Virtual Environments Research at the Georgia Tech GVU Center

by

Larry Hodges
Jay Bolter
Elizabeth Mynatt
William Ribarsky
and
Ron van Teylingen

GIT-GVU-93-45
December 1993

Graphics, Visualization & Usability Center

Georgia Institute of Technology
Atlanta GA 30332-0280

VIRTUAL ENVIRONMENTS RESEARCH AT THE GEORGIA TECH GVU CENTER

**Larry F. Hodges, Jay Bolter, Elizabeth Mynatt,
William Ribarsky, and Ron van Teylingen**

**Graphics, Visualization & Usability Center
Georgia Institute of Technology
Atlanta, GA 30332**

1. The GVU Center

The Graphics, Visualization and Usability (GVU) Center was established at Georgia Tech in 1991 in recognition of the central importance of these three disciplines to the future growth of computing.

The key emphasis of the GVU Center is effective communication of information between computers and people, as well as use of the computer to facilitate communication between individuals. This is not the domain of a single discipline, but rather draws on many diverse fields. Accordingly, the GVU Center emphasizes an interdisciplinary approach to research and education, bringing together 30 faculty and over 100 graduate students from the College of Architecture; School of Civil Engineering; College of Computing; School of Industrial and Systems Engineering; Office of Information Technology; School of Literature, Communication and Culture; School of Mathematics; Multimedia Technology Lab; and School of Psychology.

The GVU Center provides an environment which promotes learning, research, and service. In our educational role, we teach the principles and methods of computer graphics, visualization, and usability to members of the academic community ranging from undergraduate students to graduate students and other faculty. Our service mission is carried out through the Scientific Visualization Lab, a joint undertaking with Information Technology (the campus-wide computer service), to provide state of the art hardware and software capabilities to the entire Georgia Tech community. Our research spans the areas of realistic imagery, computer-supported collaborative work, algorithm animation, medical imaging, image understanding, scientific data visualization, animation, user interface software, usability, adaptive user interfaces, multimedia, stereoscopic graphics, virtual environments, image quality, and expert systems in graphics and user interfaces.

By integrating these three missions together in a single unit, the Center is developing a highly interactive and collaborative environment where researchers unfamiliar with computer graphics can come for help in integrating scientific visualization into their research work, graphics and user interface experts and graduate students can share their knowledge with one another and find new and interesting problems on which to work, and students can learn in a melting pot of closely-related ideas and collaborations between researchers from multiple disciplines.

2. The Virtual Environments Group

Within the GVU Center thirteen faculty members are currently directing virtual environment (VE) research projects. This group carries on a wide variety of work that includes many different aspects of virtual environment research. Current interests and projects include basic research investigating issues such as annotation of virtual spaces, directional sound, auditory and human performance issues, visual performance issues, interaction techniques, navigation, and software development for virtual environments. Applications research includes the use of virtual environments in medicine, architecture, education, scientific visualization, and training.

Equipment available for virtual environment research includes a number of Silicon Graphics (from Indigos to Reality Engines) and Sun workstations, two Virtual Research head-mounted displays, an Ascension Technologies Flock of Birds™ tracking system, a Polhemus tracking system, a Virtual Technologies CyberGlove™, and stereoscopic shutter systems from StereoGraphics (CrystalEyes™), Tektronix and 3DTV.

In addition to hardware specific to work in virtual environments, the GVU center also provides shared facilities that are available for VE research support. These facilities include a usability lab with one-way glass windows, three video cameras, a digital video/audio mixer, color monitor and VHS recorder/player; and an animation lab with a read/write video disk, real-time scan conversion, and equipment for video and audio editing.

The primary software for support of the creation and implementation of virtual environments is the Simple Virtual Environment (SVE) Library (Verlinden, et al., 1993b). SVE is a device independent library that provides mechanisms and software tools for VE applications. SVE separates the definition of the hierarchical structure of objects in the virtual world from the functions of the application, allowing the application programmer to concentrate only on the details of the application. Simple applications are therefore simple to program. In addition, the modularity of the SVE library allows for easy incorporation of additional functionality and allows for rapid prototyping of virtual worlds.

3. Description of Four Current Projects

As indicated in section 2, the virtual environments group represents a broad set of interests, backgrounds and projects. Due to space limitations we cannot describe all of our projects here. Instead we will provide brief overviews of four current projects as representative of the scope of our interests.

3.1 The Phobia Project

In the Phobia Project, we are designing a virtual reality system that can be used by a therapist in the treatment of acrophobia, the fear of heights (American Psychiatric Association, 1987). Behavioral therapy of acrophobia has included exposing the subject to anxiety producing stimuli. These stimuli can be generated through a variety of modalities including imaginal--subject generates stimulus via imagination (Marks & Gelder, 1965) and *in vivo*--subject is exposed to natural height situations (Biran & Wilson, 1981). Based on an initial subjective evaluation of what types of height situations cause anxiety in a patient, a therapist using an *in vivo* approach would arrange therapy sessions in which the patient goes through a process of exposure and adjustment to those situations. Patients begin with less threatening situations first and gradually work their way up a hierarchy of more anxiety producing situations. For example, if the patient is afraid to look out the window of a high building, sessions might begin by looking through a third floor window with the therapist present. In subsequent sessions the patient might move up to a window on the tenth floor. Other common locations for *in vivo* therapy could be outside stairways, balconies, bridges, and elevators.

Based on the types of height situations used for *in vivo* stimulus, we have designed a number of virtual height situations (see figures 1 & 2). We have now begun, in conjunction with Dr. Barbara Rothbaum in the Section of Psychiatry at Emory University Medical School, patient therapy sessions using the virtual height situations. In addition to Dr. Rothbaum, a local psychologist (Chris Crowe, Ph.D.) who specializes in treatment of phobias and a U.S. Army psychiatrist (Capt. James Williford, M.D.) are serving as evaluators of our environments.

Subjects were recruited through questionnaires distributed to students taking introductory psychology and computer science classes at Georgia Tech, Emory University, and Georgia State University. The questionnaires contain questions to screen students for a circumscribed fear of heights consistent with the diagnosis of simple phobia (America Psychiatric Association, 1989). Subjects are randomly assigned to either a treatment group or a wiating list control group. Waiting list subjects are assessed at pre-treatment and then again after eight weeks. Treatment group subjects are assessed at pre-treatment, revceive eight weeks of therapy, and then are assessed again after treatment.

If our approach is effective, then there are a number of advantages that would make it attractive to clinicians. Like *in vivo* therapy, virtual reality can provide stimuli for patients who cannot imagine well. Unlike the therapist-assisted *in vivo* techniques, VR therapy will be performed within the confines of a room, thus avoiding public embarrassment and violation of patient confidentiality. Virtual environments add the advantage of greater control over multiple height stimulus parameters and allow the ability to isolate which virtual height parameters are essential in generating a phobic response. Virtual reality exposure could be used as an intermediate step in preparing patients for maintenance therapy involving self-directed *in vivo* exposure.

Although the immediate goal of this project is to assess the viability of using VR as a tool in the treatment of persons with acrophobia, we believe that we are also gaining valuable insights into the nature of human experience of virtual environments. The best-known attribute of virtual reality is the sense of presence or immersion within a computer-generated world that the user experiences. It is an appeal to this sense of presence that is used to distinguish virtual reality as something different from merely a multimedia system or an interactive computer graphics display.

Our understanding of presence, however, is still primarily anecdotal in nature. We have yet to rigorously explore basic questions about immersive environments and the nature of presence (Sheridan, 1992).

> How do we identify and measure the effectiveness of those elements of a virtual environment that create an immersive experience?
>
> Are there applications for which a sense of presence is a necessary ingredient? If so, how do we identify those applications? How are they different from applications for which a more traditional display system is just as effective?

Most applications of VR (games excluded) have had the primary goal of increasing or modifying a person's *intellectual understanding* of the structure or nature of an object. Examples include use of VR for architectural walkthoughs (Brooks, 1986), molecular modeling (Bergman, et al., 1993), and medical visualization (Bajura, et al., 1992). Each of these applications may be viewed as applying a new tool (VR) in an old way (increasing intellectual understanding). The immersion experience implies an involvement with a virtual environment that goes beyond and sometimes even counter to intellectual understanding.

A person with a simple phobia suffers from a persistent fear of an object or situation. Exposure to the specific phobic stimulus almost invariably provokes an immediate anxiety response. If the simple phobia is a fear of heights, an intellectual understanding of the physics of gravity does not necessarily lessen the anxiety. Increasing a person's understanding in this case includes not only intellect, but also deeper levels of understanding that affect emotions and perceptions. The success of virtual reality in lessening patients' anxiety would provide experimental evidence that the immersion experience can affect participants at this level and that this effect can be carried over into real world experiences. *(Project Director: Larry F. Hodges)*

3.2 Speech and Text Annotation in Virtual Environments

In the Acrophobia Project, virtual reality is used to create in the patient a sense of presence or environmental immersion. Another important potential use of virtual reality is to create models that may represent real world environments (e.g. an architectural walkthrough) or scientific visualizations (e.g. a three dimensional visualizations of molecular structure). In exploring such a model, the user may want to make notes about what he or she finds. These notes can be made separately on paper, but an annotation system would allow the user to integrate those notes directly into the virtual environment. The notes could be associated with the various objects being examined, and these notes could then be available to the original user at a later time or to colleagues and subsequent users. Annotations could be of two kinds. A speech annotation would be a digital recording of the user's voice. A text annotation would be ASCII text (words, numbers, equations) that the user enters from the keyboard. Both kinds of annotations are ways of including symbolic information in the environment and of associating that information with data objects.

There has been relatively little study of how symbolic information should be created, displayed, and edited in three-dimensional environments. Probably the best-known published work comes from a team at Xerox Parc: the 3D/Rooms interface (Macinlay, 1991; Robertson, 1991). There is also a system at IPSI in Darmstadt (Hemmje, 1993). Both the Xerox Parc and the IPSI projects are information visualizers for large-scale databases. Neither of these projects uses an immersive virtual environment; instead, they visualize information in projected 3D on a conventional videoscreen. Nevertheless, the 3D/Rooms interface provides a foundation for further work. The interface has pioneered such objects as Cone Trees and Perspective Walls, in which verbal elements are deployed so that the user can see at a glance the relationships among the elements.

Annotation is visually simpler than information visualization, because with annotation we simply have to place icons or markers in the virtual space already defined by the data model. The user determines where the annotation marker should be place. In the GVU lab we have already started work this year on two prototypes, developed in collaboration with the Technical University of Delft. One prototype is a Speech Annotator, which allows the recording and placement of short spoken messages in a virtual environment, such as a model of the city of Atlanta. The messages are represented in the model by an iconic marker (such as a double pyramid). In one prototype, the annotations consist of comments or information about the various buildings and streets in Atlanta. Later, a user can browse through the city, select the markers, and playback their content (Verlinden, 1993). The second prototype, the World Processor, displays textual annotations in the form of a three-dimensional hypertext concept map (Verlinden, 1993a). This display is based upon the 2D concept maps used in some current hypertext systems, such as Storyspace and Sepia (Bernstein, 1991; Streitz, 1992). The World Processor's three-dimensional display is also similar to the Cone Tree display in 3D/Rooms. Ultimately the World Processor may also support hypertextual links between text units that may be widely separated in the virtual space.

In the next phase of our work, we are going to extend and continue development of annotation techniques in virtual environments. Our first goal is to create a more effective and user-friendly interface for speech annotation. The issues that must be explored include how to make and edit annotations and how to select and play them back. Making annotations should be easy for the user: it should be possible to make an annotation without interrupting the user's train of thought. The user should be able to annotate both various data objects and various locations in the virtual environment. The user may also wish to distinguish among various kinds of annotations. For example, in browsing the virtual environment, the user may want to see all the annotations made in the last week or all those made by a particular colleague. It should be possible to convey such information by the color or shape of the annotation marker. It may also be useful to have a time stamp or author name appear on the marker. In some respects these speech annotations resemble links in a hypertext or hypermedia application (Gloor, 1991; Lai, 1991; Bernstein, 1991; Streitz, 1993).

Our second task will be to create a similar annotation system for textual notes. Here the user would see the verbal text displayed on the screen above, below, or beside the object. We have already suggested that the speech annotations might show time or author stamps. In fact, the system can go much further and include sentences or paragraphs of text that describe the appropriate object. Creating textual annotations would require an editing facility, again similar to the editing facilities in some hypertext systems, but this editor will have to function in three dimensions--or at least display text on a two-dimensional pallet located in a 3D space. Current head-mounted displays do not typically have adequate resolution for effective text annotation. But the technologies here are improving. In The GVU lab, we have already begun experimenting with deploying text in 3D environments on a videoscreen. These experiments can eventually be transferred to immersive VR as the head-mounted display technology improves. We will also be able to apply lessons learned in the speech annotation experiments to the text annotation project as it evolves.

Finally, we hope to test the appropriateness of each annotation technique. Speech and text annotations will likely meet different needs. It may turn out, for example, that speech annotation makes it easier to record ideas and therefore that the user will make longer and more explicit annotations. But there may also be occasions when visual text is more effective. Some users may prefer speech annotation, while other prefer text. We wish to test both forms of annotations both anecdotally and through controlled experiments with interested users (e.g. engineers or scientists working with three dimensional models). The first comparative experiments will need to be done using 3-D on the videoscreen, rather than an immersive virtual environment, because text annotations may not be legible. We anticipate that subsequent experiments over the life of the project will involve immersive testing as well.

In general, we see these experiments in annotation as growing out of browsing and searching. We suspect that annotations themselves can serve as navigational tools. Speech annotation markers serve as signposts of significant locations in the environment. In the case of textual annotations, not only can the marker serve not only as a signpost, but the content of the annotation can also be searched to place the user at a relevant location in the environment. *(Project Director: Jay Bolter)*

3.3 Nonspeech and Directional Sound

In this research, our group is exploring the use of nonspeech and directional sound for improving the usability of virtual environments. Nonspeech sound refers to the use of specially designed auditory icons which convey information about objects and events in the virtual world. Directional sound refers to the creation of auditory cues which can be localized by the user giving the impression of sound coming from a certain location in space. Both of these techniques (nonspeech sound and directional sound) have been explored by our group in other, non-VR applications (Mynatt, 92; Burgess, 92). Our work with applying these techniques to virtual environments is still primarily in the design phase. In this work, our goal is to test the use of these techniques in improving the user's performance and sense of immersion in a variety of virtual environments.

Since the user can hear sounds coming from many directions and distances, the inclusion of spatial sound in a virtual environment supports the user's belief that the virtual environment is a real three-dimensional world . The first and still one of the few VR sites to use genuine spatial audio is the NASA Ames Research Center. To provide spatial audio in real-time, researchers and engineers at Ames developed a special-purpose signal processing engine called the Convolvotron (Wenzel, et al., 1988). Due to the expense of the Convolvotron, many researchers avoided exploring the uses of spatial audio for several years. Today, however, the emergence of several new spatial audio systems based on low-cost, single chip digital signal processors (Gehring, 90; Burgess, 92) leaves no excuse to not have spatial audio in a virtual world.

We developed our spatial audio system as part of an effort to create a three-dimensional auditory computing environment for blind computer users (Burgess, 92). The spatialization system operates by applying a pair of digital filters to an audio stream. By changing these filters in real time, it is possible to simulate the movement of a virtual sound source. The filters used in this system are based on a set of middle-ear recordings provided by Prof. Fredric Wightman of the University of Wisconsin at Madison (Wightman & Kistler, 1988). Additional processing is applied to the sound to generate simple environmental effects, specifically dense reverberation and atmospheric dispersion. These environmental effects help the user get a sense of distance from the sound source and also give a sense of the size of the listening area. The computational engine of the system hardware is currently an Ariel S-56x DSP board, but it is being ported to a multi-processor SPARC architecture.

Our first testbed application for adding spatial audio cues to a virtual world was rather simple - the virtual world consisted of a green field and a hand-held radio. A virtual sound source (an actual radio channel) was attached to the virtual radio. As the user moved the hand-held tracker, the sound source moved with the virtual radio. In this application, our primary goal was to determine if the visual cues in the virtual environment would improve the localization of the spatial audio source. Our results which seem to indicate that this improvement was not achieved to the level of our expectations may be due to many external factors. Hardware limitations of the helmet mounted display, such as its weight, the low resolution (equivalent to being legally blind), and the narrow field of view seemed to weaken any augmentation of the spatial cues. Likewise, the application only provided a few visual cues, with an extremely low level of detail.

Our goal is to develop new applications which test the effectiveness of spatial audio in navigating a complex virtual environment, as well as improving the level of immersion in the virtual environment. Since audio is a natural mechanism for conveying information about objects not in the visual field of view, spatial audio cues should be quite useful in assisting a person navigating a large virtual space. Likewise, the addition of spatial audio as a coherent cue in the virtual environment should produce an increased level of immersion as combined visual and auditory stimuli reinforce the perception of reality.

The other focus of our audio research is the use of auditory icons in computer interfaces and virtual environments. The motivation for auditory icons is based on the concept of everyday listening. People describe sounds relative to the objects interacting to make the sounds, not relative to classical dimensions such as pitch and duration. What this realization means to an interface designer is that we can use sounds to make users think of familiar objects in the same way as the graphical icons characterize commonplace objects (Gaver, ,1989; Gaver, 1991).. As with our work with spatial audio, we began exploring the use of auditory icons in order to create computer interfaces for blind users.

Past research in virtual environments has already indicated that the inclusion of audio cues to indicate events in virtual environments can improve the user's performance for certain tasks such as grasping or moving an object (Wenzel, et al., 1988). This benefit is due to the ease with which changes or events can be conveyed via auditory cues. Our goal is to demonstrate that the use of natural auditory cues can also increase the sense of immersion in the virtual environment. Careful sound design has much to do with the feeling of immersion while "watching" a movie. Try watching a movie, especially an action or science fiction movie, with the sound turned off. The included sound-effects add realism to the interactions between people and their surroundings even if those surroundings are purely imaginary. The goal of this work will be to achieve what the movie industry has been able to do for years, making people believe the unreal through the use of sound. *(Project Director: Elizabeth Mynatt)*

3.4 User-Customized Virtual Environments

We have recently extended Glyphmaker (Ribarsky, et al., 1993), a visualization tool originally build for a window-based environment to an immersive virtual environment. With this approach we can allow users to design their own virtual environments for visualization in a point-and-click fashion. The glyph-based visual representations are particularly useful for analyzing highly correlated and dynamic multivariate data such as might come from large scale observations or computer simulations. A variety of studies have shown that glyphs, which are graphical objects whose elements (e.g., position, size, shape, color, orientation, etc.) are bound to data, are useful in depicting multivariate data (Ribarsky, et al., 1993).

In our approach we have a binder interface that allows the user to interactively map data objects unto a chosen set of graphical objects depending on what she sees in the virtual space. Each data object typically has a position (and time as a parameter) plus other variable values for that position. Thus a data object might represent an atom, a fluid flow cell, a finite element, or an observational point. The user can remove the bindings, change the ranges of either the variables or the glyph elements, or rearrange bindings, and immediately view the results. This feature is very powerful for studying variables in detail or the correlations between two or more variables. Our initial VE implementation allows one-to-one bindings between data objects and simple glyphs such as spheres, cylinders, cuboids, arrows, and cones. In a later version we plan to allow mapping of collections of data objects onto surfaces.

Our interface design must take advantage of the features of the virtual environment while avoiding the limitations. This required a total redesign of the Glyphmaker binder interface since it did not employ natural manipulation and the sets of controls are too complex for the limited display resolution and sensor accuracy in the VE. We have come up with a virtual world binder interface which uses a virtual workbench that is normally located at the user's waist, in order to be easily reached without obstructing the user's view (Figure 3). Under the workbench are two drawers, one containing the basic glyphs and the other data file icons. Around the user, floating in the virtual space, are menus for the variables of the chosen file and for visualization control. The user can move or hide the menus or the workbench whenever he wants. On the workbench, the user can build a glyph, by using various controls, such as sliders for changing the color and transparency, and handles for scaling and orienting the glyph.

Once the user is satisfied with the appearance of the glyph, she can bind it to either end of the scale bar in the upper middle of Figure 3. Binding two different glyphs (with one or more changed attributes) to a scale bar gives us a way to choose a range for that element. Adjusting the arrows below the scale bar chooses the range for the previously attached data variable. If either arrow is moved and set beyond the end of the scale bar, the variable range on the scale bar is extended accordingly. We can view the result of this binding immediately after it has been established. Moreover we can change bindings, or turn them off or on, at will. This capability is very useful in studying complex relations between variables. The user can also step backward or forward in time, grab the data directly to reorient it for a better view or, using various tools, select regions for special attention.

In Figure 4 we show a typical data visualization in our virtual environment. Here we depict one time step from a molecular dynamics simulation run on a supercomputer. A small NaCl crystal (above) is smashing into a Ne surface. This process causes a great deal of heating (temperature is bound to sphere color in this depiction), large stresses between the atoms, and great disruption of the Ne surface. Such data is ideal for the virtual environment and for user-customized visualization. In the VE the user can easily grab or move around the data, even as it is being animated, to get a detailed view of structural change. She can turn on or off bindings to temperature, stress components, and other variables to get an idea of how they correlate with the system dynamics. The building, binding and visualization process requires a lot of grabbing, pointing and selecting. We are also developing, for example, more than one control interface for grabbing--allowing the user to either grab the data directly or manipulate it at a distance by "remote control". The latter feature may be useful when the user must "step back" from complicated data to gain an overview.

Our user-customized virtual world has one clear advantage. Since the user has great control over the level of detail (and what is detailed) in the visualization and eventually over the types of tools she employs, she can limit graphical structure and the means of interaction to retain immersion while still focusing on details she thinks are important. She can, for example, use her tools to select a spatial region in the data for removal, for isolation, or for display of only a limited number of properties. In another mode, her tools could be used to select a spatial region for replacement by a set of simple polygons. Since current VE systems and those in the foreseeable future are likely to have narrow optimal operating ranges, this ability to customize one's virtual world should provide an important level of refinement for some time.

Beyond these considerations, our organization of data into data objects is a powerful feature that we can use to reorganize the data based on our interaction with the visualization. This would be accomplished by a simple reclassification of variables within the data objects that would allow new bindings of data to glyphs based on our selection of a spatial region or any other variable range (e.g., what is in the selected region is in one class and what is outside is in another) (Foley and Ribarsky, 1993). The power here is the ability to isolate data variables at any level of detail and in many combinations for closer study. This focusing and linking feature has been recognized as quite useful for the study of multidimensional data in statistical graphics (Buja, et al., 1991). *(Project Director: William Ribarsky)*

4. List of VE Faculty and Interests

In addition to the projects described in the preceding section, the virtual environments group in the GVU Center carries on a wide variety of work that includes many other aspects of virtual reality research. Associated faculty members and their current interests include:

Jay Bolter (Professor, School of
Literature, Communication & Culture)
Annotation of virtual spaces.
E-mail: jb121@prism.gatech.edu
Phone: (404) 894-2735

David Burgess (Research Scientist,
College of Computing)
Directional sound for virtual
environments.
E-mail: burgess@cc.gatech.edu
Phone: (404) 894-3658

Gregory Corso (Associate Professor,
School of Psychology)
Auditory and human performance
issues in virtual environments.
E-mail: gc4@prism.gatech.edu
Phone: (404) 894-6772

Elizabeth Davis (Associate Professor,
School of Psychology)
Visual performance issues in virtual
environments.
E-mail: ed15@prism.gatech.edu
Phone: (404) 853-0194

Norberto Ezquerra (Associate
Professor, College of Computing)
Virtual environments for medical
applications.
E-mail: norberto@cc.gatech.edu
Phone: (404) 853-9173

James Foley (Professor, College of
Computing; Director, GVU Center)
Information presentation in virtual
environments
E-mail: foley@cc.gatech.edu
Phone: (404) 853-0671

Larry F. Hodges (Assistant Professor,
College of Computing)

Stereoscopic display, interaction techniques, and software development for virtual environments.
E-mail: hodges@cc.gatech.edu
Phone: (404) 894-8787

Tolek Lesniewski (Research Engineer, College of Architecture)
Virtual environments for architectural applications.
E-mail: tolek@archsun.gatech.edu
Phone: (404) 894-2980

Elizabeth Mynatt (Research Scientist, College of Computing)
Nonspeech audio, directional sound for virtual environments.
E-mail: beth@cc.gatech.edu
Phone: (404) 894-3658

William Ribarsky (Senior Research Scientist, Office of Information Technology)
Virtual environments for scientific visualization.
E-mail: ccsupwr@hydra.gatech.edu
Phone: (404) 894-6148

Mike Sinclair (Senior Research Scientist, Office of Interdisciplinary Programs)
Training applications of virtual environments.
E-mail: msinclai@gtri01.gatech.edu
Phone: (404) 894-4931

5. Final Comments

You may request a list of available GVU technical reports by sending e-mail to joan@cc.gatech.edu or via anonymous FTP to ftp.gvu.gatech.edu. The list of reports is located in file pub/gvu/tech-reports-list.txt . Several of the reports are also located in this directory in PostScript format. For general information on Virtual Environments research at Georgia Tech, you may contact:

Larry F. Hodges, Ph.D.
Co-Associate Director for Education
GVU Center
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
(404) 894-8787
hodges@cc.gatech.edu

The projects described here are supported in part by grants from the National Science Foundation, the Georgia Tech / Emory Biotechnology Center, and EduTech Institute.

References

American Psychiatric Association: Diagnostic and Statistical Manual of Mental Disorders, 3rd Edition, Revised (1987). Washington, D.C. American Psychiatric Association, 243-245.

Bajura, M., Fuchs, H. and Ohbuchi, R. (1992). Merging virtual objects with the real world: seeing ultrasound imagery within the patient. *Computer Graphics,* 26,2, 203-210.

Bergman, L.D., Richardson, J.S., Richardson, D.C. and Brooks, Jr.,F.P. (1993). VIEW - an exploratory molecular visualization system with user-definable interaction sequences. *Proc. SIGGRAPH 93*, 117-126.

Bernstein, M., Bolter, J. D., Joyce, M., and Mylonas, E. (1991). Architectures for Volatile Hypertext, *Proceedings of Hypertext '91* (San Antonio, December 15-18, 1991). New York: ACM, 1991, pp. 243-260.

Biran, M. and Wilson, G.T. (1981). Treatment of phobic disorders using cognitive and exposure methods: A self-efficacy analysis. *Journal of Consulting and Clinical Psychology*, 49, 886-899.

Brooks, Jr., F.P. (1986). Walkthough - a dynamic graphics system for simulating virtual buildings. *Proc. 1986 Workshop on Interactive 3D Graphics*., 9-21.

Buja, A. McDonald, J.A., Michalak, J. and Stuetzle. W. (1991). Interactive data visualization using focusing and linking. *Proceedings of Visualization '91*, San Diego, 156-163.

Burgess, D.A. (1992) Techniques for low cost spatial audio, *Proceedings of the Fifth Annual Symposium on User Interface Software and Technology (UIST '92)*, ACM, New York, 1992, pp. 53-59.

Burgess, D.A., Verlinden, J.C. (1993) A first experience with spatial audio in a virtual environment, Technical Report: GIT-GVU-93-31.

Foley, J. and Ribarsky, W. (1993). Next-generation visualization tools. *Proceedings of the ONR Darmstadt Workshop*. Springer-Verlag, In press.

Gaver, W.W. (1989) The SonicFinder: An interface that uses auditory icons. *Human Computer Interaction*, 4:67-94, 1989.

Gaver, W.W., Smith, R.B., O'Shea, T. (1991) Effective sounds in complex systems: The ARKola simulation. *Proceedings of CHI 1991*, New Orleans, April 28 - May 2, 1991, ACM, New York.

Gehring, B. (1990) *Focal Point (tm) 3D Sound User's Manual*, Gehring Research Corporation, 189 Madison Avenue, Toronto, Ontario, Canada, M5R 2S6.

Gloor, P. A. Cybermap (1991). Yet another way of navigating in hyperspace. *Proceedings of Hypertext '91* (San Antonio, December 15-18, 1991). New York: ACM, 1991, pp. 107-121.

Hemmje, M. (1993). A 3D based user interface for information retrieval systems. SIGIR '93, Pittsburgh, June 27, 1993.

Lai, P. & Manber, U. (1991). Flying through hypertext. *Proceedings of Hypertext '91* (San Antonio, December 15-18, 1991). New York: ACM, 1991, pp. 123-132.

Macinlay J. D., Robertson G. G., Card, S.K. (1991). The perspective wall: detail and context smoothly integrated. *Proceedings of CHI '91* (New Orleans, April,1991). New York: ACM, 1991.

Marks, I.M. & Gelder, M.G. (1965). A controlled retrospective study of behaviour therapy in phobic patients. *British Journal of Psychiatry*, 111, 571-573.

Mynatt, E.D., Edwards, W.K., (1992) Mapping GUIs to auditory interfaces, *Proceedings of the Fifth Annual Symposium on User Interface Software and Technology (UIST '92),* ACM, New York, 1992, pp. 61-70.

Ribarsky, W., Ayers, E., Eble, J., and Mukherjea, S. (1993). Using glyphmaker to create customized visualizations of complex data. Technical report GIT-GVU-93-25.

Robertson, G. G., Jock D. Macinlay, J. D., Card, S. K. (1991). Cone trees: animated 3D visualizations of hierarchical information. *Proceedings of CHI '91* (New Orleans, April,1991). New York: ACM, 1991.

Sheridan, T. B. (1992). Musings in telepresence and virtual presence. *PRESENCE: Teleoperations and Virtual Environments* 1,1, 120-125,

Streitz, Nobert, et al. (1992). SEPIA: A cooperative hypermedia authoring environment. *ECHT '92: Proceedings of the ACM Conference on Hypertext* (Milan, November 30-December 4, 1992). New York: ACM, 1992, pp. 11-22.

Verlinden, J., Bolter, J. D., & van der Mast, C. (1993). Voice annotation in virtual environments,. Paper to be delivered at ESS '93.

Verlinden, J., Bolter, J. D., & van der Mast, C. (1993). The world processor: an interface for textual display and manipulation in virtual reality. Technical Report: GIT-GVU-93-13. Atlanta: Georgia Institute of Technology, 1993.

Verlinden, J.C., Kessler, D. and Hodges, L.F. (1993). The simple virtual environment (SVE) Library: User's Guide. Technical Report GIT-GVU-93-24, 61 pages.

Wenzel, E.M., Wightman, F.L., Foster, S.D. (1988) A virtual display system for conveying three-dimensional acoustic information, *Proceedings Human Factors Society,* 32, 86-90.

Wightman, F.L., Kistler, D.J. (1988) Headphone simulation of free-field listening I & II: stimulus synthesis and psychophysical validation, *Journal Acoustical Society of America*, 85, 858-878.

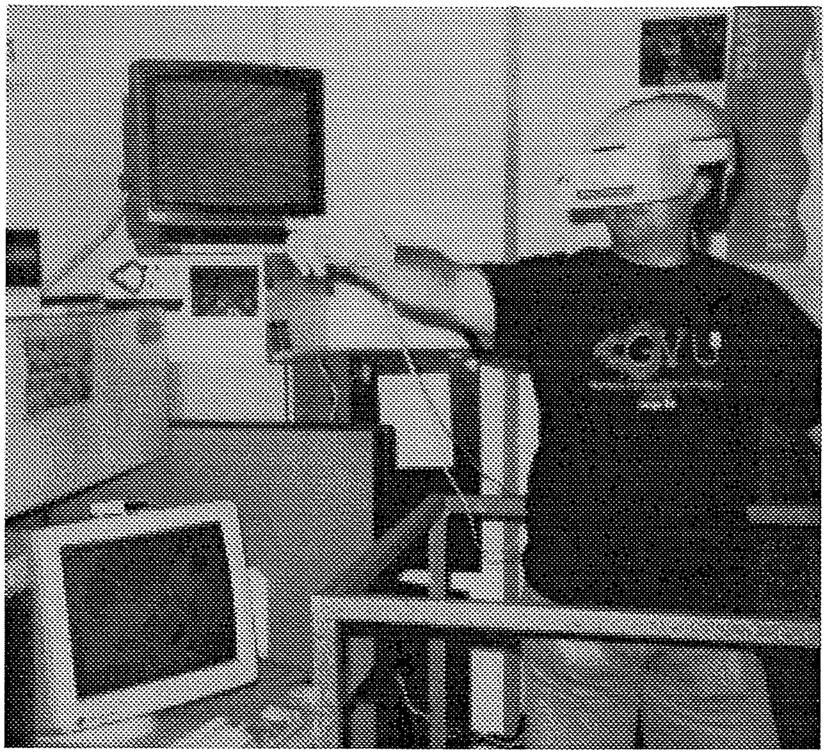

Figure 1. Riding in a virtual elevator.

Figure 2. View from the virtual elevator.

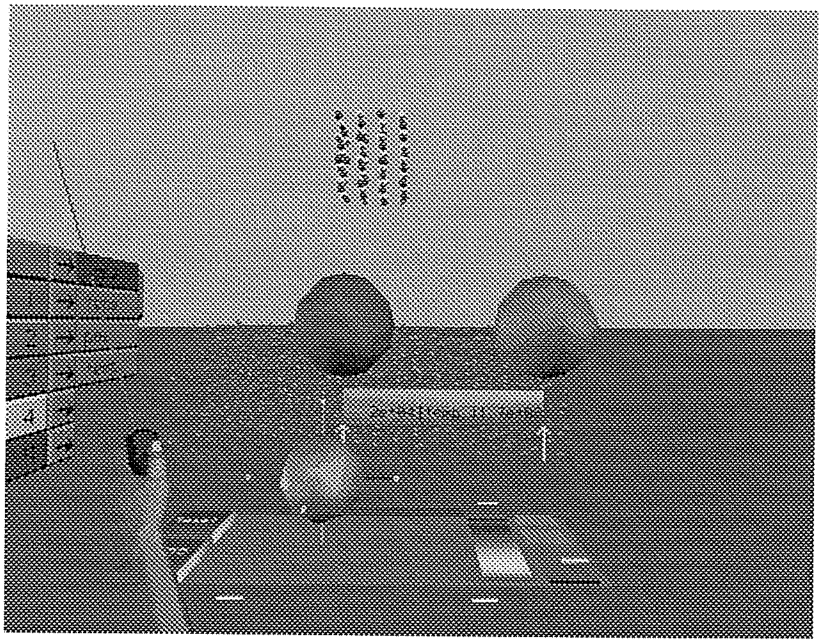

Figure 3. Virtual environment glyph binder interface for mapping data variable onto graphical elements in the virtual environment.

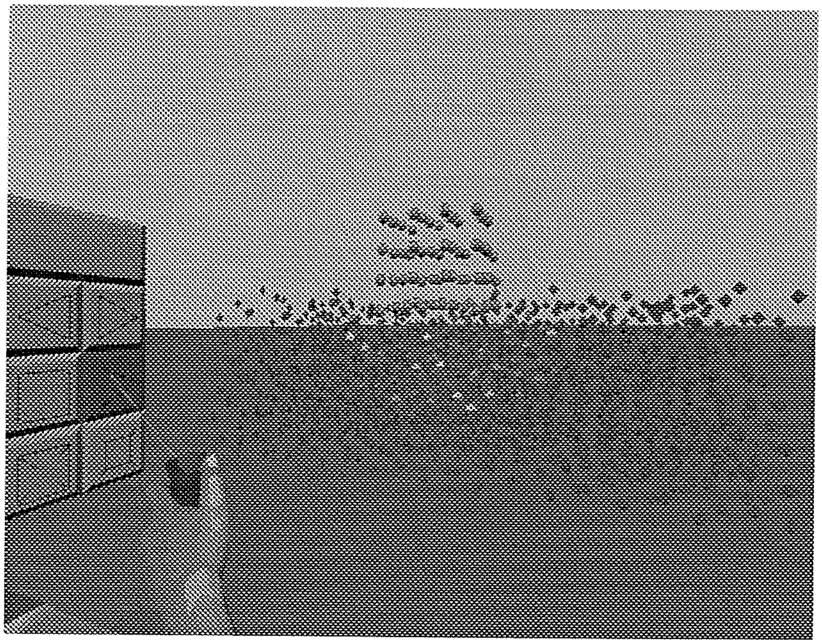

Figure 4. A time step in a molecular dynamics simulation as seen in the virtual environment. Here a NaCl cluster is smashing into a Ne surface with temperature bound to atom color. Data supplied by David Luedtke and Uzi Landman..