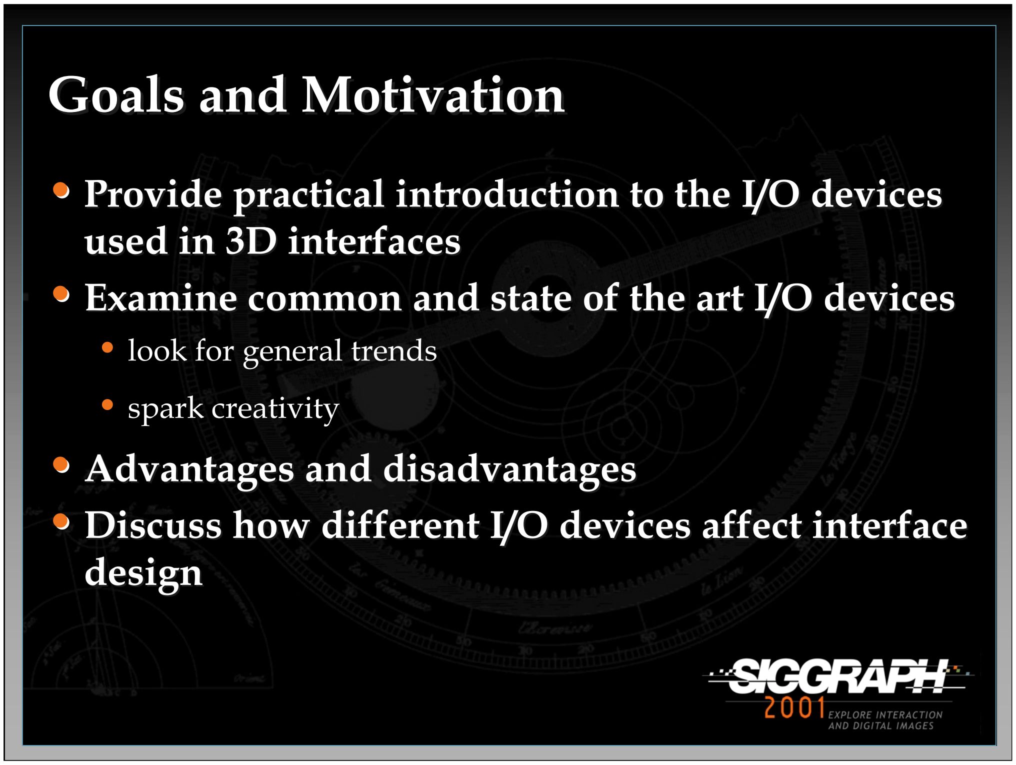

In this lecture we will discuss the various input and output devices that are used in 3D user interfaces and virtual environment applications.

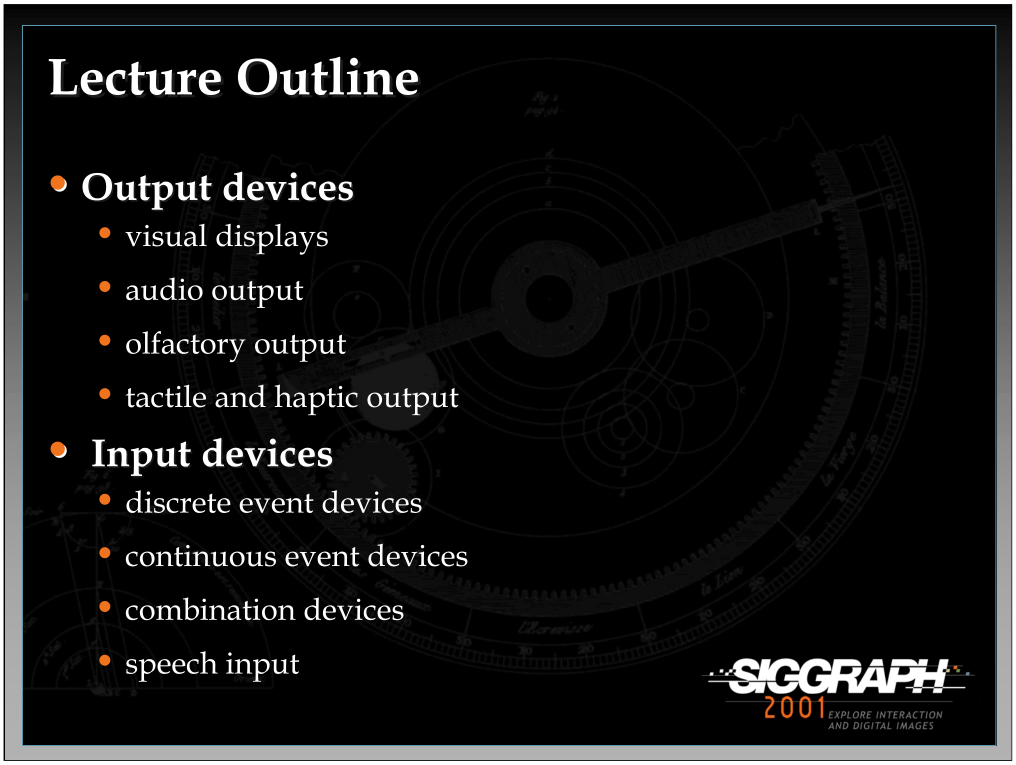

The first part of the lecture will describe a number of output devices that stimulate the human visual, auditory, haptic, and tactile systems. In the second part of the lecture, we will look at the many different ways a user can interface to a 3D world. With each device, we will discuss the advantages, disadvantages, and its effects on interface design.

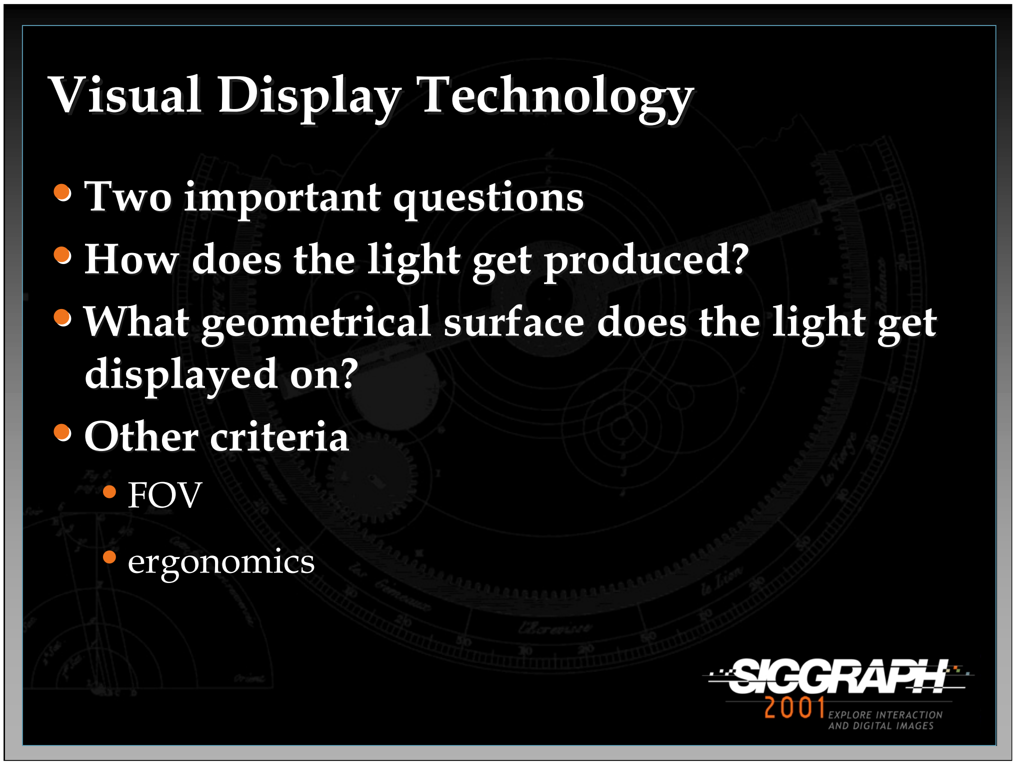

Visual display systems for virtual reality and other 3D applications have two important and interrelated components. The first is the technology underlying how the light we see gets produced; the second is the type and geometrical form of surface on which this light gets displayed. Other criteria for thinking about visual display systems include field of view (FOV) and ergonomics.

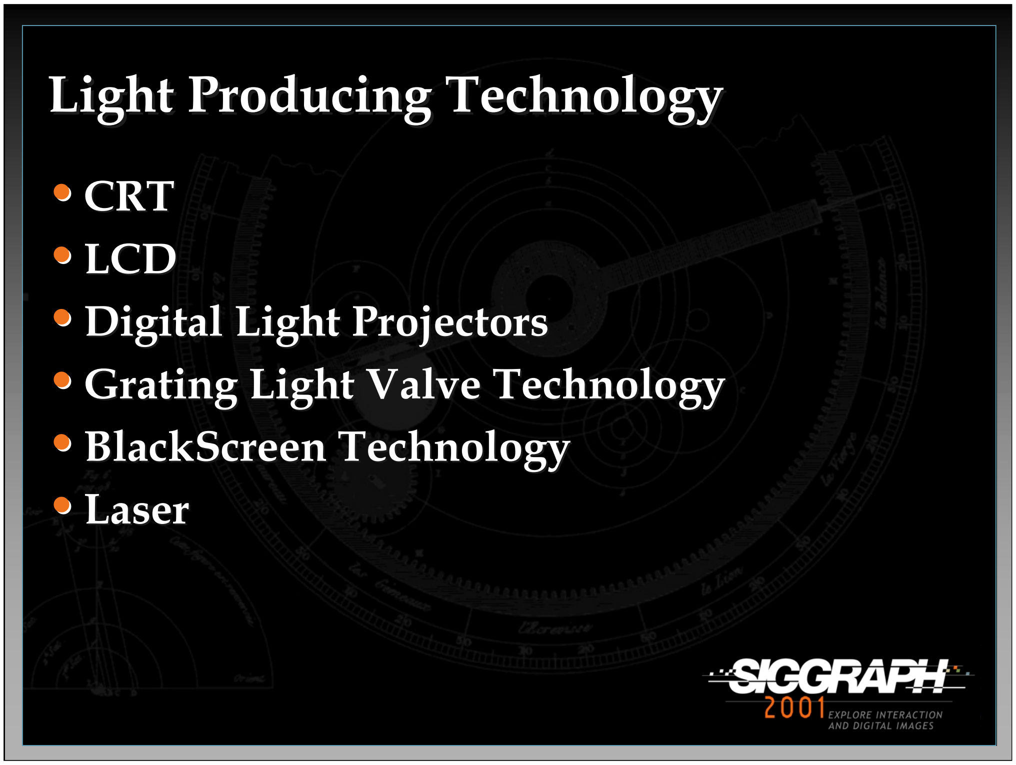

A number of different methods exist for producing the light displayed on a geometrical surface. While Cathode Ray Tubes (CRTs) and Liquid Crystal Display (LCD) panels and projectors are currently the norm in today's marketplace, newer light-producing techniques are emerging. Texas Instrument's Digital Micromirror Device (DMD) is currently available in digital light projectors. The DMD is a thumbnail-size semiconductor light switch which consists of an array of thousands of microscopic sized mirrors, each mounted on a hinge structure so that it can be individually tilted back and forth. When a lamp and projection lens are positioned in the right places, DLP processes the input video signal and tilts the mirrors to generate a digital image.

Silicon Light's Grating Light Valve technology is a micromechanical phase grating which provides controlled diffraction of incident light to produce light or dark pixels in a display system. Their approach can be used to build a 10-bit-per-pixel, high-resolution display compared with 8-bit-per-pixel LCD display. There is hope that this kind of technology may be commoditized for personal displays.

-continued on the next page

Jenmar Visual System's BlackScreen Technology (used in the ActiveSpaces telecollaboration project of Argonne National Laboratories) captures image light into a matrix of optical beads, which focus it and pass it through a black layer into a clear substrate. From there it passes into the viewing area. This screen material presents a black level undegraded by ambient light, making it ideal for use with high-luminosity projection sources and nonplanar tiled displays such as caves.

Finally, laser light is another approach to light production which projects light directly onto the retina. See the slide on Virtual Retinal Displays later in the lecture.

References:

www.dlp.com

www.siliconlight.com

www.jenmarvs.com

www.mvis.com

www.3d-perception.com

Yoder, Lars. "The Digital Display Technology of the Future" INFOCOMM'97, June 1997.

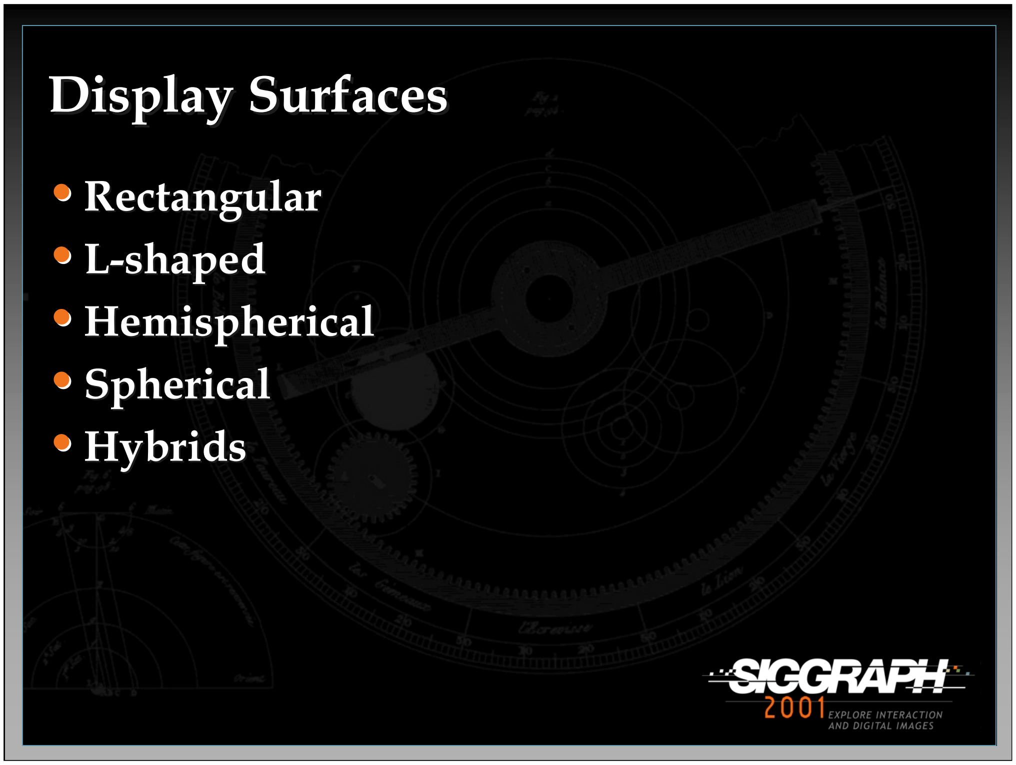

Unfortunately, no "one size fits all" display surfaces exist for virtual reality and 3D applications. Rather, many different kinds offer advantages and disadvantages. Choosing an appropriate display surface depends on the application, tasks required, target audience, financial and human resources available, and so on. In addition to traditional rectangular display surfaces, more interesting display geometries are starting to affordably emerge including hemispherical and spherical displays and those which combine different geometries together.

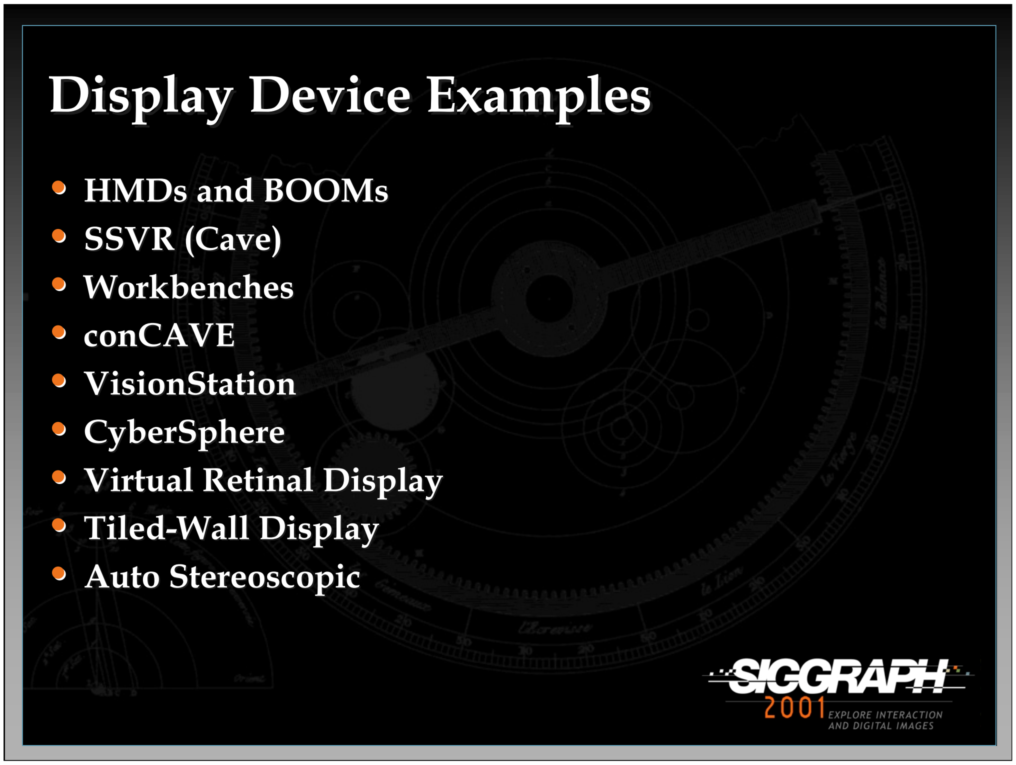

This slide shows a representative sample of the many visual display devices that exist either in the research lab or in the industrial marketplace. We will look at each example in turn and examine how these devices affect 3D interface design.

One of the most common display devices used for virtual environment applications is the head mounted display (HMD). With a tracking device attached to the device, it produces a stereoscopic view that moves relative to the user's head position and orientation. Although traditionally the user cannot naturally see the real world, cameras are sometimes mounted on the HMD which allows it to display both real world video and graphical objects. In addition, some HMDs offer see-through options. This type of technology is used in augmented reality systems.

Since each eye is presented with one screen, HMDs allow for good stereoscopic viewing. These two screens are very close to the user's eyes (1 to 2 inches). As a result, all viewable objects are behind the screen so any object clipping will appear to the user as being outside his/her field of view. A big disadvantage of HMDs is that can get heavy very quickly and, unfortunately, the higher the HMD's quality, the heavier it usually is. Although HMDs are still popular in many VR labs and entertainment centers, researchers and practitioners are rapidly moving towards projection-based display devices especially when high-resolution graphics are required.

-continued on the next page

Since the real world is completely blocked out of the user's view, interaction while wearing an HMD requires the user to have some type of graphical representation of either one or both hands or the input device used. These graphical representations can be as simple as a cube or as complicated as a hand model containing 50000 or more polygons. HMDs also put a strain on the types of input devices that can be used since the user cannot physically see the device in order to use it.

The arm mounted display shown in the picture on the right is called a BOOM developed by Fakespace. It has a counter weight on the the opposite side of the display to make the device easier to manipulate. The device also uses mechanical tracking technology to track the user's head position and orientation. The latest version of the BOOM supports resolutions of 1280x1024 pixels per eye which is better than most average quality HMDs. Since the user does not have to wear the device, it is easy to operate and allows for different users to trade places quickly. As with the HMD, providing one screen per eye allows for good stereo quality. Since the BOOM is physically attached to a large stand, the user's movement is limited. Users can move in about a six foot diameter around the center of the stand. Another disadvantage with the BOOM is the user has to have at least one hand on the device which can limit various types of two-handed interaction.

References:

www.nvis.com

www.virtualresearch.com

www.stereo3d.com

www.fakespace.com

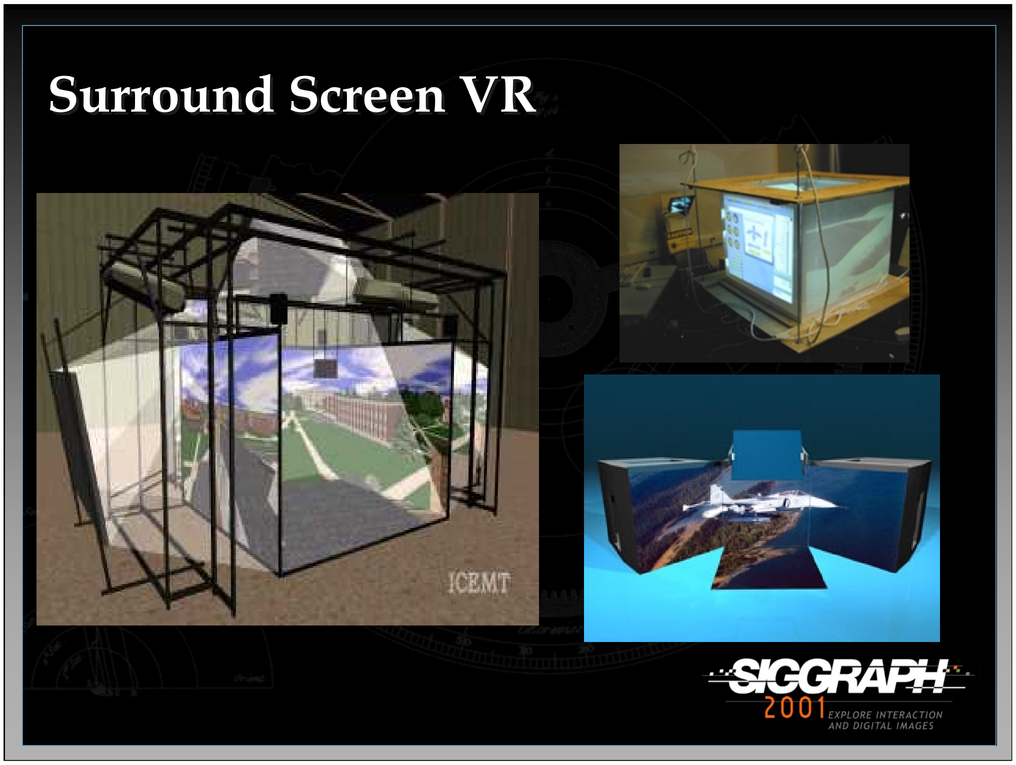

The first surround screen virtual reality system was developed by Carolina Cruz-Neira at the Electronic Visualization Laboratory at the University of Illinois at Chicago. This system was called the CAVE. Today, the term CAVE is copyrighted by Pyramid Systems (now a part of Fakespace). So the general term for such a device is a surround screen environment. These systems also go by other names such as the C2, C6, and TAN Cube and can have anywhere from three to six screens.

The figure in the upper right corner of the slide is called a Computer-driven Upper Body Environment (CUBE). It is a 360° display environment composed of four 32”X28” rear-projected Plexiglas screens. Guests stand inside the CUBE, which is suspended from the ceiling, and physically turn around to view the screen surfaces. The screens are approximately 1’ from a guest’s face and extend down to his or her midsection. It was developed at the Entertainment Technology Center at Carnegie Mellon University and represents a small-personalized version of a SSVE.

The figure in the lower right corner of the slide shows the RAVE, a reconfigurable advanced visualization environment developed by Fakespace. It is designed to be a flexible display device which can be used as a 30 foot flat wall, a 30 foot variable angle immersive theatre, a Cave-like three wall and floor immersive environment, an L-shaped cove with separate 10 foot wall, and three separate 10 foot review walls.

-continued on the next page

There are a number of advantages to using an SSVR system. They provide high resolution and a large FOV. Users only needs a pair of light weight shutter glasses for stereo viewing and have the freedom to move about the device. Additionally, real and virtual objects can be mixed in the environment and a group of people can inhabit the space simultaneously.

One of the biggest disadvantages of SSVR systems is the fact that they are so expensive and require such a large amount of physical space. Another problem with an SSVR system, as well as any projection-based display system, is stereo viewing can be problematic. When the user gets close to the display or when objects appear to be right in front of the user, it becomes more and more difficult to fuse the two images together. Eye strain is a common problem in these situations. Finally, even though multiple users can inhabit the space at one time, due to technological limitations, no more than two can be head-tracked.

Although physical objects do not have to be represented as graphical objects in SSVR systems, an important issue arises when a physical object passes in front of graphical objects that should appear in front of the said physical object. This is a common problem with any projection-based display device and can hinder the immersive experience.

In most cases the user wears a pair of shutter glasses for stereo viewing. These glasses are synched to flicker at a rate equal to the refresh rate of the graphics engine. These signals are sent to the glasses by infrared signal. So, if the signal is blocked, the shutter glasses will stop working and the stereo effect will be disrupted. As a general guideline, it is a good idea to never have the user move his/her hands or other physical objects in the line of sight of the glasses and emitters.

References:

Cruz-Neira, Carolina, Daniel Sandin, and Thomas Defanti. "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE" SIGGRAPH'93, 135-142.

http://www.etc.cmu.edu/projects/cube/index.html

www.fakespace.com

www.mechdyne.com

www.tan.de

One of the newest types of display devices is the projection-based drafting table. These devices are usually single screen and go by many different names such as Fakespace's Immersadesk and Immersive WorkBench and VersaBench (pictured on the left), the Barco Baron, and the ITI VisionMaker Digital Desk. In some cases just a single vertical screen is used. The second picture to the left shows Fakespace's Mini Workbench. A pressure sensitive display surface for 2D input is an optional feature with the Workbench. The TAN Holobench, shown in the two pictures on the right, is an L-shaped desk which provides a holographic impression to the user since objects appear to be raised above the it.

In general, workbenches provide high resolution displays, make for an intuitive display for certain types of applications (i.e 3D modeling and drafting, virtual surgery), and can be shared by several users. However, due to technological limitations, at most two users can be head tracked and these devices suffer from the same stereo problems that all rear-projected devices do.

References:

www.iti-world.com

www.barco.com

www.fakespace.com

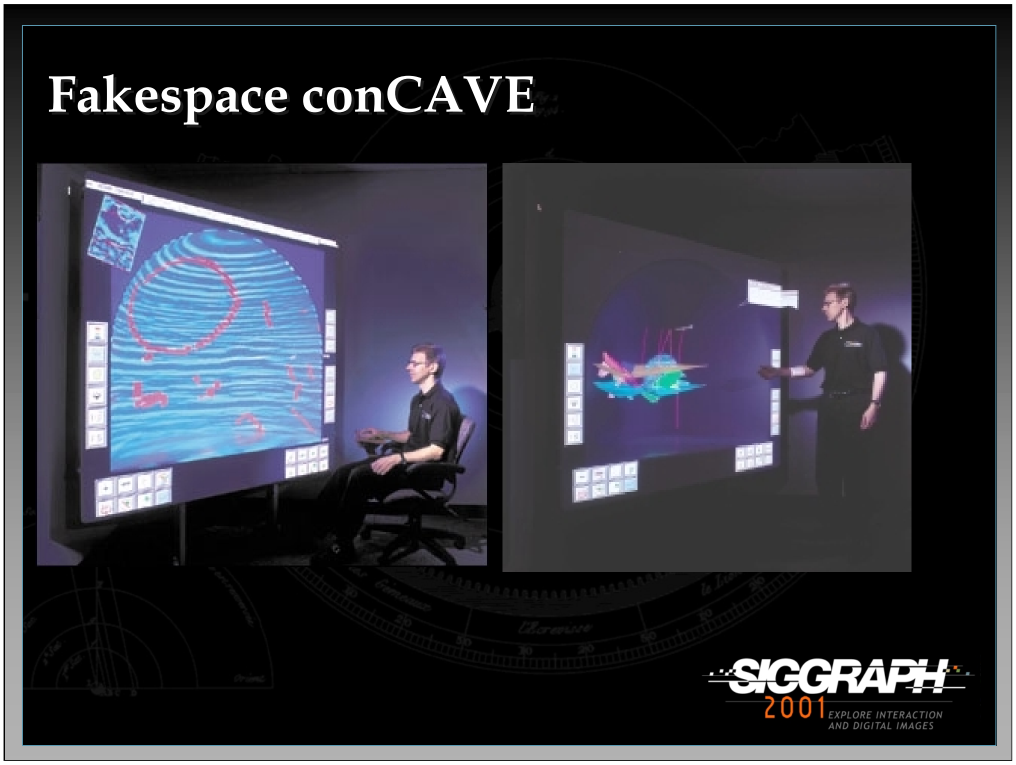

The conCAVE is a rather unique display device in that it combines flat, cylindrical, and spherical display surfaces to form one single device. The conCAVE creates spatially correct 3D "tunnel-view" images of volumetric data that extend from floor to ceiling and side to side. 3D perspective views are depth-enhanced, generating a sense of stereoscopic imagery without the need for special shutter glasses. The conCAVE also has a simple, pull down flat screen in front of the device that creates a large display for standard images such as maps, cross sections, spread-sheets or presentations. From an interface perspective, the device provides the user with both 3D interaction using a tracked paddle and 2D interaction using the virtual buttons on the front of the flat display surface.

References:

www.fakespacesystems.com

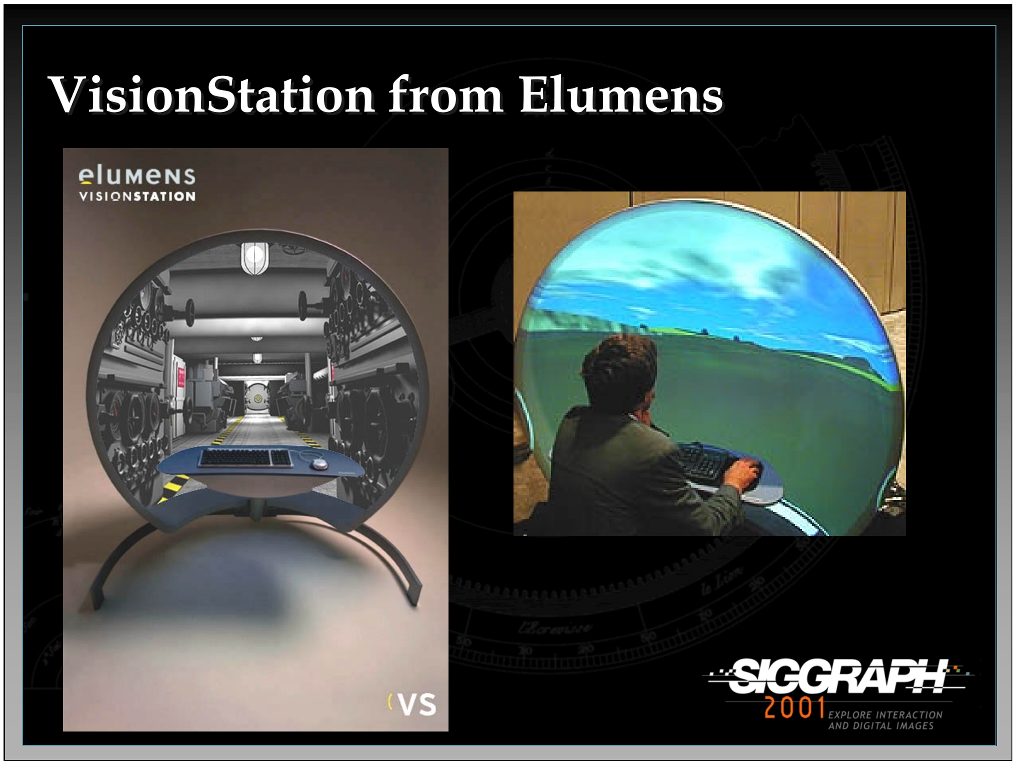

The VisionStation is a personalized device that uses a hemispherical front-projected display surface. It uses special proprietary software and optics for the projection lens to display images in a 180 by 180 degree field of view. The user sits in front a of a small table and can interact with 3D applications using keyboard and mouse or 3D input devices. One of the major problems with this device is that it is front-projected which means 3D interaction is limited since moving too close to the display surface will cast shadows on the screen.

References:

www.elumens.com

The CyberSphere is a fully spherical immersive display device prototype created by VR Systems UK and the Warwick Manufacturing Group. The system uses a large, hollow, translucent sphere (3.5 meters in diameter) supported by a low pressure cushion of air. The air cushion enables the sphere to rotate in any direction. A single user is able to enter the sphere using a closable entry hatch. Once inside, walking movements cause the large sphere to rotate. This rotational movement is transferred to a smaller secondary sphere, which is supported by means of a ring mounted upon a platform. Rotational movement of the smaller sphere is measured with rotation sensors, pushed against the circumference of the sphere with spring loaded supports. The rotation sensors send signals to the computer which update the projected images in order to give the user the illusion of walking freely through the virtual environment.

The device is still in the prototype stage and is not commercially available.

References:

www.ndirect.co.uk/~vr-systems/sphere1.htm

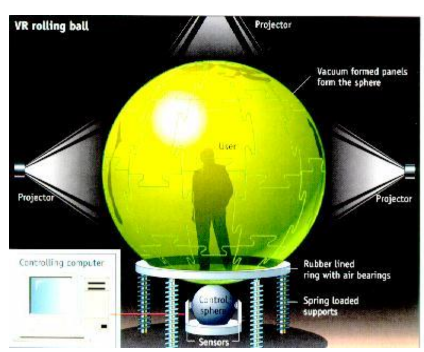

The Virtual Retinal Display (VRD) was invented at the Human Interface Technology Lab in 1991. It is based on the idea that images can be directly displayed onto the retina. With a VRD, a photon source is used to generate a coherent beam of light which allows the system to draw a diffraction limited spot on the retina. The light beam is intensity modulated to match the intensity of the image being rendered. The beam is then scanned to place each image point, or pixel at the proper position of the retina. VRDs are commercially available from Microvision and are used in augmented reality systems. For more details see the references below.

Reference:

www.hitl.washington.edu/research/vrd/

www.mvis.com

Tiled display surfaces, which combine many display surfaces and light producing devices, are becoming increasingly popular for a number of 3D and virtual environment applications. Tiled displays offer greater image fidelity than other immersive and desktop displays due to an increased number of pixels displayed over an area that fills most of a user's or group of user's FOV. The display in the figure is the Scalable Display Wall developed at Princeton University. It has a resolution of 8192 by 3064 pixels and is 18 feet long and 8 feet high. Although an intriguing display device, the tiled wall display has a number of research challenges including hardware setup, maintaining projector calibration and seamless imaging.

References:

http://www.cs.princeton.edu/omnimedia/index.html

"Special Issue on Large Wall Displays", IEEE Computer Graphics and Applications, Vol. 20, No. 4, July/Aug. 2000.

Other visual output devices use lenticular, volumetric, and holographic display technology. Most of these technologies are expensive and in the early stages of development so they are rarely utilized in mainstream 3D interfaces. However, once these technologies become more affordable and the technology has progressed sufficiently, a number of interesting interface issues will arise.

The picture on the left shows a 14 inch lenticular display developed at the Heinrich-Hertz Institute for Communication Technology in Berlin, Germany.

The picture on the right show a lenticular display developed at the Dresden University of Technology.

References:

http://at.hhi.de

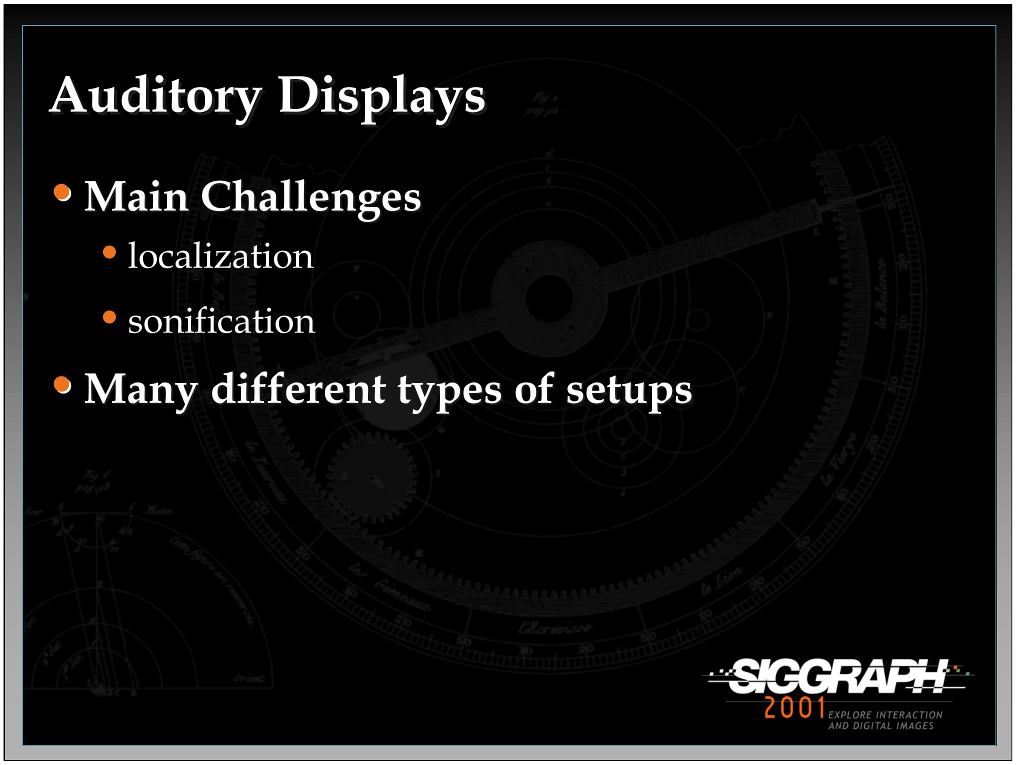

There are two different ways, localization and sonification, in which sound can be used as an output medium in virtual environment applications. In localization, the goal is to generate three dimensional sound. In sonification, the goal is to turn certain types of information into sounds.

There are a number of different ways in which an auditory system can be setup. A simple setup is to use stereo head phones. However, this restricts usage to only one person at a time. Another setup is to place speakers in certain logistic areas around the environment. This setup allows for more than one user to take part in the experience but is somewhat more complicated to setup and write software for.

Reference:

Begault, Durand R. 3D Sound For Virtual Reality and Multimedia. Academic Press, 1994.

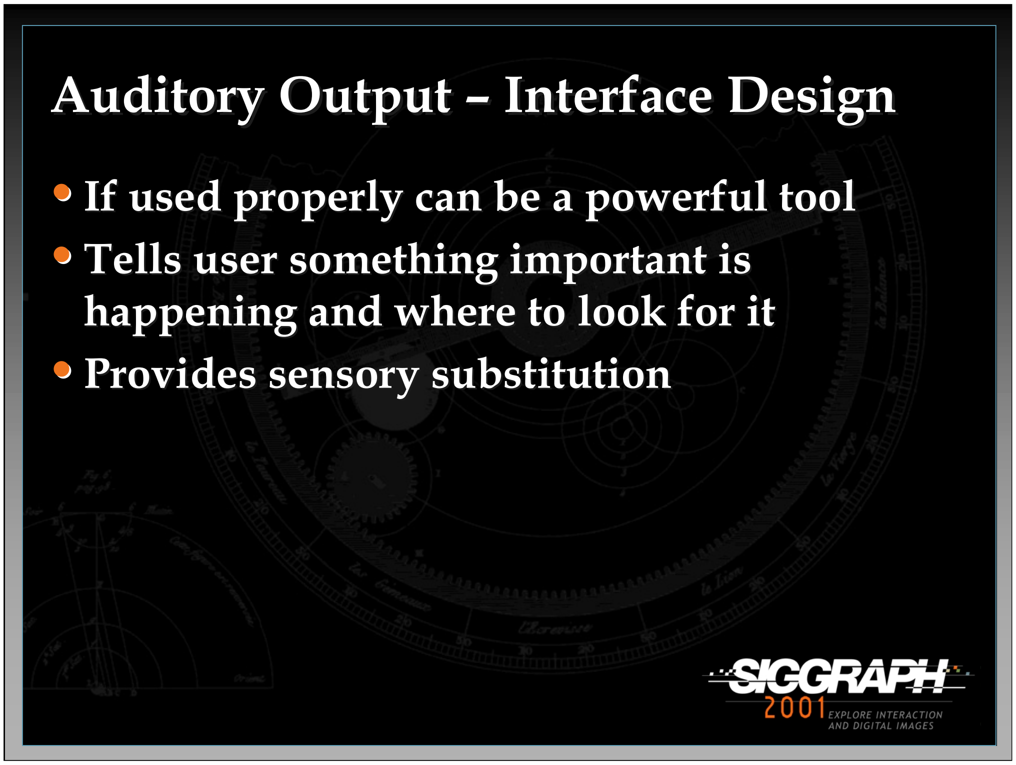

Auditory output can be very powerful when applied correctly in 3D virtual environments. It is especially useful in collaborative applications where participants can get a sense for where others are in the environment. It also can be used for sensory substitution which is important when, for example, no haptic or tactile feedback is present. A sound could substitute the feel of a button press or the moving of an object in the virtual space.

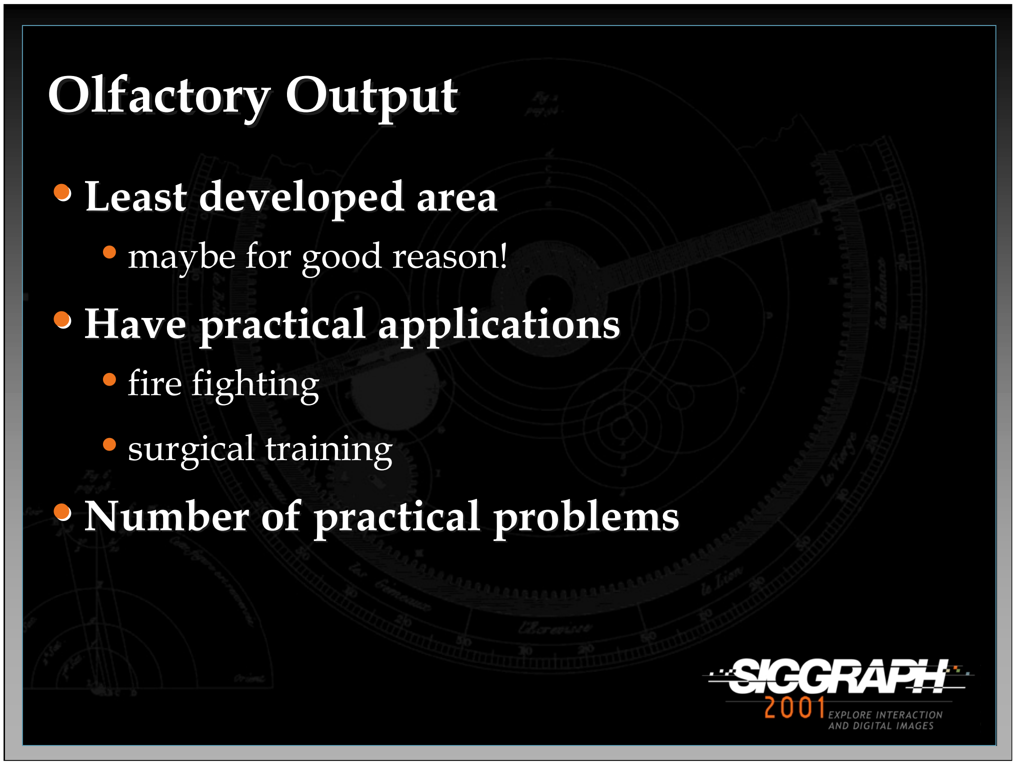

Olfactory interfaces are one of the least developed areas the virtual reality and 3D applications. There are a number of interesting design considerations when dealing with olfactory output such as odor storage and display as well as cleaning the air input and controlling the breathing space for the individual.

References:

www.hitl.washington.edu/people/tfurness/courses/inde543/reports/3doc/

Youngblut, Christine, Johnson, Rob E., Nash, Sarah H., Weinclaw, Ruth A., Will, Craig A., Review of Virtual Environment Interface Technology IDA Paper P-3186. Chapter 8, p. 209-216, http://www.hitl.washington.edu/scivw/IDA/.

Dinh, H.Q., N. Walker, L.F. Hodges, C. Song, and A. Kobayashi, “Evaluating the Importance of Multi-sensory Input on Memory and the Sense of Presence in Virtual Environments”, In IEEE Virtual Reality’99, 222-228, 1999.

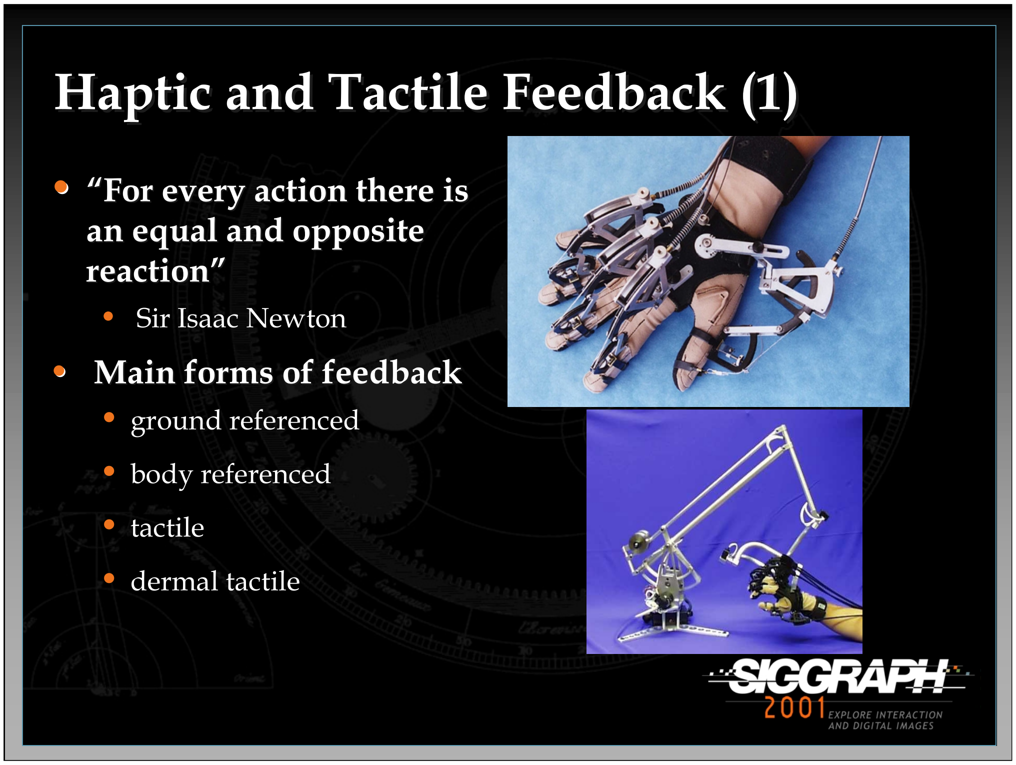

Haptics represents a critical component in virtual environment interaction. Allowing a user to touch and feel in the virtual world in the same way that they do in the physical world is extremely powerful. Unfortunately, haptic and tactile output device research is still in its early stages.

There are essentially four different methods in which haptic and tactile feedback is generated. The first method is ground-referenced feedback which creates a physical link between the user and ground with the feedback relative to a single contact point. An example shown in the bottom picture is Virtual Technologies' CyberForce. The second method is body-referenced feedback which places a device on some part of the user's body. An example of a body-referenced haptic device is Virtual Technologies' CyberGrasp which is shown in the top picture. The third method for generating feedback is tactile which uses some type of oscillatory or vibrating device to stimulate the user's tactile sense. Finally, the last method of generating feedback is via dermal tactile which stimulates the user's nerves in the fingertips.

References:

www.sensable.com

www.virtex.com

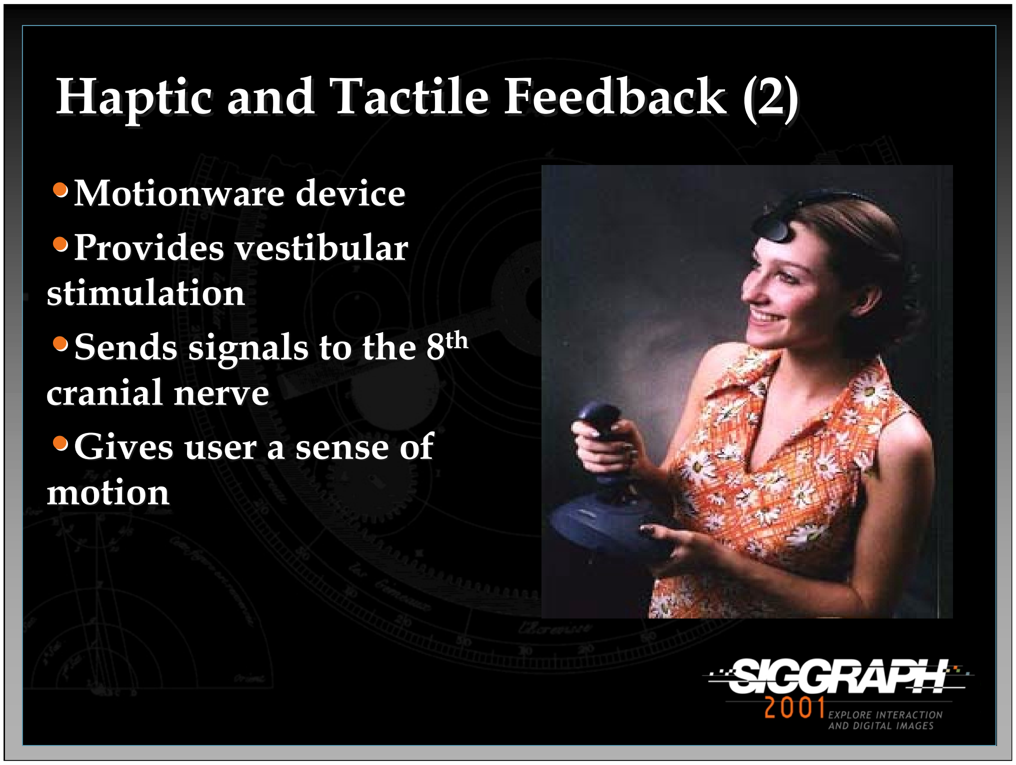

Another type of tactile feedback device is Motionware being developed by Virtual Motion. Motionware sends electrical current to the 8th cranial nerve located behind the wearer's ear. Sending these electrical signals to the 8th cranial nerve provides the user with vestibular stimulation which can mimic the sense of motion. The device will be priced at around $100.

References:

www.virtual-motion.com

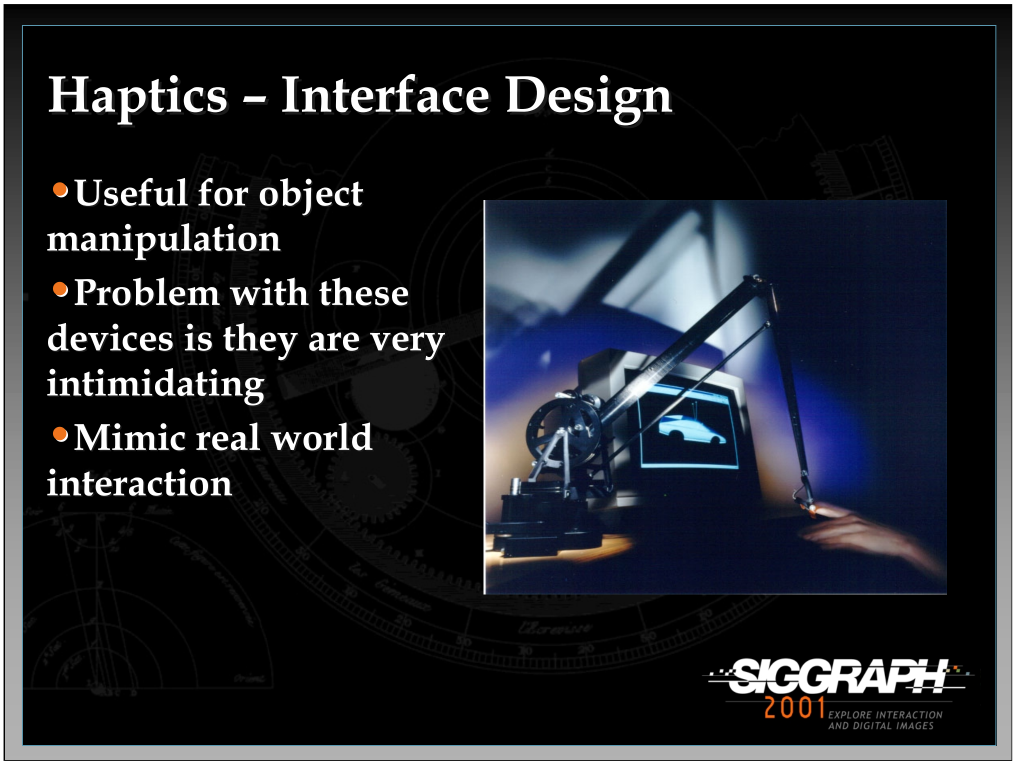

Haptic feedback from devices like the CyberGrasp are very good for grabbing objects and moving them around and they can provide a limited form of real world interaction. The main problem with these devices is that they are somewhat intimidating to the user. People are commonly afraid to put these devices on and once they do they're afraid they'll break them or get injured. The picture shows the Phantom 3.0 from Sensable Technologies. This device is less intrusive than the CyberGrasp.

References:

Burdea, Grigore C. Force and Touch Feedback for Virtual Reality. Wiley Interscience, 1996.

There are many different haptic devices that are being developed in research labs around the world. The slide shows from top to bottom, left to right the Pneumatic Master Arm from Southern Methodist University, a 5 DOF haptic device from the University of Colorado, a magnetic levitation haptic interface from Carnegie Mellon University, and a tactile display from the Karlsruhe Research Center in Germany.

References:

haptic.mech.northwestern.edu

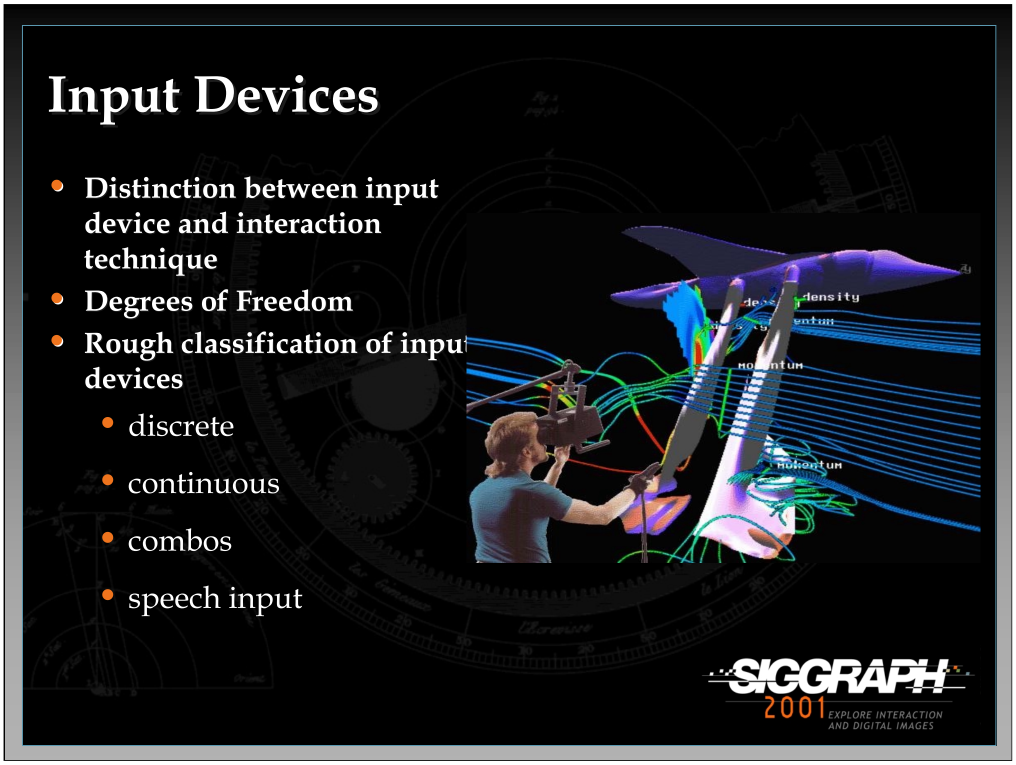

There is a distinction that must be made when we talk about input devices and interaction techniques. Input devices are just the physical tools that are used to implement various interaction techniques. In general, many different interface techniques can be mapped onto any given input device. The question is how natural, efficient, and appropriate a given input device will work with a given technique.

When talking about input devices it is convenient to talk about the degrees of freedom (DOF) that an input device has. For example, a device such as a tracker generally produces 3 position values and 3 orientation values for a total of 6 DOF. For the most part, a device with a smaller number of DOF can be used to emulate a device with a higher DOF with the addition of buttons or modifier keys.

See the papers by Shumin Zhai in the papers section of the course notes for a series of experiments that evaluate a number of the input devices presented in this part of the lecture.

Discrete input devices simply generate one event at a time based on the user. In other words, when the user presses a button an event is generated which is usually a boolean value stating whether the button was pressed down or released. The keyboard is an obvious example of a discrete input device.

The Pinch Glove system developed by Fakespace is another example of a discrete input device. These gloves had a conductive material at each of the fingertips so that when the user pinches two fingers together a electrical contact is made which generates a boolean value. There are many different pinching combinations that can be made which allows for a significant amount of input device to task mappings.

The Interaction Slippers are a custom made device which allows users to perform toe and heel tapping for invoking commands. The slippers use conductive cloth contacts and a Magellan Trackman Live! wireless mouse. See the paper, “Hands-Free Multi-Scale Navigation in Virtual Environments”, in the papers section of the course notes for more details.

The Painting Table is another example of a discrete input device that is used in the CavePainting application, a system for painting 3D scenes in a virtual environment. The device uses a set of conductive cloth contacts as well as traditional buttons and digital sliders. Users can dip the paint brush prop into the colored cups to change brush strokes. The bucket is used to throw paint around the virtual canvas.

References:

Keefe, D., Acevedo, D., Moscovich, T., Laidlaw, D., and LaViola, J. “CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience”, Proceedings of the 2001 Symposium on Interactive 3D Graphics, 85-93, 2001.

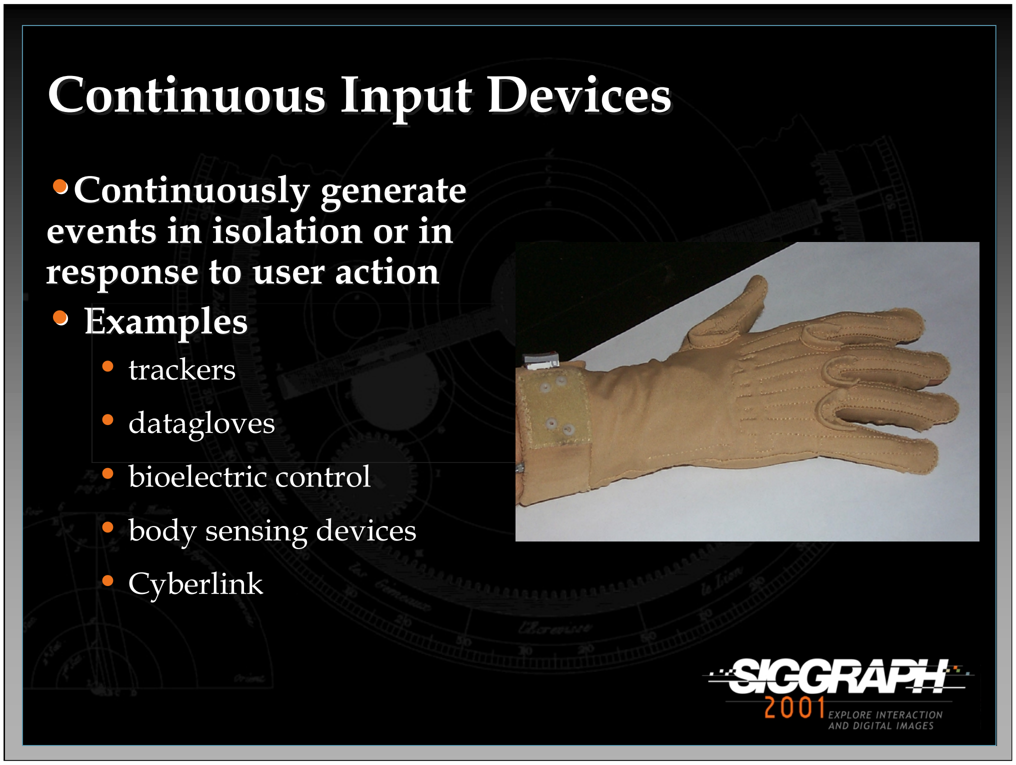

Continuous input devices generate a continual stream of events in isolation (no user manipulation) or in response to user action. For example, a tracker is a device which will continually output position and orientation records even if the device is not moving. These types of devices are important when we want to know where something is in the virtual space and we do not want to have to keep asking for it. A perfect example of this is head tracking. Two of the most common continuous devices are trackers and datagloves.

Another type of continuous input device is the Cyberlink, a brain-body actuated control technology that combines eye-movement, facial muscle, and brain wave bio-potentials to generate input signals. The Cyberlink has three sensors in a headband and its interface unit amplifies and translates the brain wave, facial muscle and eye-movement data into separate frequencies and transmits them to an PC serial port. The Cyberlink software processes and displays theses frequencies and 10 continuous command signals called Brainfingers.

References:

www.brainfingers.com

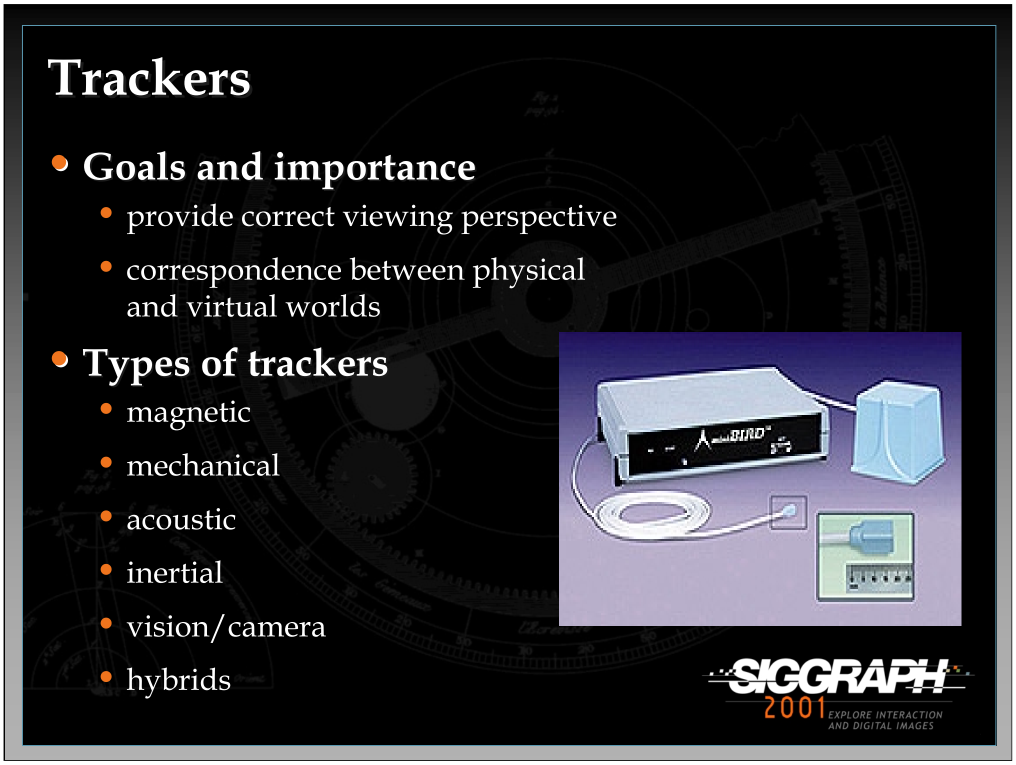

One of the most important aspects of 3D interaction in virtual worlds is providing a correspondence between the physical and virtual environments. As a result, having accurate tracking is extremely important to making the VE usable. Currently there are a number of different tracking technologies in the marketplace. The different types are shown in the slide.

Magnetic tracking uses a transmitting device that emits a low frequency magnetic field that a small sensor, the receiver, uses to determine its position and orientation relative to a magnetic source. These trackers can use extended range transmitters which increase the range of the device from around an 8 foot radius to anywhere from a 15 to 30 foot radius. The tracker shown in the picture is called the Ascension MiniBird. It uses a smaller emitter and receivers and has better accuracy than the regular system. However it's range is limited to about a 4 foot radius. It is primarily used in medical applications where range of the device is not a factor.

Mechanical trackers have a rigid structure with a number of joints. One end is fixed in place while the other is attached to the object to be tracked (usually the user's head). The joint angles are used to obtain position and orientation records. The Fakespace BOOM uses this type of tracking technology.

-continued on the next page

Acoustic tracking devices use high frequency sound emitted from a source component that is placed on the hand or object to be tracked. Microphones placed in the environment receive ultrasonic pings from the source components to determine their location and orientation. In most cases, the microphones are placed in a triangular fashion and this region determines the area of tracked space. One of the most interesting problems with this type of tracking is that certain noises such as jingling keys or a ringing phone will interfere with the device.

Inertial tracking systems use a variety of inertial measurement devices such as gyroscopes, servo accelerometers, and micro-machined quartz tuning forks. Since the tracking system is in the sensor, range is limited to the length of the cord which attaches the sensor to the electronics box. Two of the big limitations of these devices is that they only track orientation and are subject to error accumulation. The InterSense IS300 handles error accumulation by using a gravitometer and compass measurements to prevent accumulation of gyroscopic drift and also uses motion prediction algorithms to predict motion up to 50 milliseconds into the future.

Camera/vision based tracking take one or more cameras and places them in the physical environment. The cameras then grab video of the user or object to be tracked. Usually image processing techniques, such as edge detection algorithms, are used to identify the position and/or orientation of various body parts such as the head and hands. Setting up vision-based tracking systems can be difficult since there are many parameters that must be fixed in order to track the user properly. These parameters include the number of cameras, the placement of the camera, what background (what is in back of the user) is put up, and if the user will be wearing special optical tools such as LEDs or colored gloves to aid in tracking. Ascension's laserBIRD is an example of an optical tracking device. laserBIRD delivers accurate position and orientation tracking without environmental interference or distortion. Its miniaturized scanner reflects laser beams throughout the work space. Each sensor, attached to a tracked object, instantly picks up the laser beams. Signals are then directed back to the scanner's DSP electronics for processing and transmission to a host PC or workstation. Other vision-based approaches include the UNC HighBall which uses LED beacons mounted on the ceiling.

-continued on the next page

Hybrid trackers attempt to put more than one tracking technology together to help increase accuracy, reduce latency, and, in general, provide a better virtual environment experience. An example is the InterSense IS600. It combines inertial and ultrasonic tracking technologies which enables the device to attain 6 DOF. The major difficulty with hybrid trackers is that the more components added to the system, the more complex the device becomes.

Other types of hybrid tracking include the combination of video cameras and structured digital light projectors. Combining these two technologies allow for the capture of depth, color, and surface reflectance information for objects and participants in the environment. This approach is currently being used at the University of North Carolina, Chapel Hill in their Office of the Future project as well as in the National Tele-Immersion Initiative. The picture shows two user collaborating in two remote locations.

References:

www.ascension-tech.com

www.polhemus.com

www.isense.com

www.3rd-tech.com (Commercial version of the UNC HighBall Tracker)

Raskar, Ramesh, Welch, Greg, et al. "The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays" SIGGRAPH '98, ACM Press, 179-188.

Amela Sadagic et. al., " National Tele-Immersion Initiative: Towards Compelling Tele-Immersive Collaborative Environments", Medicine meets Virtual Reality 2001 conference, January 24-27, 2001.

Eye tracking systems provide applications with knowledge of the user's gaze direction. This information opens the door to a number of interesting interaction techniques such as eye directed selection and manipulation. The figure on the left shows the Eyegaze system, a non-intrusive approach which uses an infra-red source that reflects off of the pupil, developed by LC Technologies. The figure on the right shows iView, a head-mounted eye tracking device developed by SensoMotoric Instruments.

References:

www.eyegaze.com

www.smi.de

Data gloves measure finger movement of the hand by using various kinds of sensor technology. These sensors are embedded in the glove or placed on top of the glove, usually on the back of the hand. The number of sensors in the glove depends on the manufacturer. Virtual Technologies' CyberGlove has either 18 or 22 sensors which can measure at least 2 joints in each finger, wrist roll and yaw, and others. These types of gloves are commonly used for hand gesture and posture recognition which can be applied to a variety of different interface techniques in virtual environments. Fifth Dimension Technologies (5DT) offers gloves that have either 5 sensors, one for each fingertip or 16 sensors, 2 for each finger and abduction between fingers. 5DT also has wireless versions of each glove.

References:

www.virtex.com

www.5dt.com

A recent development at NASA Ames Research Center is a bioelectric input device which reads muscle nerve signals emanating from the forearm. These nerve signals are captured by a dry electrode array on the arm. The nerve signals are analyzed using pattern recognition software and then routed through a computer to issue relevant interface commands. The figure on the left shows a user entering numbers on a virtual numeric keypad while the figure on the right shows a user controlling a virtual 757 aircraft.

References:

Jorgensen, Charles, Kevin Wheeler, and Slawomir Stepniewski. Bioelectric Control of a 757 Class High Fidelity Aircraft Simulation, http://ic.arc.nasa.gov/publications/index.html, 1999.

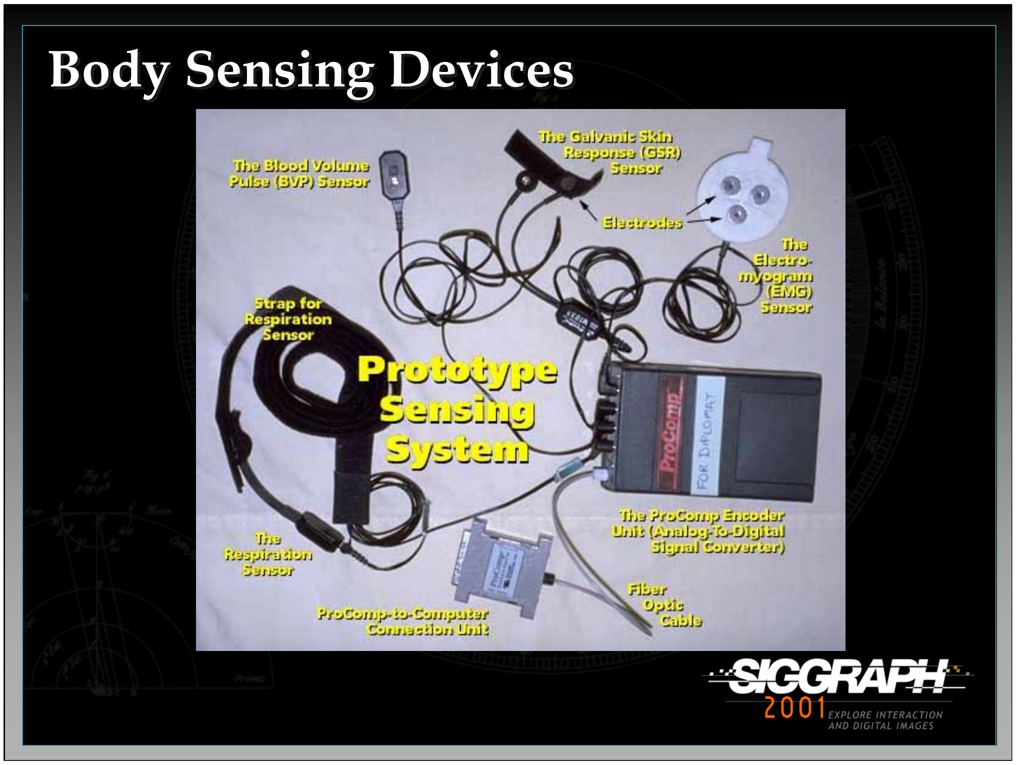

The MIT Media Lab's affective computing group has developed a Prototype Physiological Sensing System which includes a Galvanic Skin Response sensor, a Blood Volume Pulse sensor, a Respiration sensor, and an Electromyogram. By using this prototype, interface developers can monitor a user's emotional state to dynamically modify an application's interface to better fit the user's needs.

References:

http://www.media.mit.edu/affect/

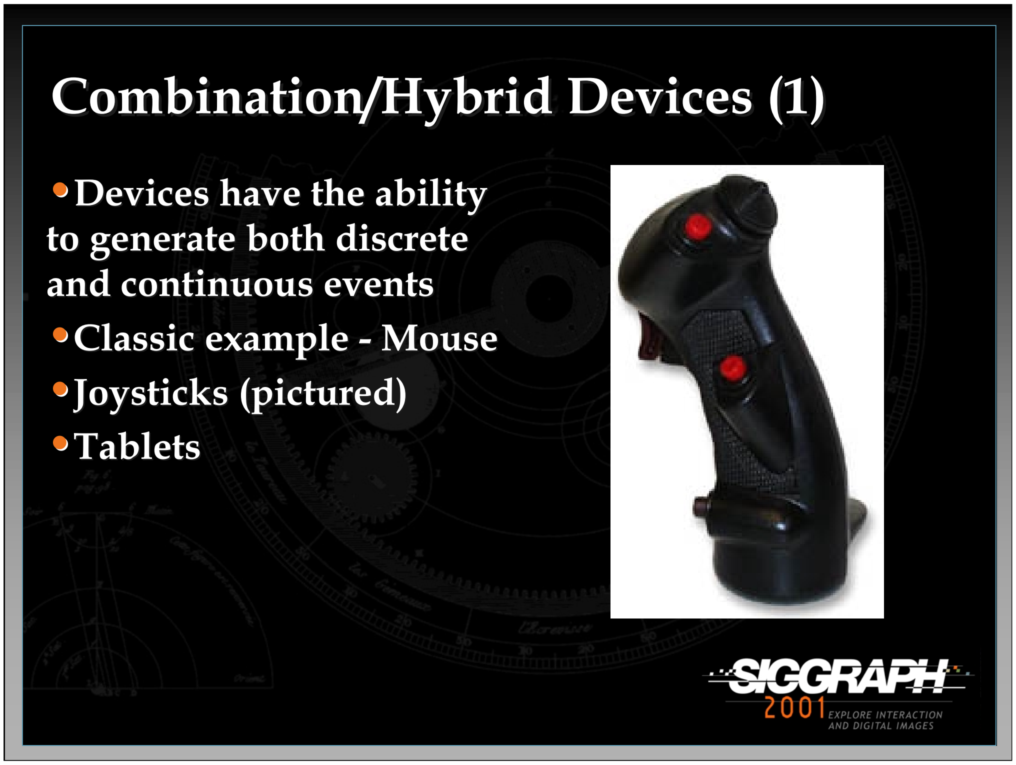

A combination/hybrid input device combines both discrete and continuous event generating devices to form a single device that is more flexible. Two of the most common hybrid devices are the joystick and mouse. Another device in this category is the pen-based tablet. Pen-based tablets are becoming more and more popular in virtual environment applications because they give the user the ability to interact in 2D which provides a useful combination in certain interfaces. The figure shows the SpaceStick developed by MUSE Virtual Presence.

References:

www.vrweb.com

www.wacom.com

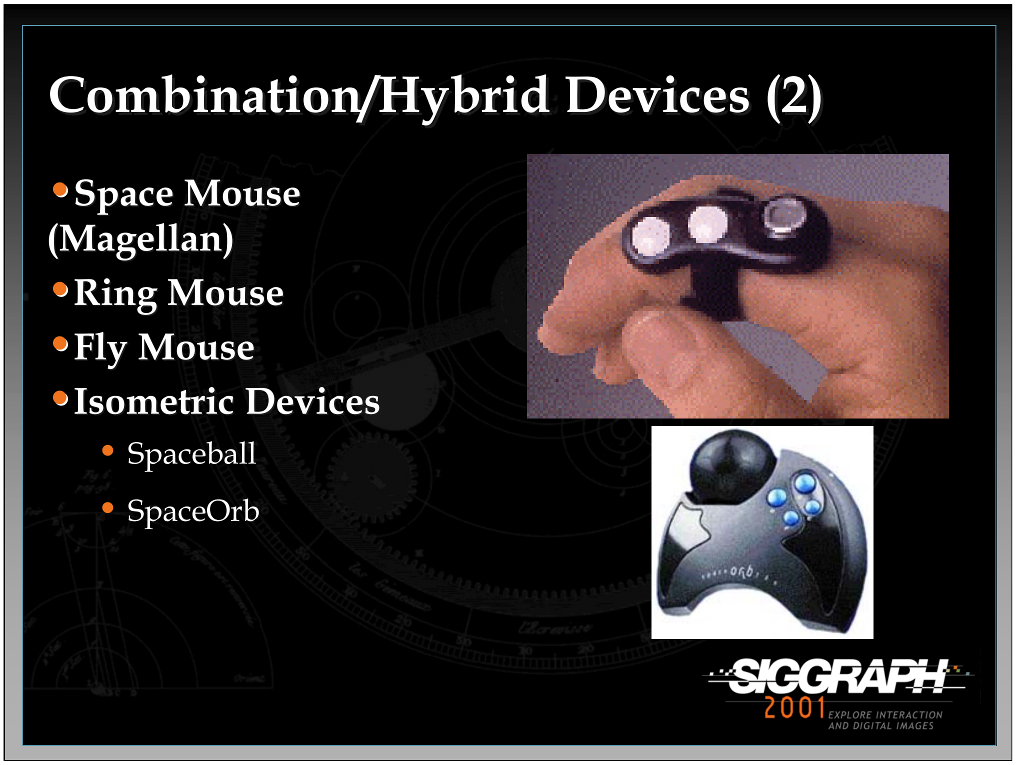

The Space Mouse (Magellan) is a 6 DOF input device originally designed for telerobotic manipulation. Slight pressure of the fingers onto the cap of the Magellan generates small deflections in X, Y, and Z, which can move objects in 3D space. With slight twists of the cap, rotational motions are generated. It also has a series of buttons which will generate discrete events. The Ring Mouse (top picture) is a small device worn on the user's finger which uses ultrasonic tracking. It also has two buttons for generating discrete events. The main advantages of this device is that it is wireless and inexpensive. The Fly Mouse is a 3D mouse that also uses ultrasonic tracking. This device has five buttons instead of two and also can be used as a microphone.

Another type of input devices are isometric which have a large spring constant so they cannot be perceptibly moved. Their output varies with the force the user puts on the device. A translation isometric device is pushed while a rotation isometric device is twisted. A problem with these devices is that users may tire quickly from the pressure they must apply in order to use them. The bottom figure is a picture of the SpaceOrb, an isometric device from Labtec priced at approximately forty dollars.

References:

www.spacemouse.com, www.labtec.com, www.pegatech.com
www.qualixdirect.com/html/3d_mouse_and_head_tracker.html

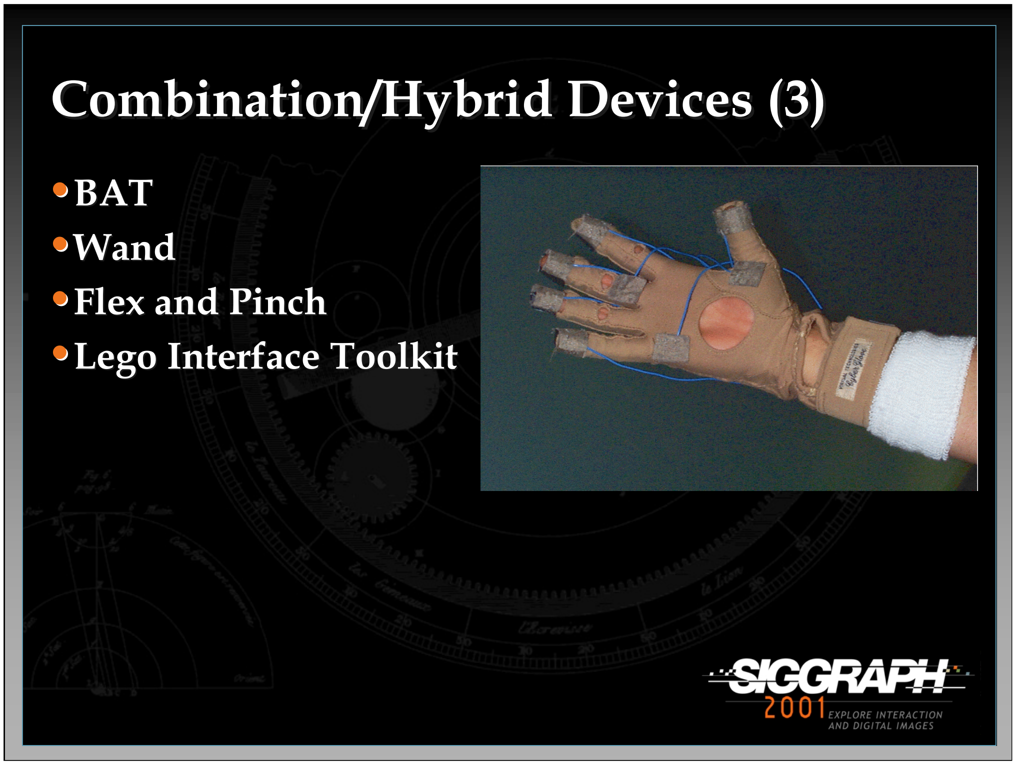

The BAT is a device that was developed by Colin Ware in the late 1980's. It essentially is just a tracking device with three buttons attached to it. It's similar to the other 3D mice mentioned in the previous slide except it is rather easy to build one with a few electrical components (provided you have the tracking device). The Wand is a device that is commonly seen in SSVR environments. It is simply a more elegant version of the BAT that is commercially developed. The Flex and Pinch input system is a custom built device which takes the functionality of the Pinch Glove system and combines it with the bend sensing technology of a data glove. The pinch buttons are made from conductive cloth and can be placed anywhere on the bend sensing glove. The Lego Interface Toolkit is a rapid prototyping system for physical interaction devices in immersive environments. It utilizes Lego bricks because they are easily obtained and support a variety of physical configurations.

References:

Ware, Colin and Danny R. Jessome. "Using the Bat: A Six Dimensional Mouse for Object Placement." Proceedings of Graphics Interface'88, 119-124.

LaViola, Joseph and Robert Zeleznik. "Flex and Pinch: A Case Study of Whole-Hand Input Design for Virtual Environment Interaction." Proceedings of the IASTED International Conference on Computer Graphics and Imaging '99, 221-225.

Ayers, Matthew and Robert Zeleznik. "The Lego Interface Toolkit." Proceedings of User Interface Software and Technology, 1996, 97-98.

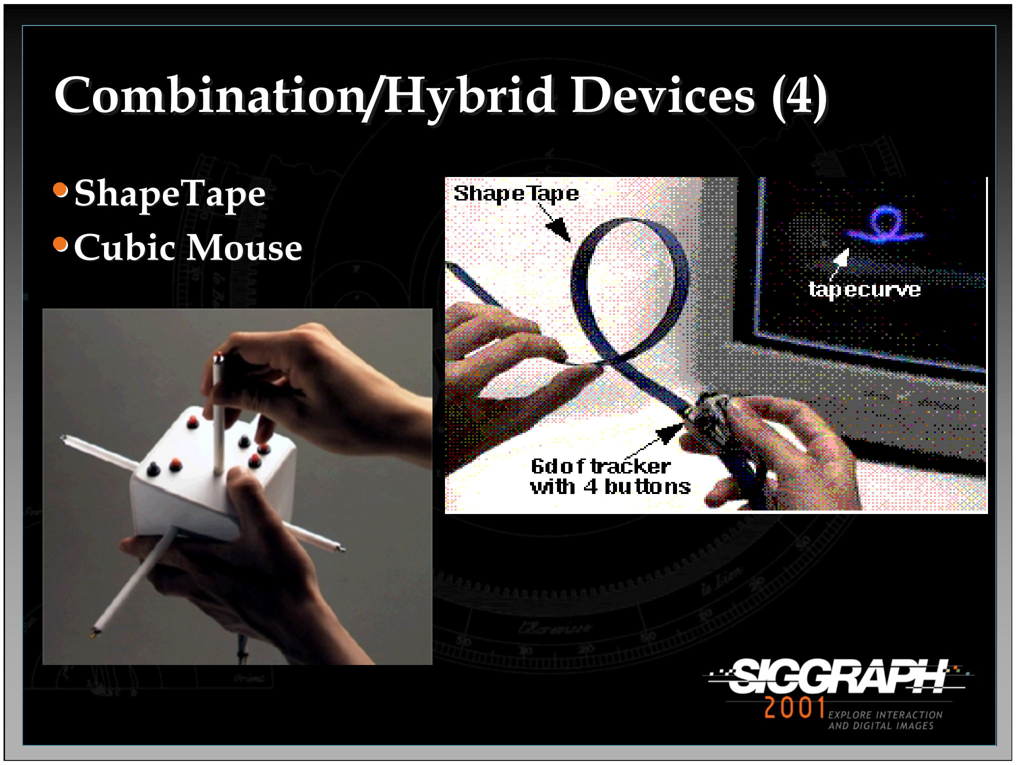

ShapeTape is a continuous bend and twist sensitive strip which encourages two-handed manipulation. A BAT is attached and the tool (shown in the figure on the right) is used for creating and editing curves and surfaces along with cameral control and command access. ShapeTape senses bend and twist with two fiber optic sensors at 6cm intervals.

The Cubic Mouse (shown in the figure on the left) is an input device developed at GMD that allows users to intuitively specify three-dimensional coordinates in graphics applications. The device consists of a box with three perpendicular rods passing through the center and buttons for additional input.

References:

Balakrishnan, Ravin, George Fitzmaurice, Gordon Kurtenbach, and Karan Singh. “Exploring Interactive Curve and Surface Manipulation Using a Bend and Twist Sensitive Input Strip” Proceedings of the 1999 Symposium on Interactive 3D Graphics, 111-118, 1999.

Frohlich, Bernd, John Plate. “The Cubic Mouse: A New Device for Three-Dimensional Input”, Proceedings of CHI2000, 526-531, 2000.

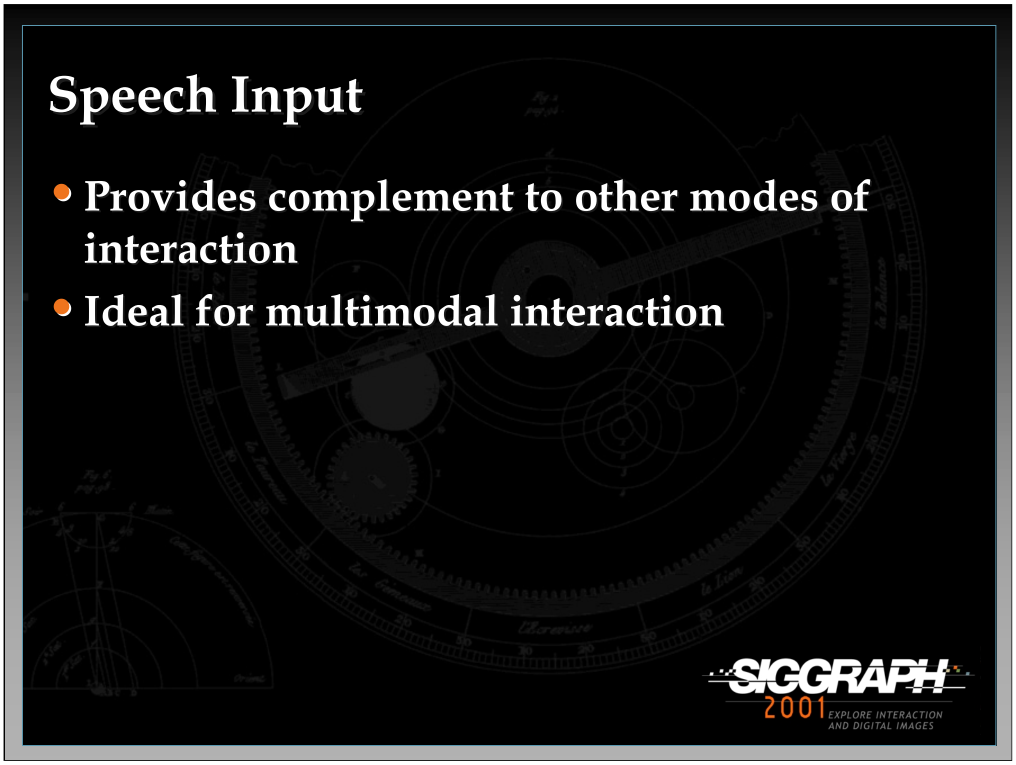

Speech input provides a nice complement to other input devices. As a result, it is a natural way to combine different modes of input (e.g. multimodal interaction) to form a more cohesive and natural interface. In general, when functioning properly speech input can be a valuable tool in virtual environment applications especially when both of the user's hands are occupied. There are many issues to consider when dealing with speech input besides what speech recognition engine to use. There are tradeoffs that must be made when dealing with speech input. An important issue is where the microphone is to be placed. Ideally, a wide area mike would be best so that the user does not have to wear a headset. Placing such a microphone in the physical environment could be problematic since it might pick up noise from other people or machines in the room. One of the big problems with using speech input is having the computer know when to and not to listen to the user's voice. Often, a user is conversing with a collaborator with no intention of issuing voice commands but the applications "thinks'' the user is speaking to it. This misinterpretation can be very problematic. One of the best ways to avoid this problem is to use an implicit or invisible push-to-talk scheme. A push-to-talk scheme lets the user tell the application when he/she is speaking to it or someone else. In order to keep the naturalness of the speech interface, we do not want to have to add to the user's cognitive load. The goal of implicit push-to-talk is to imbed the "push'' into existing interaction techniques so the user does not have the burden of remembering to signal the application that a voice command is about to be issued.

References:

LaViola J. Whole-Hand and Speech Input in Virtual Environments, Master's Thesis, Brown University, Dept. of Computer Science, December 1999.

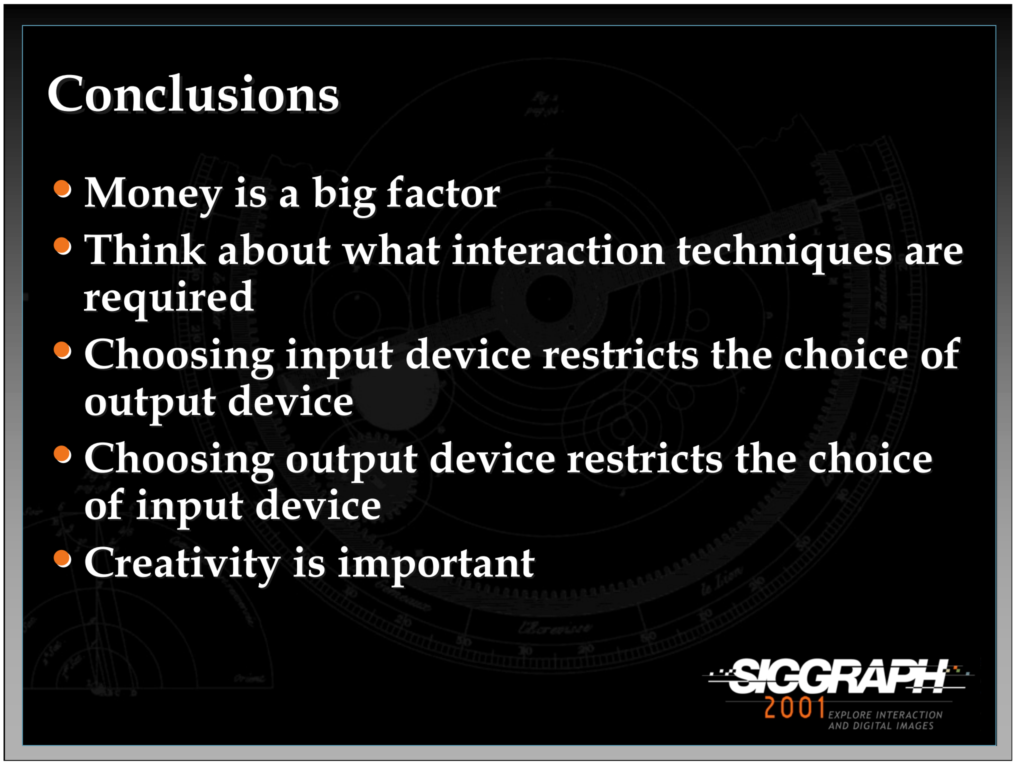

Obviously when choosing input and output devices for creating virtual environment applications and systems, money is a big issue. However, getting the most expensive I/O devices does not necessarily guarantee that the VE will be usable. In general, when selecting I/O device, think about what the user is going to be doing in the VE and what sorts of interaction techniques will be required. At that point, thinking about the physical devices that are best suited for the required techniques.

Finally, none of the input or output devices described in this lecture are perfect. As a result, there is a lot or research left to be done to develop better I/O devices. Creativity is important when thinking about them. If you can't find a commercially available device to suit you needs then build one that will.

General Reference:

Carolina Cruz-Niera, "Applied Virtual Reality." Course #14. Siggraph 1998.

Youngblut, C. R.E. Johnson, S.H. Nash, R.A. Wienclaw, and C.A. Will, "Review of Virtual Environment Interface Technology." Technical Report IDA Paper P-3186, Log:H96-001239. Institute for Defense Analysis. 1996.