Immersive Projection Technology for Visual Data Mining

Edward J. WEGMAN and Jürgen SYMANZIK

The PlatoCAVE, the MiniCAVE, and the C2 are immersive stereoscopic projection-based virtual reality environments oriented toward group interactions. As such they are particularly suited to collaborative efforts in data analysis and visual data mining. In this article, we provide an overview of virtual reality in general, including immersive projection technology, and the use of stereoscopic displays for data visualization. We discuss design considerations for the construction of these immersive environments including one-wall versus four-wall implementations, augmented reality, stereoscopic placement, head tracking, the use of LCD devices, polarized light stereo, voice control, and image synchronization.

Key Words: C2; CAVE; IPT; MiniCAVE; PlatoCAVE; Virtual reality; Visualization; VR; VRGobi.

1. INTRODUCTION

In Book 7 of *The Republic* (Lee 1987), Plato introduces the Allegory of the Cave. Written in the form of a narrative by Socrates in dialog with Glaucon, Plato's older brother, the Allegory of the Cave is ostensibly about the enlightenment of philosopher-kings by progressive revelation of truth. However, it is also a very evocative articulation of Plato's metaphysical ideas. In brief, Plato held that all objects we see in our ordinary existence are but pale reflections of real, but not accessible *ideal forms*. These ideal forms characterize the essence of the objects we see in real life. That is, the table we see is a imperfect version of some real, extant ideal table. Indeed all tables are imperfect versions of this ideal table. The Allegory of the Cave is an attempt to explain this idea.

> Imagine an underground chamber like a cave, with a long entrance open to the daylight and as wide as the cave. In this chamber are men who have been prisoners there since they were children, their legs and their necks being so fastened that they can only look straight ahead of them and cannot turn their heads. Some way off, behind and higher up, a fire is burning, and between the fire and the prisoners and above them runs a road, in front of which a curtain–wall

---

Edward J. Wegman is Professor, Center for Computational Statistics, George Mason University, MS 4A7, 4400 University Drive, Fairfax, VA 22030-4444 (E-mail: ewegman@gmu.edu). Jürgen Symanzik is Professor, Utah State University, Department of Mathematics and Statistics, Logan, UT 84322–3900 (E-mail: symanzik@sunfs.math.usu.edu).

Journal of Computational and Graphical Statistics, Volume 11, Number 1, Pages 1–26

> has been built, like the screen at puppet shows between the operators and their audience. . . . our prisoners could [not] see anything . . . except for the shadows thrown by the fire on the wall of the cave opposite them." (Lee 1987, pp. 256).

Plato argues that what we perceive as reality are as the shadows on the wall of the cave. We are constrained by the physical limitations of our bodies and senses from seeing the ideal form and see only the shadows of this deeper reality. We have often thought of this Allegory of the Cave when trying to develop visualization tools for viewing multidimensional objects. Indeed, we are constrained by our senses bound to the three- (four-) dimensional world from easily seeing into higher dimensions. Just as with Plato, we can see only two- or perhaps three-dimensional shadows (projections) of real higher dimensional objects. This Allegory of the Cave actually motivated many of the visualization environments that we created and which are the subject of this article.

An interesting feature of the Cave Allegory is that there are multiple prisoners who, while they cannot see each other, see the same images and who can interact verbally with each other. In other words, Plato's Cave was perhaps the first collaborative environment.

In Section 2 of this article, we provide a basic introduction, including a historic overview, on anaglyphs and stereoscopic displays, *Virtual Reality* (VR), the PlatoCAVE, other CAVEs, *Immersive Projection Technology* (IPT), and main design considerations underlying these systems. In Section 3, we present aspects of data visualization and exploration via stereoscopic displays and VR hardware. In Section 4 we focus on the MiniCAVE, a particular IPT environment developed at George Mason University. We finish with an outlook on future developments that might be useful for data visualization and visual data mining in Section 5. Some of the immersive environments described in this article have reached nominal cost levels, thus making them accessible for general academic and commercial usage.

2. BASICS OF ANAGLYPHS, STEREOSCOPIC DISPLAYS, VR, CAVES, AND IPT

2.1 ANAGLYPHS AND STEREOSCOPIC DISPLAYS

The desire to capture and display the three-dimensional real world in its entirety and not only as a two-dimensional projection dates back far into the Middle Ages where two pictures on wood constructions have been drawn, allowing each eye to view only one of these pictures. The idea to represent the three dimensions of the real world just on paper using anaglyphs first appeared in 1853. A German teacher, Wilhelm Rollmann, initially described the effect of stereoscopic graphics drawn in red and green colors that are looked at with the naked eye (Rollmann 1853a); that is, what is now called free-viewing stereoscopic images. Later the same year, Rollmann (1853b) described the effect of looking at such colored pictures using filter glasses of corresponding complementary colors. Eventually, this work has been judged by Vuibert (1912) and Rösch (1954) as the birth of anaglyphs. Independently from Rollmann, the French teacher Joseph Charles d'Almeida proposed in 1858 to use light of different colors to produce anaglyphs. However, the name anaglyphs was introduced by the French Ducos du Hauron not earlier than in 1891. Vuibert (1912) is one of the earliest books that entirely deals with anaglyphs.

More than two decades later, Köhler, Graf, and Calov (1938) succeeded in printing anaglyphs of high quality. Their so-called *Plastoreoskop-Verfahren* almost perfectly prevents the remainder of an additional noise picture. The main application of anaglyphs at that epoch was in geometry; see, for example, Graf (1938, 1941).

Rellensman and Jung (1939), Rellensmann (1940), and Linhard (1940) dealt with early applications of anaglyphs for mining and related applications. The Hungarian Pál is credited with the distribution of anaglyphs all over Europe in a variety of sciences. His books (Pál 1961, 1974) contain anaglyphs from fields such as geometry, mechanical engineering, architecture, chemistry, and spatial mathematical problems. Also the anaglyphs in Fejes Tóth (1965) were created by Pál.

The first textbook, according to its author, that deals only with anaglyphs is Mucke (1970). In this book, applications are given for ten subdisciplines within mathematics and natural sciences. Mucke and Simon (1966, 1967) focus entirely on anaglyphs for use in geometry. Schmidt (1977) is a late book of high quality anaglyphs dealing with the same topic.

Architecture is an obvious field for the use of anaglyphs. Probably one of the technically best anaglyph pictures printed so far is Schwenkel's drawing of the buildings constructed for the 1972 Olympic Games in Munich (Schwenkel 1972). During the 1960s, anaglyphs were introduced to chemistry. Examples are the textbooks of Holleman and Wiberg (1963) and Klages (1965). Anaglyphs have many potential uses in cartography, biology, medicine, and for CAD applications. One of the more recent textbooks that contain anaglyphs is Vince (1995).

The work of Burkhardt should be mentioned here, too. His work (Burkhardt 1963, 1972, 1974) covers technical aspects and problems of printed anaglyphs such as optimal colors, best filter glasses, and so on.

The introduction of colored computer monitors made it possible to simply transfer red-green anaglyphs from paper to the computer monitor. However, early restrictions were the limited number of colors that could be displayed on computer monitors (sometimes only eight) and difficulty finding matching filter glasses that do not leave an unwanted ghost image and otherwise do not remove too much brightness for just one eye. An intermediate step to modern shutter glass technology between 1986 and 1988 were the Tektronics LCD polarized monitors for both PCs and UNIX systems, where a second screen was mounted in front of the monitor to produce the stereoscopic effect.

With further progress in *cathode ray tube* (CRT) display technology (both in resolution and number of displayable colors) and speed of computer hardware, it eventually became possible to display and alternate two fully colored images on a computer monitor at a high frequency. With additional hardware—that is, shutter glasses that are synchronized with the computer monitor—it finally becomes possible to display the left image on the screen only when the left shutter glass is open and the right image only when the right shutter glass is open. When this alternation of left-eye and right-eye images is done at a high frequency, typically 120 Hz or above (i.e., 60 left-eye views and 60 right-eye views per second), the human brain combines the left and right images that are only partially visible into a three-dimensional, flicker-free image within the human brain. The use of such time-multiplexed stereoscopic displays is the most commonly used technique when we think of stereoscopic CRT displays these days.

The mathematics underlying anaglyphs and stereoscopic displays can be found in Hodges (1992) and Wegman and Carr (1993), for example. In addition, Hodges (1992) also described other, less frequently used stereoscopic display systems.

2.2 VR

Many different definitions of the term *virtual reality* (VR) can be found throughout the literature. Cruz-Neira (1993) summarized the following definitions:

- "VR is the body of techniques that apply computation to the generation of experientially valid realities." (William Bricken);
- "Virtual reality is the place where humans and computers make contact." (Ken Pimentel, Kevin Teixeira);
- "VR has to do with the simulation of environments." (Gregory Newby);
- "VR provides real-time viewer-centered head-tracking perspective with a large angle of view, interactive control and binocular display." (Daniel Sandin);
- "Virtual reality refers to immersive, interactive, multisensory, viewer-centered, three-dimensional computer generated environments and the combination of technologies required to build these environments." (Carolina Cruz–Neira);
- "An experience in which a person is surrounded by a three-dimensional computer-generated representation and is able to move around in the virtual world and see it from different angles, to reach into it, grab it, and reshape it." (Howard Rheingold); and
- "Virtual reality is a media to recreate the world in which we live and to create illusions of new and yet unknown worlds." (Anonymous).

Additional terms related to *virtual reality* are *artificial reality*, *augmented reality*, *virtual environments*, and *cyberspace*. Some people use these terms interchangeably for VR, while other people make clear distinctions between each of them.

Even though there exist different definitions of VR, there is little doubt that the origin of VR dates back to 1965 when Ivan Sutherland proposed the *ultimate display* (Sutherland 1966). Between 1966 and 1970 Sutherland also built and refined the *Sword of Damocles* (Sutherland 1968) at the MIT Lincoln Laboratory, at Harvard University, and at the University of Utah. This device is considered to be the first *head mounted display* (HMD). It consisted of two cathode ray tubes that were mounted alongside each of the user's ears and additional hardware that was suspended from the ceiling by a mechanical arm to measure the user's head position and orientation. In 1971 Frederick Brooks at the University of North Carolina developed the *GROPE–II System*, a molecular docking tool for chemists, which used the *Argonne remote manipulator* (ARM), one of the first force-feedback devices. A more sophisticated device, the *GROPE–III* molecular modeling system which was developed by Brooks and his students, was described by Ouh-Young et al. (1988). In 1985 Thomas Zimmerman at VPL Research designed the *DataGlove*, a device that is capable of measuring the degree to which each of the user's fingers is bent. Another VR device, the *binocular omni-orientation monitor* (BOOM) was commercialized in 1989 by Fake Space Labs. The BOOM is a small box that contains two *cathode ray tubes* (CRTs) which can be viewed through two eye holes. The box is attached to a mechanical arm that measures its position and orientation while the user moves it around to explore the virtual world.

A brief chronology of further events that influenced the development of VR can be found in Cruz-Neira (1993). A more detailed overview can be found in Pimentel and Teixeira (1995) or Vince (1995), for example.

2.3 THE PLATOCAVE

The virtual environments described in this and the next section began with emergence of virtual reality hardware in the late 1980s. Much of the initial work at George Mason University was based on trial and error without carefully thought-out design considerations, but based on the conviction that some of the VR technology must be useful for data visualization. In what follows, we describe the assemblage of commercial products into systems; this is not intended to endorse such products, but only to give instances of products which could be used.

Beginning in 1991, researchers at George Mason University acquired a number of virtual reality hardware components and began a series of experiments in visual data analysis (Wegman and Luo 1994; Wegman, Poston, and Solka 1996; Wegman, Luo, and Chen 1998; Wegman et al. 1999). Initially HMDs were purchased but they were unsuitable for extended use, primarily because they only had a low resolution and they were cumbersome to wear. They were also not conducive to group interaction and were quickly discarded as a technology for visual data analysis in favor of higher resolution, projection-based technology, leading to the PlatoCAVE.

The PlatoCAVE is an immersive display system installed in a room approximately 20 feet on a side. It consists of a Stereographics CRT projection system driven by a Silicon Graphics workstation. The original installation, dating from late 1991, was driven by a Silicon Graphics Crimson VGXT. This was subsequently replaced with a Silicon Graphics Onyx with Reality Engine^2 graphics, and in 1997, with a Silicon Graphics Onyx II with Infinite Reality Engine graphics. The Stereographics projection system is a reworked Electrohome CRT projector with capability for using Crystal Eyes shutter glasses for stereoscopic display. The green phosphor decay on the original Electrohome system is too slow to support the required 120 frames per second needed for the stereoscopic effect. Slow decay yields optical crosstalk between the left-eye image and the right-eye image and destroys the stereoscopic effect. The green phosphor tube is replaced with a less bright, but faster decaying green phosphor tube. The Crystal Eyes shutter glass are lightweight and contain active electronics which alternately obscure left and right eyes at a rate of 60 frames per second for each eye. They are controlled by an infrared sending unit which synchronizes the shutters with the projected image.

The images are projected on a wall of the PlatoCAVE and are approximately 15 feet in diagonal measurement and span the entire wall from floor to ceiling. The remaining three walls of the environment are painted in a charcoal gray to help focus attention on the bright wall. The system is augmented by five-channel audio. Typically the room is dimly lit by incandescent lamps on dimmers with just enough light to navigate. The effect is that the projection wall seems to drop away and yields a stereoscopic, three-dimensional, full-color image. Coupled with a suitable audio soundtrack, the effect can be quite stunning.

2.4 CAVEs AND IPT

In approximately the same time frame as the PlatoCAVE was being developed at George Mason University—independently and unknown to one another—Carolina Cruz-Neira and her colleagues developed a much more ambitious visualization environment at the Electronic Visualization Lab (EVL) of the University of Illinois in Chicago, known simply as the "CAVE" (Cruz-Neira et al. 1992; Cruz-Neira, Sandin, and DeFanti 1993; Cruz-Neira et al. 1993; Cruz-Neira 1995; Roy, Cruz-Neira, and DeFanti 1995). The abbreviation "CAVE" stands for *CAVE Audio Visual Experience Automatic Virtual Environment*.

The CAVE is a projection-based VR system where the illusion of immersion is created by projecting stereoscopic computer graphics into a cube of projection screens that (completely) surround the viewer. Many of the salient features of the CAVE and PlatoCAVE are the same; for example, the use of Crystal Eyes Stereographics' LCD shutter glasses. The CAVE initially was driven through multiple Silicon Graphics VGX workstations, each attached to a rear-projection display for each of its walls. A Silicon Graphics Personal Iris workstation served as a master controller for the entire system. Communication among workstations took place via Ethernet. The first CAVE environment was essentially a cube approximately 10 feet on a side. Using mirrors for folding the light path and rear projection, the CAVE projects stereoscopic images on the front and two side walls. Using direct down projection, the fourth projector is used for a stereoscopic image on the floor.

In addition to the three walls and floor, the CAVE typically features head-tracking of one user so that the computed viewpoint is dynamically adjusted according to the location of the tracked head within the 10^3 cubic foot volume. The dynamic head tracking allows the user to see his or her entire environment from the correct perspective and thus creates a compelling illusion of reality.

A detailed technical comparison of the CAVE with other display devices for VR such as CRT, BOOM, and HMD can be found in Cruz-Neira et al. (1992). The CAVE is an easy-to-learn, high-resolution VR interface that is superior to these devices in particular because of its full field-of-view, its visual acuity, and the lack of intrusion. It requires only very lightweight, unrestrictive equipment to be worn that does not make the user feel uncomfortable. Moreover, the CAVE allows multiple viewers to share the same virtual environment at the same time, enhancing the visual experience and making discussion possible while inside this environment. Overall, the CAVE is a very helpful tool for collaborative work. It might also be helpful for a new user to join a guide—that is, an expert navigator—in the CAVE and get introduced to the particular problem before exploring the virtual environment him/herself.

Carolina Cruz-Neira moved to Iowa State University in 1995 where she was involved in the development of a second, larger CAVE-like environment known as the "C2." During the last five years, the CAVE has been commercialized and many clone installations have been made worldwide; for example, three- and five-sided CAVEs at Fraunhofer-IGD at

Darmstadt, Germany (Unbescheiden 1998); the "CUBE" at the University of Stuttgart, Germany (Rantzau et al. 1998); the "CABIN" and the "CoCABIN" at the Universities of Tokyo and Tsukuba, Japan (Hirose et al. 1998); the six-surface, fully immersive CAVE at the Royal Institute of Technology, Sweden (Ihrén and Frisch 1999); the spatially reconfigurable "CABANA" at HRL Laboratories at Malibu, CA (Daily et al. 1999); and the six-sided "C6" at Iowa State University, Ames, IA (Bernard, Vance, and Cruz-Neira 1999).

The CAVE and its successors belong to *immersive projection technology* (IPT) systems where the user is visually immersed within the virtual environment. In addition to the "traditional" IPT systems such as HMD and BOOM—and CAVE—another recent addition to the family of IPT systems are desk-like *responsive workbenches* (RWBs) such as the "PanoramaBench" (Ramshorn 1998) and the "ErgoDesk" (Forsberg, LaViola Jr., and Zeleznik 1998). Comparisons and advantages of different IPT systems can be found in Bullinger, Riedel, and Breining (1997) and Eckel et al. (1997), for example. In recent years, the CAVE and RWB idea has spawned a conference known as the International Immersive Projection Technology Workshop (Bullinger and Riedel 1997; Cruz-Neira and Riedel 1998; Bullinger and Riedel 1999; Cruz-Neira, Riedel, and Rössler 2000).

2.5 DESIGN CONSIDERATIONS

Both the CAVE and the PlatoCAVE are visually impressive, powerful environments. The design of a collaborative environment, however, has implications on what features are desirable and, sometimes, the more ambitious environment is actually not the more desirable one.

2.5.1 One Wall Versus Four Wall

Although the multiwall CAVEs are fully immersive environments, one-wall installations such as the PlatoCAVE may nearly achieve the same effect. In the PlatoCAVE installation at George Mason University, the screen is about 15 feet in diagonal measurement and nine feet tall. Thus, the screen is approximately 12 feet across. Normal human peripheral vision subtends an angle of approximately 150° horizontally and 135° vertically (Wegman and Carr 1993, Section 6.6). A simple trigonometric calculation shows that by standing just under 6.4 feet from the screen, one's whole visual field will be occupied by the stereoscopic image. Standing in the middle of the room (about 10 feet away from the screen) still makes the stereoscopic image by all odds the dominant feature. Our conclusion is that very little is lost by restricting attention to a large one-wall projection screen. Indeed, stereo objects in peripheral vision can be distracting and perhaps be detrimental to group interaction. IMAX theatres are successful examples of huge (curved) projection screens for a large number of viewers. The movie *Captain Eo*, featuring Michael Jackson, is one example how the entertainment industry might use huge one-walled stereoscopic displays for large audiences.

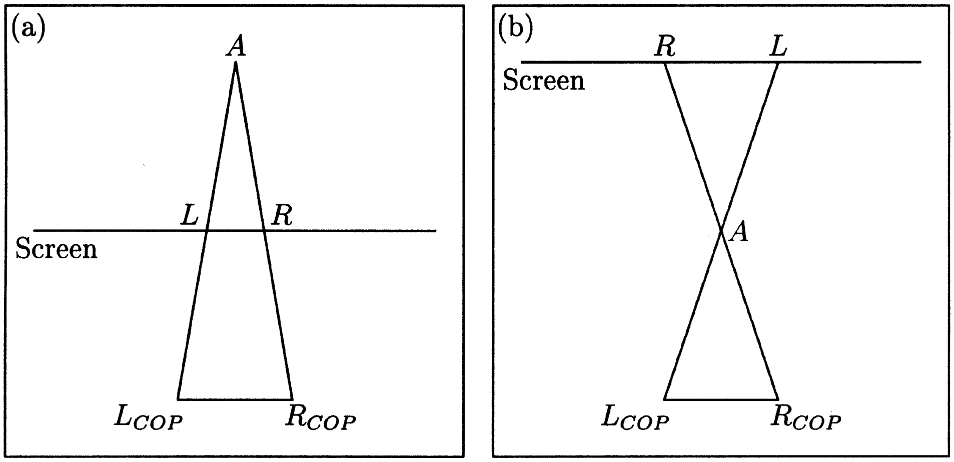

Figure 1. L_{COP} is the left center of projection (left eye) while R_{COP} is the right center of projection (right eye). In (a), A is the object viewed behind the screen. L and R are, respectively, the left and right eye images of A projected on screen. In (b), A is the object in front of the screen. Notice that in (a) and (b) there are two similar triangles, but the relative positions of L and R are interchanged in (a) and (b).

2.5.2 Augmented Reality

Both the CAVE and the PlatoCAVE share the desirable feature that the Crystal Eyes shutter glasses allow the viewer to see not only other participants in the immersive environment, but also allow the viewer to see his or her own body. This contrasts with virtual reality using HMDs or BOOMs. In the latter, so-called *avatars* must be used to represent other participants and iconic representations of body parts such as disembodied hands are used. The projection-based IPT environments are probably more correctly called *augmented reality* rather than *virtual reality* in the sense that not only does the viewer see the virtual environment, but he or she also sees the real people in that virtual environment. This is clearly a great asset in a collaborative environment.

2.5.3 Stereoscopic Placement

An interesting issue is the placement of the stereoscopic images. In the CAVE-like environments, the developers of applications have tended to place images in front of the screen, while in the PlatoCAVE environment, we have tended to place images behind the screen as shown in Figure 1. Figure 1(a) shows the placement of an object at point A behind the screen. Here L_{COP} is the left *center of projection* (COP), that is, the position of the left eye and R_{COP} is the right center of projection, that is, the position of the right eye. L and R are, respectively, the projections of A on the screen for the left eye and the right eye. The distance between L and R is the *stereoscopic disparity*. Notice that in Figure 1(a) there are two similar triangles, ALR and $AL_{\mathrm{COP}}R_{\mathrm{COP}}$. These similar triangles can be used to compute the stereoscopic disparity. Figure 1(b) shows the placement of an object at point A in front of the screen. Interestingly enough the ALR and the $AL_{\mathrm{COP}}R_{\mathrm{COP}}$ triangles are still similar and can also be used to compute the stereoscopic disparity. Indeed the formulas remain the same except for changes in sign. For example, Hodges (1992) and Wegman and Carr (1993, sec. 6.1), derived the formula in detail. However, here we are interested only in a qualitative description.

Two features are of significance. The left eye image of A and the right eye image of A depend on whether the object A is in front of the screen or behind the screen. They are simply interchanged depending on position. However, it is clear that the angle $L_{\mathrm{COP}}AR_{\mathrm{COP}}$ is much larger when the object A is in front of the screen. This angle is the *angular parallax* and it is well known that stereo images are harder to fuse with larger parallax angles (Wegman and Carr 1993, Section 6.2). A second phenomena involves synchronization of focus and parallax. Our eyes converge inwards as objects become closer. Their lenses also obviously focus more closely as objects become closer. These two effects are synchronized for normal vision. However, for stereoscopic displays, they must be somewhat decoupled. The lenses must focus on the screen while the parallax must be maintained for the virtual position of the object. The decoupling of the focus and parallax is one reason why some individuals have difficulty with stereoscopic displays. In any case, as objects are farther away, both the demand on the eyes for parallax compensation and for focusing distance become less. Thus, a strategy that places objects behind the screen makes the stereoscopic visualization much easier. As we shall see in the next subsection it also eases the criticality of the center of perspectivity. Thus, we argue that placement of the objects to be visualized near or behind the screen creates an easier visualization environment and enhances group interaction.

2.5.4 Head-Tracking

Head-tracking is usually accomplished by the placement of a magnetic sending unit on the Crystal Eyes shutter glasses and a corresponding sensing unit somewhere in the volume of the CAVE or PlatoCAVE environment. In the PlatoCAVE, we use a system called Flock-of-Birds although there are a number of similar units commercially available. In VR systems involving HMDs, it is useful to track head position in six degrees of freedom, that is, spatial x, y, z positions and x, y, z angles. In IPT environments only three degrees of freedom are required, that is, spatial x, y, z positions. In head-tracked environments, the COP is dynamically adjusted according to head location. The implication of this is that the viewer whose head is tracked will see exactly the right stereo perspective. However, all other viewers will see a possibly greatly distorted stereo perspective as shown in Figure 2. Objects will be substantially displaced for nonhead-tracked observers. In addition, there is normally substantial spatial distortion for the nonhead-tracked observers. Of course, straight lines will also not appear straight across the corners of projection walls in CAVE-like environments. Finally, as the tracked observer moves, the stereoscopic distortion will change for the nontracked observer even if the nontracked observer remains stationary. The only recourse in the head-tracked environment is for all observers to move around in a tight cluster.

The alternative to head-tracking is to select a nominal viewpoint and compute the stereoscopic image from the nominal viewpoint. Obviously, this approach must be implemented for IMAX theatres. In the PlatoCAVE, we typically select a position about the center of the room as a nominal viewpoint. In addition, we typically place the object behind the screen.

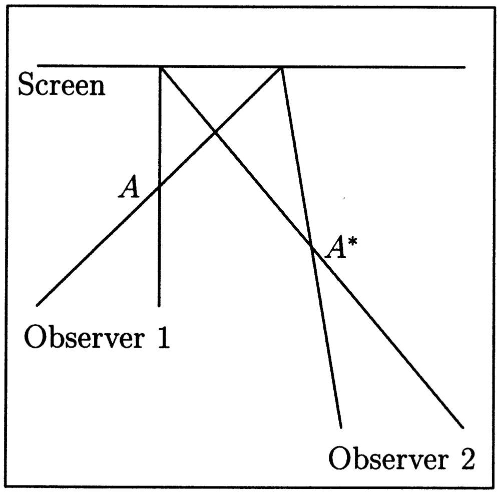

Figure 2. If Observer 1 is head-tracked, a virtual object in position A will be perceived by Observer 2 in position A. The scale of stereoscopic disparity is greatly exaggerated in this figure. However, Observer 2 will see a considerably distorted image.*

Everyone except an individual standing at the nominal viewpoint will, of course, experience some degree of distortion. However, by positioning an object behind the screen, the distortion tends to be minimized and, because the viewpoint is fixed, images remain stable for all observers independent of the movement of any one observer. Thus, for a collaborative immersive environment, placement of the image behind the screen and not implementing head-tracking are strategies which enable most effective group interaction.

3. DATA VISUALIZATION AND EXPLORATION VIA STEREOSCOPIC DISPLAYS AND VR HARDWARE

3.1 Dynamic Statistical Graphics

Dynamic statistical graphics enables data analysts in all fields to carry out visual investigations leading to insights into relationships in complex data. Dynamic statistical graphics involve methods for viewing data in the form of point clouds or modeled surfaces. Higher dimensional data can be projected into one-, two-, or three-dimensional planes in a set of multiple views or as a continuous sequence of views which constitutes motion through the higher dimensional space containing the data.

There is a strong history of statistical graphics research on developing tools for visualizing relationships between many variables. Much of this work is documented in videos available from the American Statistical Association Statistical Graphics Section Video Lending Library at http://www.bell-labs.com/topic/societies/asagraphics/library/index.html. Additional material on statistical graphics can also be found in journals such as *Journal of Computational and Graphical Statistics* and in *Computing Science and Statistics*, the proceedings from the Interface conferences. The following paragraphs serve only as a basic overview for readers unfamiliar with dynamic statistical graphics but they are not intended as a full introduction into this topic.

A video clip of the successive stages in a multidimensional scaling algorithm (Kruskal 1970) is one of the first examples how to apply dynamic statistical graphics. A second example by Chang (1970) shows an interactive search for a structured two-dimensional projection in five dimensions where three of the five dimensions are noise. PRIM-9 (Picturing, Rotation, Isolation and Masking in up to 9 dimensions), documented by Fisherkeller, Friedman, and Tukey (1974a, 1974b), is the landmark example of early dynamic statistical graphics. Projections formed the fundamental part of the visualization system and were complemented with isolation and masking. A good explanation of the importance of projection as a tool for visualizing structure in high-dimensional data can be found in Furnas and Buja (1994).

One major breakthrough in using projections for visualizing higher dimensions was made by Asimov (1985) in his work on the grand tour. The grand tour, further exploited by Buja and Asimov (1986), in an abstract sense shows a viewer all possible projections in a continuous stream (which could be considered to be moving planes through p-space). Several possibilities for "showing all possible projections" were explored in the original work, but the most successful method to arise from it is based on interpolating between random planes. Another common approach to displaying high-dimensional data can be found in Becker and Cleveland (1987) where data is plotted in a scatterplot matrix; that is, a matrix of pairwise scatterplots. Users can do linked brushing among the plots—that is, mark points with different symbols and colors—while this information is also immediately displayed in all related (linked) plots.

3.2 USE OF ANAGLYPHS AND STEREOSCOPIC DISPLAYS

Since the introduction of PRIM-9, most interactive and dynamic statistical graphics have been restricted to display at most two dimensions at a time. However, there have been some approaches to display statistical data in three or more dimensions. Stereoscopic displays and anaglyphs have been used within statistics by Carr, Littlefield, and Nicholson (Carr, Littlefield, and Nicholson 1983; Carr and Littlefield 1983; Carr and Nicholson 1985; Carr, Nicholson, Littlefield, and Hall 1986). In particular anaglyphs can be considered as an important means to represent three-dimensional pictures on flat surfaces. They have been used in a variety of sciences (described in Section 2.1 of this article) but they found only little use in statistics. One of the first implementations of red-green anaglyphs in statistical software was the "real-time rotation of three-dimensional scatterplots" in Mason Hypergraphics, described by Bolorforoush and Wegman (1988, p. 125). Independent from the work of Carr, Wegman, et al., interactive statistical anaglyph programs also have been developed by Hering, Symanzik, and von der Weydt (Hering and von der Weydt 1989; Hering and Symanzik 1992; Symanzik 1992; Symanzik 1993a, 1993b; Hering 1994).

Other than the major efforts by the previously described research groups, the literature reveals only little use of anaglyphs within statistics, although statistical datasets and problems often are well suited to be displayed using anaglyphs. Wegman and DePriest (1986) is one of the rare sources in statistics where anaglyphs are used in the papers of Banchoff (1986), Carr et al. (1986), and Gabriel and Odoroff (1986). Moreover, Wegman and DePriest (1986) seems to be the first *statistical* reference where colored (red-green) anaglyphs have been published in print. Huber (1987, p. 450), also noticed the lack of anaglyphs within statistics. ("*Statisticians still lag behind other scientists in their use of stereo pairs. [...] The statisticians' efforts (e.g., the plates in Wegman and DePriest 1986) come late and pale in comparison.*")

Instead of further using anaglyphs, researchers at George Mason University in the late 1980s were intrigued by the possibility of using time-multiplexed stereoscopic displays to visualize more complex structures. Tektronics made a stereoscopic monitor which used a liquid crystal polarizing filter in front of the screen. Coupled with polarizing glasses and driven by a Silicon Graphics 4D–120/GTX workstation, it was possible to produce true stereoscopic, full color images on a workstation monitor as early as 1989. However, while this was effective for a single observer, it was not particularly useful as a collaborative environment.

Four major applications within statistics which capitalize on stereoscopic displays have been developed at George Mason University. *ExplorN* includes stereoscopic rotating and scalable scatterplots animated with a grand tour. *Mason Ridge* includes stereoscopic density plots and density contours and features use of rendering and transparency to visualize complex abstract functional surfaces. *3-D MRI* uses stereoscopic displays, rendering and transparency to visualize solid voxel data such as MRI, PET, and ultrasound data. *Stereo SkyFly* is an adaptation of SGI software that allows a flythrough of an elevation database using texture mapping on the triangulated surfaces derived from the elevation data. The adaptation allows not only a stereoscopic view, but also real-time access to the elevation database.

3.3 USE OF VR AND IPT HARDWARE

Iowa State University and George Mason University probably are the two leading academic institutions with respect to the use of VR and IPT hardware for data visualization. While Section 4 is devoted to one major development at George Mason University, the MiniCAVE, this section summarizes the work done at Iowa State University.

In addition to the work on VR-based data visualization conducted at these two universities, independent work also has been conducted elsewhere, for example, at Georgia Tech and the Delft Technical University, The Netherlands, resulting in the "Virtual Data Visualizer" (van Teylingen, Ribarsky, and van der Mast 1997), and at the University of South Carolina, using the *Virtual Reality Modeling Language* (VRML) for VR applications on the World Wide Web (Rossini and West 1998).

The use of Iowa State University's C2 for statistical visualization is based on the framework of three-dimensional projections of p-dimensional data, using as a basis the methods developed and available in XGobi (Swayne, Cook, and Buja 1998). XGobi uses multiple linked windows to display scatterplots of multiple views of high-dimensional data. Some of XGobi's main features are interactive and linked brushing, identification, scaling, univariate, bivariate, trivariate plot modes, and the grand tour. This approach has a history of development throughout dynamic statistical graphics research, as discussed in Section 3.1. The implementation of these (or similar) features in the C2 resulted in VRGobi. The main difference between XGobi and VRGobi is that the XGobi user interface is rather like a desktop with pages of paper whereas VRGobi is more like having the whole room at the user's disposal for the data analysis.

VRGobi and the statistical visualization in the C2 have been extensively explored and documented in the literature (Swayne, Cook, Kohlmeyer, and Cruz-Neira 1996; Symanzik et al. 1997; Cook et al. 1997; Cook et al. 1998; Nelson, Cook, and Cruz-Neira 1998, 1999; Cook 2001). Main developers of VRGobi, over time, were Di Cook and Carolina Cruz–Neira, with major contributions by Brad Kohlmeyer, Uli Lechner, Nicholas Lewin, Laura Nelson, and Jürgen Symanzik. Additional information on VRGobi can be found at http://www.icemt.iastate.edu/research/visualization/statistics/index.html.

The initial implementation of VRGobi contains a three-dimensional grand tour (Asimov 1985; Buja and Asimov 1986; Buja, Asimov, and Hurley 1989). The basic idea of a grand tour is that a continuous sequence of projections is shown to the user. In XGobi the sequence contains two-dimensional projections of the data, and in VRGobi the projections are three-dimensional. Continuity of motion allows the user to mentally make connections between different views of the data, and taking arbitrary three-dimensional projections can expose features of the data not visible in one-dimensional or two-dimensional marginal plots.

One of the most difficult developments for VRGobi was the user interface (and not the statistical display components). Although it is relatively simple to create popup menus that allow to select colors and symbols for brushing in a desktop environment, designing an appealing and operational three-dimensional interface for the C2 was a real challenge. Eventually, four main components make up VRGobi: the viewing box, the three-dimensional control panel, the variable spheres, and possibly a map view.

The viewing box is delimited by a wireframe cube with the variable spheres at the vertices. The viewing box can be reshaped or rotated within the C2 environment. One part of the viewing box is a speed pole that allows the user to manually control the speed of motion during a tour with a gloved hand that is displayed as a floating hand and forearm.

The points displayed within the viewing box are painted in different colors and glyph types. The palette of colors available for painting is shown to the right of the viewing box. There are four different symbol types (sphere, cube, pyramid, and star) and the size of the symbol can be changed by using a resizing pole next to the color palette. There are two different shapes of brush that can be used (sphere and rectangular prism) and that can be resized easily.

Variable spheres show the contribution to the projection of each variable and also allow the user to select and deselect variables from being included in the tour. Interaction with the control tools is reinforced by sound feedback.

A three-dimensional map view allows the user to explore data in its spatial context within VRGobi, similar to the ArcView/XGobi link (Cook, Majure, Symanzik, and Cressie 1996, 1997) for the desktop.

IPT environments are remarkably different from display devices that are commonly available for data analysis. They extend beyond the small gain of one more dimension of viewing space, to being a completely defined "real" world space. In VRGobi, the temptation is to grab the objects or climb a mountain in the map view and to step aside when a point approaches our face during the grand tour. The objects surround the viewer and it is possible to walk through the data.

3.4 EXPERIMENTS COMPARING IPT AND SCREEN-BASED VR PERFORMANCE

Certainly, at some time in the near future, one should conduct larger studies to determine statistically (and significantly) whether VR and IPT in fact are more powerful than desktop–based software tools with respect to efficiency of discovery and Gestalt synthesis of statistical data. So far, only few experiments have been conducted where human performance in VR/IPT environments has been compared with human performance on a workstation display. Unfortunately, most of these experiments consisted of only a few human participants due to cost and time constraints.

In Nelson et al. (1998, 1999), experiments were conducted on structure detection—that is, visualization—and ease of interaction. With only 15 human test subjects, one cannot expect to obtain statistically significant results. However, at least these experiments showed that there was a clear trend that the test subjects performed considerably better on visualization tasks in the C2 than with XGobi on the workstation display. In contrast, interaction tasks such as brushing, provided better results for the workstation. However, subjects with some limited VR experiences already performed considerably better on the interaction tasks in the C2 than subjects with no prior VR experience, suggesting that there is some learning needed to effectively use the VR hardware.

In Deisinger, Cruz-Neira, Riedel, and Symanzik (1998) experiments were conducted with subjects using HMDs, monitors, and IPT environments. Again, with only 18 participants in two experiments (6 and 12 participants), no statistically significant result has been obtained. However, the IPT environment gave inexperienced users the best feeling of immersion and was liked most by all test participants.

Bullinger, Riedel, and Breining (1998) report on perception issues using HMDs, BOOMs, and IPT environments, in particular on experiments to guess dimensions of objects using these three types of VR hardware. While the average measured performance was slightly better for BOOMs than for IPT environments (followed by HMDs), 71% of the test subjects subjectively ranked their performance in the IPT environment best, again supporting the better subjective perception of IPT environments. While measured performance in IPT environments currently does not always relate to the subjectively anticipated performance (some problems, for example, arise due to imprecise VR tracking devices; see Czernuszenko, Sandin, and DeFanti 1998), one of the research goals in VR is the design of adequate and better usable 3D-interfaces to overcome some of the current hardware restrictions.

Dede, Salzman, Loftin, and Ash (2000) described *Project ScienceSpace*, a collection of three artificial worlds implemented at the PlatoCAVE and other VR hardware at George Mason University. NewtonWorld deals with laws of motion. In MaxwellWorld, users build electrostatic fields and, in PaulingWorld, users learn about molecular structure and chemical bonding. Dede and colleagues conducted extensive series of experiments with physics instructors and high school students. Overall, they found "immersive VR technology promising for learning complex science" (Dede et al. 2000, pp. 404). Based on their findings with respect to complex physical scenarios and the promising results reported by Nelson et al.

(1998, 1999) on statistical applications, one could easily hypothize that a larger study would also show that complex statistical applications would considerably benefit from the use of IPT environments when compared to standard desktop-based tools.

4. THE MINICAVE ENVIRONMENT

Certainly the CAVE and its successors and even the PlatoCAVE are costly affairs and beyond the reach of most academic and industrial groups. In addition, the CRT-based projection systems tend to be relatively dim and notoriously temperamental. Simply the process of heating up causes their convergence to deteriorate so that reconvergence and refocusing are required periodically throughout a session. The hardware typically is not portable so that the components must be permanently installed. With this in mind, researchers at George Mason University considered the possibility of converting the CRT-based projection system to a LCD-based projection system.

The decision to do so was cemented with the release of MATLAB 5.0 in 1996. This software was installed on a then state-of-the-art 200 megahertz Pentium Pro running Windows 95 and also on a Silicon Graphics Onyx with Reality Engine^2. We ran the benchmarks on both computers and, much to our surprise, neither machine dominated. A $3,000 Pentium Pro personal computer held its own against a $120,000 SGI workstation. Coupled with the fact that the Crystal Eyes shutter glasses had Windows NT drivers available, we launched on a project to downsize the PlatoCAVE into what we had designated as the MiniCAVE. Whereas the hardware for the MiniCAVE installation cost initially about $175,000, a full CAVE environment including mechanical supports for the screens, projectors, and mirrors, the projectors and the computers price out closer to $1,000,000. In addition, a full CAVE environment demands considerably more space than the MiniCAVE.

4.1 DESCRIPTION OF THE MINICAVE ENVIRONMENT

The MiniCAVE is co-located with the PlatoCAVE in the same room at George Mason University, but it is PC-based. Initially we began with a 333 megahertz Pentium II machine running Windows NT. After an unfortunate loss, this machine was replaced by a dual 450 megahertz Pentium III currently running Windows NT 4.0. The present machine uses an Intergraph Intense 3D AGP graphics card with 32 megabytes of memory. The card supports OpenGL and α-channel. The commodity chips such as those in the Intel series have in many senses overtaken the specialty RISC-chips found in workstations and certainly the commodity graphics cards compete favorably with the specialty graphics processor once available only through such vendors as SGI. The SGI–based applications mentioned in Section 3.2 such as ExplorN were based on the proprietary SGI GL standard, but the release of the OpenGL standard makes it relatively simple to port such applications to a PC environment. We were able to do this with the ExplorN and the Stereo SkyFly applications and using the Windows NT drivers, we were also able to integrate the Crystal Eyes shutter glasses into the PC environment.

The monitor/projector end proved to be somewhat more difficult. Most monitors designed for PCs are unable to cope with the high frame rates required for the Crystal Eyes shutter glasses. Initially we were forced to use a SGI monitor attached to the PC. However, we are currently using a Dell monitor quite successfully. We were able to connect the NT-based PC to the CRT projection system which made the PlatoCAVE essentially transparent to the driving computer. The attempt to use LCD projectors with the Crystal Eyes technology met with no success. We have since used polarized light LCD-based projectors. A description is given in the following.

4.2 DESIGN CONSIDERATIONS FOR THE MINICAVE

Obviously many of the same considerations outlined for the CAVE and the PlatoCAVE carry over directly to the MiniCAVE environment. However, the transition to PC-based computers carries additional considerations and, incidentally, additional opportunities.

4.2.1 LCD Devices

LCD (and so-called DLV) projectors have considerable appeal. They are inexpensive, bright, portable, and rugged. However, as they are typically configured, they are unsuitable for time-sequential stereoscopic projection such as required by the Crystal Eyes technology. As mentioned earlier, in CRT-based projectors, the brightness of a phosphor begins to decay immediately after the electron has struck the phosphor. The implication is that by the time the electron beam scans the first line of a CRT, the brightness of the last line has already decayed below normal perceptual levels. Thus, as the shutter glasses switch from right eye to left eye, the right-eye image has already decayed, there is no optical crosstalk, and the stereoscopic illusion is preserved. It is only the persistence in the visual system coupled with the extremely high frame rate that allows us to see flicker-free stereoscopic images.

In contrast, LCD systems have no need to decay. Thus, as a right-eye image is scanned, the appropriate pixel will be turned on and remains turned on until the same pixel is scanned again. This lack of decay is one reason the LCD projectors are correspondingly brighter than the CRT projectors. The implication of this, however, is that as shutter glasses switch from right eye to left eye, the entire right eye-image is still visible until the pixel is scanned. Thus, both right-eye and left-eye images are available for a portion of the time the left eye sees the image. Similarly both images are available for a portion of the time the right eye sees the image. The consequence due to persistence in the visual system is that independent of the shutter glasses, both eyes simultaneously see the left-eye and the right-eye images and the stereoscopic illusion is destroyed.

One obvious solution to the above problem is to induce an artificial decay into the LCD projection system. This has the advantage of restoring the stereoscopic illusion, but has the disadvantage of causing the projector to be dark for the majority of the frame. The disadvantage in terms of brightness of the images would be considerable. Such a scheme would require a rework of the electronics of the projector. Demand for stereoscopic displays using time-sequential shutter glasses does not seem to be sufficient to encourage projector makers to make such changes. An alternate solution is to use polarized light stereo. Incidentally although LCD projectors typically do not sync to frame rates much above 90 frames per second, this is not an inherent limitation of LCD technology. The Crystal Eyes glasses are LCD-based and sync to such high frame rates.

4.2.2 Polarized Light Stereo

Using polarized light stereo is the alternate solution we have used. Polarizing the light from the projector cuts down substantially on the projector's brightness, but still retains the other advantages of LCD projectors. A company called VREX sells an all-in-one stereoscopic polarized light projector. However, an alternate solution is to use two projectors (for each wall) each with a polarizing filter. The latter solution has the advantage of doubling the relative brightness. Polarization carries with it its own advantages and disadvantages. Polarizing glasses are inexpensive, passive, and lightweight, a decided advantage in a group-user environment. However, ordinary painted surfaces such as we had been using in the PlatoCAVE are diffuse reflecting surfaces which destroy the polarization. Thus, a special screen is required. While not terribly expensive, it is a slight added cost. The polarizing screens add to the brightness as light lost through diffusion is now reflected back to the viewer. However, the brightness is more dependent on the viewing position. In the VREX system, the polarization is vertical for one channel and horizontal for the other. This means eyes must be maintained in a horizontal plane. When the head rotates, more optical crosstalk is seen. Even with the head in a perfectly level situation, there is somewhat more optical crosstalk than with the Crystal Eyes shutter glasses. In short, there are pros and cons to both systems. LCD/polarized light stereo is brighter, more rugged, and less expensive than CRT/shutter glasses stereo.

A small word on rear projection systems is in order. Screens for rear projection of polarized light stereo are available. However, they tend to be considerably more expensive than what one might use for rear screen projection for a CRT projector. We have not implemented a four-wall MiniCAVE at this point. However, our plans have certainly been to scale the physical size of a four-wall MiniCAVE from a 12-foot cube to perhaps a 6-foot cube. Such a four-wall MiniCAVE would certainly offer the possibility of being relative portable. The smaller scale would make the requirement for head tracking less severe since the ability to be distant from the nominal center of projectivity would be greatly reduced. The higher cost of polarized light rear projection screens would also suggest a greatly reduced size would be in order.

4.2.3 Speech Recognition

Standard desktop computing has used the so-called desktop metaphor. The screen is essentially viewed as a virtual desktop with windows viewed as virtual pieces of paper. The mouse control becomes essentially an extension of the user's hands for moving the pieces of paper about and popping one or the other to the top of the heap. The pull-down menus are a simple extension of pointing. The desktop metaphor is useful in a single user PC or workstation environment because the screen is essentially a two-dimensional surface. However, the environments such as CAVE, PlatoCAVE, or MiniCAVE are essentially three-dimensional environments and the practicality of the desktop metaphor fails.

In a three-dimensional environment attempts have been made to port the desktop metaphor to the three-dimensional setting. Interaction with menus has been attempted by the use of data gloves. Whether the menus are two or three dimensional, the extension seems forced and cumbersome. Disembodied two-dimensional tool bars floating in a three-dimensional image are unnatural. Moreover, as often as not, the method for interacting with them is by means of a data glove which has the same tracking mechanism as the head tracking discussed in Section 2.5.4. There is a latency in tracking which implies that hand position sensed by the viewer and hand position as computed by the computer are not precisely the same. The effect is that one often misses the tool bar as one tries to press one of its buttons. Furthermore, unlike with the mouse, there is no tactile feedback and even sound feedback is not always sufficient. This is a fairly unsatisfactory situation. In particular for VRGobi in the C2, where a large amount of software development time has been spent on the user interface, it is very difficult to select a different brush color or resize the brush. Simple commands such as "Select green brush color" or "Resize brush to 80%" would help considerably to simplify the use of VRGobi.

For these reasons we started to incorporate voice-control into the MiniCAVE environment. The idea to develop voice (and gesture) interactive systems has been exploited for more than 20 years. One of the earlier systems was the "Put That There" system developed at MIT (Schmandt and Hulteen 1981). A basic, nontechnical introduction to speech recognition can be found in Markowitz (1996) or Rodman (1999), for example. Katagiri (2000) and DeMori and Suen (1985) are examples of collections of highly technical research papers on speech recognition. These latter two books provide some deeper insight into the mathematical and statistical aspects related to speech recognition.

Much of the technical detail how to use speech recognition in the MiniCAVE is contained in Wegman et al. (1999) which we will not repeat here. A short sketch is in order, however. There are two principal commercially available speech recognition software products, one by IBM and a second by Dragon. Although we experimented with both, we ultimately settled on Dragon Dictate. Both programs support multiple users and a fairly large vocabulary and do not require any knowledge on the underlying mathematical and statistical foundations of speech recognition. However, the IBM product is intended for continuous speech and actually suffers in performance when used for single, isolated words. Both products require some training and neither performs well out of the box. In general, however, we were interested in a more limited vocabulary and a broader array of speakers. This is true because when using the software in the MiniCAVE environment, we were interested in a limited vocabulary of single-word utterances. We linked the voice control to the Stereo SkyFly application and required only such utterances as "engage", "faster", "stop", "reverse", "left", "right", and so on. This scheme has been quite successful and works not only with American accents, but also German, Chinese, and South American accents. It is a simple leap to see how this could be applied to data analysis, data mining, and visualization software.

4.2.4 Synchronization

The following discussion applies only to multiwall CAVE-like environments where one computer only controls the image on one wall. Temporal synchronization is necessary to align images generated by autonomous, multiple computers so that the viewer is confronted with a continuous display which mimics the real world. Synchronization is needed at two levels. First, synchronization must be achieved so that images displayed by the projector are in sufficiently close temporal alignment so that blending of the images is achieved as perceived by the human visual system. This synchronization requires that the images displayed by each projector be no more than 1/100th of 1 second delayed from fastest to slowest image. We refer to this type of synchronization as “image lock.” Second, when using Crystal Eyes shutter glasses, synchronization between the time-sequential images for left eye and right eye is required so that all projectors display left-eye information simultaneously and similarly display right-eye information simultaneously. The required synchronization is within approximately 1/150th of 1 second. We designate this type of synchronization as “stereo lock.”

A synchronization signal generator is achieved by subprocesses running under a multiprocessor operating system on multiple independent computers communicating via Ethernet or similar computer networking scheme with speed capabilities of at least two megabits per second. One of the independent computers for the system is designated as the “master” and the others are designated as the “slaves.” The stereo lock is achieved by the master computer broadcasting a message via the computer network connection to each of the slave computers indicating which of the left-eye or right-eye images are to be displayed. This message needs to contain only a single bit of information plus routing overhead which is limited to a single packet of information. A packet containing 64 bytes or 512 bits would be available in less than 3/10,000th of 1 second on a two megabit per second computer network, easily within the 1/150th of 1 second requirement for stereo lock. The image lock synchronization works by having each slave computer reporting to the master computer when the slave computer has finished computing its current frame. Until each slave (and master) have completed computing the corresponding current frame, all computers display and redisplay the previous frame. When the master computer has received messages from each slave computer that the next frame is computed, and when the master computer itself has completed the next frame computation, the master computer broadcasts a signal to all slave computers to display the next frame. The next frame packet is similar in size to the stereo lock packet so that switching to the next frame can occur within the same 3/10,000th of 1 second time scale. The computation time of individual frames may vary depending on complexity of the image from 1/15th of one second to 1/150th of one second.

4.3 PRESENT STATE OF THE MINICAVE

The one-wall MiniCAVE with speech recognition has been implemented on a dual 450-megahertz Pentium III machine at George Mason University. We use the polarized light LCD projector with both front and rear projection. We have ported the Stereo SkyFly application to the NT environment and enhanced its capabilities to read in real elevation databases. (The original two-dimensional SGI application created a synthetic elevation database on the fly. The latter is easier to achieve since the synthetic database depends only on CPU and active memory, while real elevation databases require reading from a disk.) A version of ExplorN has also been ported to the MiniCAVE environment. In addition, a commercial version of visual data mining software which has the working name of CrystalVision has been developed. CrystalVision also runs in the MiniCAVE environment.

5. FINAL THOUGHTS

Much of the motivation for the PlatoCAVE and the MiniCAVE environments at George Mason University has been to provide a collaborative environment which facilitates group dynamics. We have given more than 100 demonstrations in our environment to groups ranging widely from medical doctors, technically oriented statisticians, engineers, potential donors to the university, prospective high school students and their parents, and groups of handicapped individuals. The success of the environment as a collaborative tool has been thoroughly demonstrated. Software systems capable of supporting stereoscopic images such as *ExplorN*, *Mason Ridge* and *3-D MRI* have been used extensively and lead to an excellent technical interaction and often animated discussions.

Because of the low cost of a one-walled MiniCAVE, approximately $6,000 for the computer and $13,000 for the polarized light stereo projectors, this technology is feasible on a much broader scale than more elaborate immersive systems such as the CAVE. Although this is not implemented at this stage, a natural future development is the linkage of multiple MiniCAVEs via the Internet. A handicap of course is that part of the advantage of augmented reality aspects—for example, seeing other participants—is lost and the introduction of avatars would be inevitable. However, the synchronization technology mentioned in Section 4.2.4 would be applicable and would lead to a reasonable speed even with PC-based computing.

In addition to our eyes and our voice, we might want to use some additional senses, that is, ears and hands, in future IPT environments for visual data mining. Early work to present information via sound has been conducted (Bly 1981) and might be worth revisiting in a CAVE-like environment. And certainly force-feedback devices and haptic devices might be used to explore data such as higher dimensional densities. Other VR interfaces are expected to be widely used in the near future (van Dam 1997) and might also become a factor how to explore data more effectively in a VR environment.

Cook (2001) lists three fields, “environmental studies, especially data having a spatial component; shape statistics; and manufacturing quality control,” that would most benefit from IPT environments. Certainly, recent experimental desktop links of virtual reality and visualization software with spatial statistical applications such as the links between ViRGIS and $\mathrm{RA_3DIO}$ with XGOBI (Symanzik, Pajarola, and Widmayer 1998; Schneider, Stamm, Symanzik, and Widmayer 2000) would benefit considerably when being conducted in an IPT environment. In addition to Cook’s fields, we think that medical, genetic, and biological statistical data would also considerably benefit when being explored in an IPT environment. Certainly, data from different military scenarios might also be explored in a CAVE-like environment.

It is worth mentioning in closing that, while part of the motivation for the original PlatoCAVE came from Plato’s Allegory of the Cave, another more contemporary literary work also motivated the development of our environments, Gene Roddenberry’s *Star Trek*. Indeed, the view screen on the bridge of the various Enterprises and Voyager has similar capabilities to the PlatoCAVE including voice controlled interaction. Similarly, AT&T Labs have installed so-called Power Walls which, although not stereoscopic, have about the same group dynamic character to them.

ACKNOWLEDGMENTS

The work of Wegman was supported by the Army Research Office under Grant DAAG55-98-1-0404, by the Office of Naval Research under Grant DAAD19-99-1-0314 administered by the Army Research Office, by the National Science Foundation under a Group Infrastructure Grant DMS-9631351, and the Defense Advanced Research Projects Agency under Agreement 8905-48174 with The Johns-Hopkins University. The original installation of the PlatoCAVE was funded by a grant from the Office of Naval Research. Upgrading to an environment with the Silicon Graphics Onyx with RE^2 graphics was supported by a DURIP grant awarded from the Army Research Office. The most recent upgrades and the development of the MiniCAVE were funded by matching funds to ARO and NSF grants.

Symanzik contributed to the VRGobi development at Iowa State University as a research assistant from 1996 to 1997 and to the MiniCAVE development at George Mason University from 1997 to 1999. His work at Iowa State University was partially supported by an EPA grant under cooperative agreement # CR822919 and his work at George Mason University was partially supported by the NSF through the Group Infrastructure Grant DMS-9631351.

Technical development of the PlatoCAVE and the MiniCAVE owes much to many individuals. Principal among these are Qiang Luo and Xiaodong Fu. Also contributing substantially to the development of the MiniCAVE were P. Vandersluis, A. Dzubay, R. Wall, R. Moustafa, F. Camelli, N–A. Khumbah, and Y. Zhu.

We thank our many co-authors of conference papers and journal articles underlying this work, in particular Dan Carr, Di Cook, Carolina Cruz–Neira, and Qiang Luo. A lot of the material presented in this current paper goes back to our collaborators over the last 15 years. Thanks are also due to the anonymous reviewers, associate editor, and editor of *JCGS* for their helpful comments to improve this article.

CITED SOFTWARE

CrystalVision: Wegman, E. J., Luo, Q., and Fu, X. (2000), copyright (c) 2000, a commercial package for data visualization running on PC's and UNIX, based on ideas originally implemented in ExplorN.

ExplorN: Carr, D. B., Luo, Q., Wegman, E. J., and Shen, J. (1992), copyright (c) 1992, a UNIX package for Silicon Graphics workstations incorporating scatterplot matrices, stereo ray glyph plots, parallel coordinates, and the d-dimensional grand tour.

Mason Hypergraphics: Wegman, E. J. and Bolorforoush, M. (1988), a PASCAL program for IBM RT under the AIX operating system and for MS–DOS machines, with red–green anaglyphs, considered to be the predecessor of Mason Ridge and ExplorN.

Mason Ridge: Luo, Q. and Wegman, E. J. (1991), copyright (c) 1990, 1991, a UNIX package for Silicon Graphics workstations for two- and three-dimensional density rendering using stereoscopic displays, transparency and lighting models and for ridge estimation.

Stereo SkyFly: Luo, Q. and Fu, X. (1998), a UNIX package for Silicon Graphics Workstations and a Windows NT package allowing a stereoscopic flythrough of an elevation database using texture mapping on triangulated surfaces.

3-D MRI: Luo, Q. and Wegman, E. J. (1993), copyright (c) 1993, a UNIX package for Silicon Graphics workstations for reconstructing two-dimensional MRI slices into a three-dimensional voxel image with the capability of rendering isodensity contours in transparent stereoscopic displays.

[Received October 1999. Revised May 2001.]

REFERENCES

Asimov, D. (1985), “The Grand Tour: A Tool for Viewing Multidimensional Data,” *SIAM Journal on Scientifc and Statistical Computing*, 6, 128–143.

Banchoff, T. F. (1986), “Visualizing Two-Dimensional Phenomena in Four-Dimensional Space: A Computer Graphics Approach,” in *Statistical Image Processing and Graphics*, eds. E. J. Wegman and D. J. DePriest, New York: Marcel Dekker, pp. 187–202.

Becker, R. A., and Cleveland, W. S. (1987), “Brushing Scatterplots,” in *Dynamic Graphics for Statistics*, eds. W. S. Cleveland and M. E. McGill, Belmont, CA: Wadsworth and Brooks/Cole, pp. 201–224.

Bernard, J., Vance, J., and Cruz-Neira, C. (1999), “The Making of the C6,” in *3. International Immersive Projection Technology Workshop*, eds. H.-J. Bullinger and O. Riedel, Berlin: Springer, pp. 11–21.

Bly, S. (1981), “Presenting Information in Sound,” in *Proceedings: Human Factors in Computer Systems*, March 15–17, 1982, Gaithersburg, MD, ACM, pp. 371–375.

Bolorforoush, M., and Wegman, E. J. (1988), “On Some Graphical Representations of Multivariate Data,” in *Proceedings of the 20th Symposium on the Interface between Computing Science and Statistics*, eds. E. J. Wegman, D. T. Gantz, J. J. Miller, Alexandria, VA: American Statistical Association, pp. 121–126.

Buja, A., and Asimov, D. (1986), “Grand Tour Methods: An Outline,” *Computing Science and Statistics*, 17 , 63–67.

Buja, A., Asimov, D., and Hurley, C. (1989), “Methods for Subspace Interpolation in Dynamic Graphics,” Technical Memorandum, Bellcore.

Bullinger, H.-J., Riedel, O., and Breining, R. (1997), “Immersive Projection Technology—Benefits for the Industry,” in *1. International Immersive Projection Technology Workshop*, Berlin: Springer, pp. 13–25.

——— (1998), “Perception in Different IPT-Systems,” in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, Ames, IA: Iowa State University (CD).

Bullinger, H.-J., and Riedel, O. (eds.) (1997), *1. International Immersive Projection Technology Workshop, 14./15*, Berlin: Springer.

——— (eds.) (1999), *3. International Immersive Projection Technology Workshop, 10/11*, Berlin: Springer.

Burkhardt, R. (1963), “Untersuchungen zur Verbesserung von Anaglyphenbildern,” *Bildmessung und Luftbildwesen*, 31, 190–195 (in German).

——— (1972), “Die farbliche Bemessung von Anaglyphenbrillen,” *Die Farbe*, 21, 26–32, (in German).

——— (1974), “Studie über Anaglyphen–Druckfarben,” *Die Farbe*, 23, 161–172 (in German).

Carr, D. B., and Littlefield, R. J. (1983), “Color Anaglyph Stereo Scatterplots—Construction Details,” in *Proceedings of the 15th Symposium on the Interface between Computer Science and Statistics*, ed. J. E. Gentle, New York: North-Holland, pp. 295–299.

Carr, D. B., Littlefield, R. J., and Nicholson, W. L. (1983), “Color Anaglyph Stereo Scatterplots—Construction and Application,” in *1983 Proceedings of the Section on Statistical Computing*, Alexandria, VA: American Statistical Association, pp. 255–257.

Carr, D. B., and Nicholson, W. L. (1985), “Evaluation of Graphical Techniques for Data in Dimensions 3 to 5: Scatterplot Matrix, Glyph, and Stereo Examples,” in *1985 Proceedings of the Section on Statistical Computing*, Alexandria, VA: American Statistical Association, pp. 229–235.

Carr, D. B., Nicholson, W. L., Littlefield, R. J., and Hall, D. L. (1986), “Interactive Color Display Methods for Multivariate Data,” in *Statistical Image Processing and Graphics*, eds. E. J. Wegman and D. J. DePriest, New York: Marcel Dekker, pp. 215–250.

Chang, J. (1970), *Real-Time Rotation*, ASA Statistical Graphics Video Lending Library (contact: dfs@research.att.com).

Cook, D. (2001), “Virtual Reality: Real Ponderings,” *Chance*, 14, 47–51.

Cook, D., Cruz-Neira, C., Kohlmeyer, B. D., Lechner, U., Lewin, N., Nelson, L., Olsen, A., Pierson, S., and Symanzik, J. (1998), "Exploring Environmental Data in a Highly Immersive Virtual Reality Environment," *Environmental Monitoring and Assessment*, 51, 441–450.

Cook, D., Cruz-Neira, C., Lechner, U., Nelson, L., Olsen, A., Pierson, S., and Symanzik, J. (1997), "Using Dynamic Statistical Graphics in a Highly Immersive Virtual Reality Environment to Understand Multivariate (Spatial) Data," in *Bulletin of the International Statistical Institute, 51st Session Istanbul 1997, Proceedings Book 2*, pp. 31–34.

Cook, D., Majure, J. J., Symanzik, J., and Cressie, N. (1996), "Dynamic Graphics in a GIS: Exploring and Analyzing Multivariate Spatial Data Using Linked Software," *Computational Statistics: Special Issue on Computeraided Analysis of Spatial Data*, 11, 467–480.

——— (1997), "Dynamic Graphics in a GIS: More Examples Using Linked Software," *Computers and Geosciences: Special Issue on Exploratory Cartographic Visualization*, 23, 371–385; paper, CD, and http://www.elsevier.nl/locate/cgvis.

Cruz-Neira, C. (1993), "Virtual Reality Overview," *SIGGRAPH '93 Course Notes #23*, pp. 1–18.

——— (1995), "Projection-based Virtual Reality: The CAVE and its Applications to Computational Science," PhD thesis, University of Illinois at Chicago.

Cruz-Neira, C., Leigh, J., Papka, M., Barnes, C., Cohen, S. M., Das, S., Engelmann, R., Hudson, R., Roy, T., Siegel, L., Vasilakis, C., DeFanti, T., and Sandin, D. J. (1993), "Scientists in Wonderland: A Report on Visualization Applications in the CAVE Virtual Reality Environment," in *IEEE 1993 Symposium on Research Frontiers in Virtual Reality*, pp. 59–66.

Cruz-Neira, C., and Riedel, O. (eds.) (1998), *2nd International Immersive Projection Technology Workshop*, May 11–12, 1998, Ames, IA: Iowa State University (CD).

Cruz-Neira, C., Riedel, O., and Rössler, A. (eds.) (2000), *4th International Immersive Projection Technology Workshop*, June 19–20, 2000, Ames, IA: Iowa State University (CD).

Cruz-Neira, C., Sandin, D. J., and DeFanti, T. A. (1993), "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," in *ACM SIGGRAPH '93 Proceedings*, Anaheim, CA, pp. 135–142.

Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., and Hart, J. C. (1992), "The CAVE: AudioVisual Experience Automatic Virtual Environment," *Communications of the ACM*, 35, 64–72.

Czernuszenko, M., Sandin, D., and DeFanti, T. (1998), "Line of Sight Method for Tracker Calibration in Projection-Based VR Systems," in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, May 11–12, 1998, Ames, IA: Iowa State University (CD).

Daily, M. J., Sarfaty, R., Jerald, J., McInnes, D., and Tinker, P. (1999), "The 'CABANA': A Re-Configurable Spatially Immersive Display," in *3. International Immersive Projection Technology Workshop*, eds. H.-J. Bullinger and O. Riedel, Berlin: Springer, pp. 123–132.

Dede, C., Salzman, M., Loftin, R. B., and Ash, K. (2000), "The Design of Immersive Virtual Learning Environments: Fostering Deep Understandings of Complex Scientific Knowledge," in *Innovations in Science and Mathematics Education—Advanced Designs for Technologies of Learning*, eds. M. J. Jacobson and R. B. Kozma, Mahwah, NJ: Lawrence Erlbaum Associates, pp. 361–413.

Deisinger, J., Cruz-Neira, C., Riedel, O., and Symanzik, J. (1997), "The Effect of Different Viewing Devices for the Sense of Presence and Immersion in Virtual Environments: A Comparison of Stereoprojections Based on Monitors, HMDs, and Screen," in *Advances in Human Factors/Ergonomics (21B). Design of Computing Systems: Social and Ergonomic Considerations*, eds. M. J. Smith, G. Salvendy, and R. J. Koubek, Amsterdam: Elsevier Science, pp. 881–884.

DeMori, R., and Suen, C. Y. (eds.) (1985), *New Systems and Architectures for Automatic Speech Recognition and Synthesis*, Berlin: Springer.

Eckel, G., Göbel, M., Hasenbrink, F., Heiden, W., Lechner, U., Tramberend, H., Wesche, G., and Wind, J. (1997), "Benches and Caves," in *1. International Immersive Projection Technology Workshop*, eds. H.-J. Bullinger and O. Riedel, Berlin: Springer, pp. 59–63.

Fejes Tóth, L. (1965), *Reguläre Figuren*, Budapest: Akadémiai Kiadó, (in German).

Fisherkeller, M., Friedman, J. H., and Tukey, J. (1974a), *PRIM-9: An Interactive Multidimensional Data Display and Analysis System*, ASA Statistical Graphics Video Lending Library (contact: dfs@research.att.com).

——— (1974b), “PRIM-9: An Interactive Multidimensional Data Display and Analysis System,” Technical Report SLAC-PUB-1408, Stanford Linear Accelerator Center, Stanford, CA.

Forsberg, A. S., LaViola, Jr, J. J., and Zeleznik, R. C. (1998), “ErgoDesk: A Framework for Two- and Three-Dimensional Interaction at the ActiveDesk,” in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, Ames, IA: Iowa State University (CD).

Furnas, G., and Buja, A. (1994), “Prosection Views: Dimensional Inference Through Sections and Projections,” *Journal of Computational and Graphical Statistics*, 3, 323–385.

Gabriel, K. R., and Odoroff, C. L. (1986), “Illustrations of Model Diagnosis by Means of Three-Dimensional Biplots,” in *Statistical Image Processing and Graphics*, eds. E. J. Wegman and D. J. DePriest, New York: Marcel Dekker, pp. 257–274.

Graf, U. (1938), *Sphärische Geometrie, Trigonometrie der Ebene und Kartenentwürfe*, Leipzig: Quelle and Meyer (in German).

——— (1941), “Konstruierte Anaglyphen,” *Bildmessung und Luftbildwesen*, 16 , 59–66; published as a supplement to *Allgemeine Vermessungs–Nachrichten*, 53, (in German).

Hering, F. (1994), “Anaglyphs in Statistics—A Tool for Interactive Multivariate Data Analysis,” in *SoftStat ’93—Advances in Statistical Software 4*, ed. F. Faulbaum, New York: Gustav Fischer, pp. 277–283.

Hering, F., and Symanzik, J. (1992), “Anaglyphen 3D—Ein Programm zur interaktiven Anaglyphendarstellung,” Forschungsbericht 92/1, Fachbereich Statistik, Universität Dortmund, (in German).

Hering, F., and von der Weydt, S. (1989), “Interaktive Anaglyphendarstellungen als Hilfsmittel zur Analyse mehrdimensionaler Daten,” Forschungsbericht 89/7, Fachbereich Statistik, Universität Dortmund, (in German).

Hirose, M., Ogi, T., Yamada, T., Tanaka, K., and Kuzuoka, H. (1998), “Communication in Networked Immersive Virtual Environments,” in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, May 11–12, 1998, Ames, IA: Iowa State University (CD).

Hodges, L. F. (1992), “Tutorial: Time-Multiplexed Stereoscopic Computer Graphics,” *IEEE Computer Graphics and Applications*, 12, 20–30.

Holleman, A. F., and Wiberg, E. (1963), *Lehrbuch der anorganischen Chemie (57.–70. Auflage)*, Berlin: de Gruyter (in German).

Huber, P. J. (1987), “Experiences with Three-Dimensional Scatterplots,” *Journal of the American Statistical Association*, 82, 448–453.

Ihrén, J., and Frisch, K. J. (1999), “The Fully Immersive CAVE,” in *3. International Immersive Projection Technology Workshop, 10./11. May 1999, Center of the Fraunhofer Society Stuttgart IZS*, eds. H.-J. Bullinger and O. Riedel, Berlin: Springer, pp. 59–63.

Katagiri, S. (ed.) (2000), *Handbook of Neural Networks for Speech Processing*, Norwood, MA: Artech House.

Klages, F. (1965), *Einführung in die organische Chemie (2. Auflage)*, Berlin: de Gruyter, (in German).

Köhler, O., Graf, U., and Calov, C. (1938), *Mathematische Raumbilder*, Berlin: Dreyer (in German).

Kruskal, J. B. (1970), *Multidimensional Scaling*, ASA Statistical Graphics Video Lending Library (contact: dfs@research.att.com).

Lee, D. (translator) (1987), *Plato—The Republic*, London: Penguin Books.

Linhard, M. (1940), “Die räumliche Darstellung von physikalisch—chemischen Prozessen, insbesondere von Schmelz—und Reduktionsvorg. angen nach dem Anaglyphenverfahren,” *Metall und Erz*, 37, 228–231 (in German).

Markowitz, J. A. (1996), *Using Speech Recognition*, Upper Saddle River, NJ: Prentice Hall.

Mucke, H. (1970), *Anaglyphen Raumzeichnungen*, Leipzig: Teubner (in German).

Mucke, H., and Simon, H. (1966), *Anaglyphen zur darstellenden Geometrie: Zweitafelprojektion*, Berlin: Volk und Wissen (in German).

——— (1967), *Anaglyphen zur darstellenden Geometrie: Dreitafelprojektion, Axonometrie, Zentralperspektivische Kartenprojektion*, Berlin: Volk und Wissen Volkseigener Verlag (in German).

Nelson, L., Cook, D., and Cruz-Neira, C. (1998), “XGobi vs the C2: An Experiment Comparing Data Visualization in an Immersive Virtual Environment with a Workstation Display,” in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, Ames, IA: Iowa State University (CD).

——— (1999), "XGobi vs the C2: Results of an Experiment Comparing Data Visualization in a 3-D Immersive Virtual Reality Environment with a 2-D Workstation Display," *Computational Statistics: Special Issue on Interactive Graphical Data Analysis*, 14, 39–51.

Ouh-Young, M., Pique, M., Hughes, J., Srinivasan, N., and Brooks, F. P. J. (1988), "Using a Manipulator for Force Display in Molecular Docking," in *1988 IEEE International Conference on Robotics and Automation*, Philadelphia, PA, April 24–29, 1988, Washington, DC: Computer Society Press, pp. 1824–1829.

Pál, I. (1961), *Darstellende Geometrie in Raumbildern (1. Auflage der deutschen*, München: Lindauer, (in German).

——— (1974), *Raumgeometrie in der technischen Praxis*, Budapest: Akadémiai Kiadó (in German).

Pimentel, K., and Teixeira, K. (1995), *Virtual Reality through the New Looking Glass* (2nd ed.), New York: McGraw-Hill.

Ramshorn, C. (1998), "The PanoramaBench," in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, Ames, IA: Iowa State University (CD).

Rantzau, D., Frank, K., Lang, U., Rainer, D., and Wössner, U. (1998), "COVISE in the CUBE: An Environment for Analyzing Large and Complex Simulation Data," in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, Ames, IA: Iowa State University (CD).

Rellensmann, O. (1940), "Die Bedeutung des Anaglyphen-Raumbildes für das Grubenrisswesen im Erzbergbau," *Metall und Erz*, 37, 225–228 (in German).

Rellensmann, O., and Jung, H. (1939), "Das Anaglyphen-Raumbild als Hilfsmittel für die Veranschaulichung von bergmannischen Aufgaben," *Glückauf*, 75, 530–533 (in German).

Rodman, R. D. (1999), *Computer Speech Technology*, Norwood, MA: Artech House.

Rollmann, W. (1853a), "Notiz zur Stereoskopie," *Annalen der Physik und Chemie*, 89, 350–351 (in German).

——— (1853b), "Zwei neue Stereoskopische Methoden," *Annalen der Physik und Chemie*, 90 , 186–187 (in German).

Rösch, S. (1954), "100 Jahre Anaglyphen," *Die Farbe*, 3, 1–6 (in German).

Rossini, A. J., and West, R. W. (1998), "Virtual Reality and Statistical Research," *Computing Science and Statistics*, 29, 215–219.

Roy, T., Cruz-Neira, C., and DeFanti, T. A. (1995), "Cosmic Worm in the CAVE: Steering a High Performance Computing Application from a Virtual Environment," *Presence: Teleoperators and Virtual Environments*, 4, 121–129.

Schmandt, C., and Hulteen, E. A. (1981), "The Intelligent Voice-Interactive Interface," in *Proceedings: Human Factors in Computer Systems*, March 15–17, 1982, Gaithersburg, MD: ACM, pp. 363–366.

Schmidt, R. (1977), *Darstellende Geometrie mit Stereo-Bildern*, Berlin: Bauverlag (in German).

Schneider, M., Stamm, C., Symanzik, J., and Widmayer, P. (2000), "Virtual Reality and Dynamic Statistical Graphics: A Bidirectional Link in a Heterogeneous, Distributed Computing Environment," in *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'2000)*, Las Vegas, Nevada, June 26–29, 2000, Volume IV, CSREA Press, pp. 2345–2351.

Schwenkel, D. (1972), "Räumlich Perspektive Darstellung von Bauwerken durch Automatisch Gezeichnete Anaglyphenbilder," *Bildmessung und Luftbildwesen*, 40, 144–147 (in German).

Sutherland, I. E. (1966), "The Ultimate Display," in *Proceedings of the IFIP Congress 65*, New York City, May 24–29, 1965, Volume 2, ed. W. A. Kalenich, Washington, DC: Spartan Books, pp. 506–508.

——— (1968), "A Head-Mounted Three-Dimensional Display," in *AFIPS Conference Proceedings*, Volume 33, Part One, 1968, Fall Joint Computer Conference, Washington, DC: Thompson Book Company, pp. 757–764.

Swayne, D. F., Cook, D., and Buja, A. (1998), "XGobi: Interactive Dynamic Graphics in the X Window System," *Journal of Computational and Graphical Statistics*, 7, 113–130.

Symanzik, J. (1992), "Computerdarstellungen von Anaglyphen—Ein Hilfsmittel der multivariaten Statistik," Diplomarbeit, Fachbereich Informatik, Universit. at Dortmund, (in German).

——— (1993a), "Anaglyphen 3D—A Program for theInteractive Representation of Three-Dimensional Perspective Plots of Statistical Data," in *Information and Classification. Concepts, Methods and Applications. Proceedings of the 16th Annual Conference of the Gesellschaft für Klassifikation e. V.*, eds. O. Opitz, B. Lausen, and R. Klar, Berlin: Springer, pp. 384–389.

——— (1993b), "Three-Dimensional Statistical Graphics based on Interactively Animated Anaglyphs," in *1993*

Proceedings of the Section on Statistical Graphics, Alexandria, VA: American Statistical Association, pp. 71–76.

Symanzik, J., Cook, D., Kohlmeyer, B. D., and Cruz-Neira, C. (1996), "Dynamic Statistical Graphics in the CAVE Virtual Reality Environment," Working Paper for the Dynamic Statistical Graphics Workshop, Sydney, July 7, 1996. Department of Statistics, Iowa State University, Ames, IA, Preprint 96-10, available at http://www.math.usu.edu/~symanzik/papers/papers.html.

Symanzik, J., Cook, D., Kohlmeyer, B. D., Lechner, U., and Cruz-Neira, C. (1997), "Dynamic Statistical Graphics in the C2 Virtual Reality Environment," *Computing Science and Statistics*, 29, 41–47.

Symanzik, J., Pajarola, R., and Widmayer, P. (1998), "XGobi and XploRe Meet ViRGIS," in *1998 Proceedings of the Section on Statistical Graphics*, Alexandria, VA: American Statistical Association, pp. 50–55.

Unbescheiden, M. (1998), "Evaluation of VR-Technology: Applications using 3- and 5-Sided CAVES," in *2nd International Immersive Projection Technology Workshop*, eds. C. Cruz-Neira and O. Riedel, May 11–12, 1998, Ames, IA: Iowa State University, (CD).

van Dam, A. (1997), "Post-WIMP User Interfaces," *Communications of the ACM*, 40, 63–67.

van Teylingen, R., Ribarsky, W., and van der Mast, C. (1997), "Virtual Data Visualizer," *IEEE Transactions on Visualization and Computer Graphics*, 3, 65–74.

Vince, J. (1995), *Virtual Reality Systems*, Wokingham, UK: ACM Press/Addison-Wesley.

Vuibert, H. (1912), *Les Anaglyphes Géométriques*, Paris: Librairie Vuibert, (in French).

Wegman, E. J., and Carr, D. B. (1993), "Statistical Graphics and Visualization," in *Handbook of Statistics, Vol. 9: Computational Statistics*, ed. C. R. Rao, Amsterdam: North Holland/Elsevier Science Publishers, pp. 857–958.

Wegman, E. J., and DePriest, D. J. (eds) (1986), *Statistical Image Processing and Graphics*, New York: Marcel Dekker.

Wegman, E. J., and Luo, Q. (1994), "Visualizing Densities," Technical Report 100, Center for Computational Statistics, George Mason University, Fairfax, VA, available at http://www.galaxy.gmu.edu/stats/techrepo.html.

Wegman, E. J., Luo, Q., amd Chen, J. X. (1998), "Immersive Methods for Exploratory Analysis," *Computing Science and Statistics*, 29, 206–214.

Wegman, E. J., Poston, W. L., and Solka, J. L. (1996), "Immersive Methods for Mine Warfare," in *MASEVR '95: Proceedings of the Second International Conference on the Military Applications of Synthetic Environments and Virtual Reality*, pp. 203–218.

Wegman, E. J., Symanzik, J., Vandersluis, J. P., Luo, Q., Camelli, F., Dzubay, A., Fu, X., Khumbah, N.-A., Moustafa, R. E. A., Wall, R. L., and Zhu, Y. (1999), "The MiniCAVE—A Voice-Controlled IPT Environment," in *3. International Immersive Projection Technology Workshop*, eds. H.-J. Bullinger and O. Riedel, Berlin: Springer, pp. 179–190.