

The Virtual Venue: User-Computer Interaction in Information-Rich Virtual Environments

Abstract

We present a virtual environment application that allows users to access embedded information within an immersive virtual space. Due to the richness and complexity of this environment, efficient and easy-to-use interaction techniques are a crucial requirement. The "Virtual Venue" seamlessly combines both two- and three-dimensional interaction techniques into a single system and utilizes previously reported as well as novel techniques that fit the task of information access. We present tools for user control of the system, travel through the environment, and information retrieval, as well as authoring tools for the creation of information-rich virtual environments. A usability study and its results are also presented and discussed. The study indicates that the use of abstract information that is tightly coupled to the virtual environment can be quite successful in enhancing the relevance of both the environment and the information. Results also show that the set of well-constrained interaction techniques presented here are usable and efficient for information retrieval.

1 Introduction and Related Work

Many successful immersive virtual environment (VE) applications can be characterized by a low degree of interaction between the user and the system. For example, architectural walkthrough (e.g., Brooks et al, 1992) requires only that the user has some technique for moving through the space. VEs used for exposure therapy (Hodges, Rothbaum, Kooper, Opdyke, Meyer, North, de Graff, & Williford, 1995) may contain no interactive techniques whatsoever, with the exception of head tracking. It is hoped, however, that useful applications of VEs exist in which the user does more than look at the environment and move through it. Research is underway for applications such as immersive modeling

(e.g., Bowman, 1996) and surgical simulation (e.g., Billingham, Savage, Oppenheimer, & Edmond, 1996), but these more complex systems have not yet seen much practical use.

We propose that information retrieval, for educational or instructional purposes, is an interactively complex application of immersive VEs that will be of great usefulness. The ways in which people access information have rapidly moved from the simple printed page to the use of audiovisual materials to the World Wide Web and multimedia software titles, and the use of VEs as information spaces may be the logical next step. Information retrieval is now characterized by nonlinear access (e.g., hyperlinks), the use of multiple information types, and a greater degree of user control over the presentation and content of the information.

It seems natural, then, that we should ask whether nonspatial information might be enhanced for the user if presented and accessed in the context of an immersive, three-dimensional virtual environment, since VEs offer a context for information and a natural browsing metaphor (head tracking) that cannot be replicated in nonimmersive environments. If so, then appropriate interaction techniques will be necessary to give the user control over the system and easy access to this *embedded information*.

The use of 3-D environments as information spaces is not a new idea. A good deal of work in 3-D desktop in-

Doug A. Bowman

bowman@cc.gatech.edu

Larry F. Hodges

hodges@cc.gatech.edu

College of Computing
Graphics, Visualization, and Usability Center
Georgia Institute of Technology
Atlanta, GA 30332-0280

Jay Bolter

jay.bolter@lcc.gatech.edu

School of Literature, Communication, and Culture
Graphics, Visualization, and Usability Center

formation systems has been done. For example, Robertson, Card, and Mackinlay (1993) organized information spatially and allowed users to view, access, and navigate through the information using a variety of techniques (cone tree, perspective wall, point-of-interest navigation, etc.). Another application of this type is Semnet (Fairchild, Poltrock, & Furnas, 1988), which provides 3-D representations of information databases in which spatial relationships relate to semantic relationships in the information. Pieces of information are represented as graphical objects in a 3-D structure, and object attributes such as size and color are used to indicate characteristics of the data.

Much of the previous work in the use of immersive VEs as information spaces has focused on *information visualization*. One class of systems deals with scientific visualization in VR, such as the virtual wind tunnel (Bryson & Levit, 1992), the Virtual Data Visualizer (Ribarsky, van Teylingen, & van der Mast, 1995), Cosmic Explorer (Song & Norman, 1993) and ScienceSpace (Dede, Salzman, & Loftin, 1996). These applications allow the user to visualize animations of abstract objects that represent scientific data.

Another category of systems, similar to the desktop applications discussed above, present abstract *database visualizations* (e.g., Benford, Snowdon, & Mariani, 1995; Fairchild, 1993; Risch, May, Thomas, & Dawson, 1996). Such applications attempt to organize a complex dataset into an understandable visual representation, which can then be navigated and accessed by the viewer. Both of these types of environments present abstract information by converting it to a perceptual form using color, shape, texture, or animation to represent information attributes.

Our work has a somewhat different focus. We are researching ways that symbolic data relating to an environment may itself be embedded into that environment. For example, semantic information related to a college campus may be included with a virtual model of the campus, so that selecting a building reveals information such as its name, the departments it houses, office hours, etc. In this way, users can form links between the perceptual data and the symbolic data that relates to it (Bolter,

Hodges, Meyer, & Nichols, 1995). Learning and relevance of the information may be enhanced for the user because of this coupling, and, similarly, enjoyment of the 3-D environment could be augmented since additional information about the environment is available. In the sense that we are presenting additional information about an environment within that environment, this is analogous to work in augmented reality (e.g., Feiner, MacIntyre, & Seligmann, 1993), where virtual information is overlaid onto a view of the physical world.

In this paper, we present an example of such an “information-rich” virtual environment, along with various techniques, tools, and metaphors for simple and efficient interaction between the user and the system. These include a handheld menu system, a constrained travel technique, spatial hyperlinks, audio help, and three metaphors for information access. Although our goal is the development of VEs with a high level of informational content, we feel that the use of appropriate interaction techniques for information access and system control is just as important to the success of the application as the information content itself, and therefore deserves as much consideration. Thus, we wished to produce a synthesis of both existing and novel interaction techniques that fit the task of information gathering.

Based on these goals, our research has focused on two questions. Can the use of immersive VR actually enhance the presentation of symbolic information? If so, what interaction techniques can be used to retrieve information easily and to indicate the relationship between the information and the virtual environment? As a first step in answering these questions, we also discuss the results of a usability study performed on the system.

2 The Virtual Venue

2.1 Environment

The “Virtual Venue” is a VR application in which users are immersed in a 3-D model of the Georgia Tech Aquatic Center (Figure 1). The Aquatic Center has




Figure 1. *The Virtual Venue environment*

hosted worldclass competitions in the sports of swimming, diving, synchronized swimming, and water polo.

The Center contains two pools: a diving well and a swimming pool. It also has a diving tower complex, with five diving platforms and four springboards. Spectator seating runs the length of the two pools, and the entire complex is open to the outside on the ends, providing views of the Georgia Tech campus and downtown Atlanta.

The software is built upon the Simple Virtual Environment (SVE) Toolkit (Kessler, Kooper, Verlinden, & Hodges, 1994), and runs on a Silicon Graphics Crimson with Reality Engine. Tracking is performed by a three-tracker Polhemus Fastrak, and users wear a Virtual Re-

search VR4 head-mounted display (HMD). A three-button joystick and/or a stylus are used for input.

Within the system, we have embedded multiple forms of symbolic and perceptual information. These include text, audio, imagery, 3-D animation, and “experiential” information. Users of the Virtual Venue can move about the Center and obtain information about the building itself, the events that have been held there, and the sports of swimming and diving.

2.2 Handheld Menu System

A simple way for the user to control system actions is a necessity in such a highly interactive application. The



Figure 2. *Two Interaction Techniques for Hand-held Menus: above) joystick, next page) pen and tablet*

tool should be consistent, always accessible, nonintrusive, and easy to use. To meet these needs, we have developed a handheld menu system for our virtual environment.

The use of commands in general and menus in particular within a VE has been the subject of some debate. Some feel that all interaction in the VE should be “natural,” or similar to the physical world (Nielsen, 1993). In our experience, however, we have found that natural metaphors are not appropriate for abstract commands, and that overuse of real-world techniques can lead to system clutter and reduced user efficiency due to the multiple metaphors and tools needed to perform various tasks (Bowman & Hodges, 1995). A good deal of re-

search into virtual menus has been performed (Bowman, 1996; Ferneau & Humphries, 1995; Jacoby & Ellis, 1992). The current approach borrows from the system described by Angus and Sowizral (1995) in its use of the “pen and tablet” metaphor.

In the Virtual Venue, the user sees a virtual tablet (Figure 2), on which are printed various menu items. By selecting a widget on the tablet, users either change to a different menu or issue a command to the system. Menus are used for travel through the environment, controlling the actions of a simulated diver, obtaining text information, and setting system controls. The menus are arranged in a hierarchical structure with only two levels in most places, and three levels maximum.



Fig. 2 (Cont'd).

The menu hierarchy is shown in Figure 3. We have developed two separate techniques for navigating the menus and selecting menu items.

If only two trackers are available (one being used for the HMD position), a modified joystick is used to control the menu system (Figure 2, page 481). The joystick is tracked in 3-D space and has three buttons that correspond to selecting the current item, scrolling down through the menu, and returning to the next-level menu. A virtual representation of the joystick appears in the VE, and the tablet appears to be attached to the top of the joystick. Feedback is provided such that, when a button is pressed, the user sees the graphical thumb move to press that button on the virtual joystick.

If three trackers are available, we use a pen-and-tablet metaphor (Figure 2, page 482). Users hold a physical tablet and stylus, both of which are tracked in 3-D space (Figure 4). The user touches the desired widget with the stylus and presses the stylus button to select it. Using the same technique, the user presses a back arrow to return to the previous level.

This handheld menu system achieves all of the goals listed above. It is consistent, since all actions the user takes to control the application are performed using the same technique. The user carries the menu in her hand, so commands are always accessible. The tablet may always be placed to the side or in the user's lap, so it does not obscure the environment. Finally, the system is

Travel	Information
Diving Platform	About Center
Seating Area	Center Events
Starting Block	Swim Records
Overhead View	Dive Records
Dives	Controls
Pike Dive	Sound On
Twist Dive	Sound Off
Somersault	Follow Diver
Set Dive Speed	Don't Follow
one-fifth	
one-third	
one-half	
full speed	

Figure 3. Hierarchy of Menu Items in the Virtual Venue

simple to learn and easy to use, as we will show in the usability study presented below. The pen-and-tablet metaphor also has the advantage that it can mimic any 2-D interface that can be controlled with a single-button mouse, since users can “click” or “drag” on the physical 2-D surface of the tablet. This physical feedback provides an important constraint: if the pen is touching the tablet, users are assured that a menu item will be selected, and otherwise the pen may be used for other system functions. The physical surface of the tablet provides information that allows users to access the menu easily (by touching the stylus to the tablet) and to move between menu items quickly (by dragging the stylus across the tablet surface).

2.3 Constrained Travel

A second issue concerns the movement of users around the Virtual Venue. Most VR applications allow the user to travel freely throughout the environment, either “walking” in two dimensions (constrained to a ground plane) or “flying” in all three dimensions (Mine, 1995). Often, however, users who have freedom of movement become lost or disoriented, or miss important features of the environment (Bowman & Hodges, 1995). Since our system was not designed to allow users




Figure 4. Physical Devices used in the Virtual Venue (pen & tablet metaphor)

to examine every detail of the 3-D model, we implemented a more constrained method of travel called the *flying chair* metaphor.

Users sit in a physical chair (Figure 4) with their feet above the ground so that the illusion of a flying chair with no support below it can be maintained. In the VE, a virtual chair and legs may be seen, providing users with a sort of “virtual body” even though their legs and feet are not actually being tracked. It has been shown that the use of a virtual body can lead to an increased sense of immersion (Barfield, Zeltzer, Sheridan, & Slater, 1995). Seven locations were chosen in the VE that allowed interesting views of the Aquatic Center and related to information we wished to present. By selecting a menu item or spatial hyperlink (see below) on the tablet, users can travel to one of these positions. The system could teleport users instantly to the new location, but it has been shown (Bowman, Koller, & Hodges, 1997) that teleportation is correlated with increased user disorientation. Therefore, the motion is smoothly animated in a

straight line from the current location to the destination, and the virtual chair smoothly rotates to face the most interesting view at the new position. Views were chosen so that traveling between any two views did not cause the user to fly through any solid objects.

Using this metaphor, we greatly reduce the possibility that users will become lost or disoriented. They may feel somewhat restricted, since they are only able to travel to a set of discrete locations, but, if these predefined positions are chosen wisely, users will be able to obtain views of all the interesting parts of the environment, as well as to visit each place where embedded information may be retrieved. The use of constrained travel also allows better information access. By remaining oriented in the VE, finding information is simpler, and users can more easily relate the symbolic information to the environment and vice versa.

2.4 Spatial Hyperlinks

Hyperlinks are used on the World Wide Web and in multimedia applications to link related information that is not necessarily presented in a linear structure. Since we were using text information with many interrelated topics, hyperlinks were a natural addition to our set of interaction techniques. We have also extended a specialized hyperlink metaphor that seems to be especially effective in an immersive VE: the *spatial hyperlink*.

Spatial hyperlinks relate text information back to the environment in which the user sits. For example, suppose the user is reading a page of information about diving world records, which mentions a record for the one-meter springboard event. To relate this information back to the environment, the user needs to know where the one-meter springboard is, what a diver would see when preparing to dive from it, etc. By selecting a spatial hyperlink over the words "one-meter springboard," the user is automatically transported, via the flying chair, to the diving board within the environment.

The spatial hyperlink idea has been used to a certain extent in some multimedia applications (e.g., a CD-ROM describing the contents of a museum), and in VRML worlds on the World Wide Web. However, these

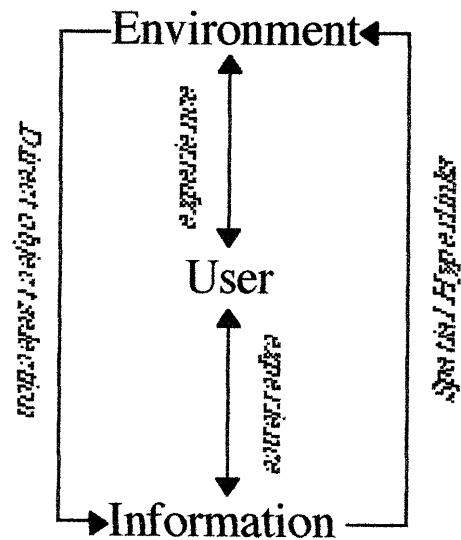


Figure 5. Information access methods can enhance user perception of the relationship between the environment and the abstract information by revealing tight couplings between spatial and abstract data.

hyperlinks usually instantly transport the user to another room or another world, with no indication of intervening context. On the other hand, in the Virtual Venue, the user travels smoothly from the current location to the new one, so that the user gains information about the relative location of the object to which the link points.

Thus, with a spatial hyperlink, the user does not simply access new information. Rather, he actually experiences a new location in the virtual world. The text information becomes more closely related to the 3-D environment, and the user's experience is enhanced. (See Figure 5.)

2.5 Audio Help

Our environment is designed for instructional purposes, so the target audience consists largely of first-time users. The system will probably be the first VR experience for many of these users as well. Since users are expected to be quite active and not simply viewing their surroundings, we felt it would be useful to provide a help system to guide users through the application. We

chose to present help through the audio modality, so that users could look at the relevant parts of the environment or interface while listening. This would not be possible with a text-based help system. Audio help appears in two contexts.

First, there are a number of *help icons* scattered throughout the environment. These are simply cubes with a recognizable icon on each side. They are positioned to be seen easily from the known positions of the flying chair. By selecting the icon, users hear context-sensitive help on relevant actions they can take while at that location. For example, a help icon placed above the diving platform instructs users how to see a simulated diver perform. (See front cover.)

Second, *automatic help* is provided for spatial hyperlinks when they are the focus of the user's attention. When the user touches the link with the stylus, it provides visual feedback and plays a short message explaining the action that will be taken if the link is selected.

This audio help is nonintrusive: it allows users to continue interacting with the system while it is being played. Users can also disable audio at any time if they do not desire to hear the help. This help system provides the right balance of prompting users when they need it while not distracting them from the other interesting features of the environment or the embedded information.

2.6 Information Content

The Virtual Venue as an information space has a rich variety of content. Information was gleaned from Web pages, newspaper articles, record books, and drawings. We attempted to include a wide range of information types to enhance the virtual Aquatic Center.

The most ubiquitous category of information is simple text. Text is available describing the building, the swimming pools, the seating capacity, swimming and diving world records, technology used in the Center, and so on. Some of the text includes spatial hyperlinks as described earlier. Text is presented to the user directly on the interface tablet. In this way, the text is always available no matter where the user is in the environment. The fact that the tablet is tracked also provides an important

benefit: the tablet may be brought closer if the user wishes the text to be larger for readability.

We also make use of a good deal of audio in the Venue environment. Besides the previously discussed audio help, we also use nonspatial audio for two other purposes. First, environmental sounds are played when appropriate. These include a cheering crowd in the grandstand and a splash when the diver dives into the pool. Second, many objects in the VE are annotated with audio. When the user selects such an object, she hears its name spoken aloud.

Image information is also available. Specifically, we have collected various photographs of the Aquatic Center and divers and placed them on the "scoreboard" in the Center. When the user selects the scoreboard, the image there (along with a caption) changes. By continuing to select the scoreboard, the user can see a virtual slide show of the available photos.

As we have mentioned previously, we have also included a physically simulated diver (Wooten & Hodgins, 1996) in the environment (front cover). The diver performs three different dives, and his speed may be controlled by the user with the handheld menu. This is an example of *animated 3D information*, which is an information class that is difficult to present in more traditional settings. It is especially effective in an immersive virtual environment, because the user can examine the animation from different points of view in order to gain knowledge about the dive.

Finally, we have embedded what we call *experiential information* within the Virtual Venue. This type of information allows the user not only to view but also to experience or "feel" informational content. For example, the user may choose to follow the diver from the starting position on the ten-meter platform all the way down into the pool. In this way, the user can examine the dive more closely. More importantly, however, the user experiences something of the speed of the dive and the force of gravity pulling the diver toward the water. Also, there are spatial hyperlinks within the text describing world records in swimming. Selecting these links cause the user to "go swimming," allowing him to experience the distance involved in a given event, such as the 100-meter freestyle.

2.7 Information Access

Given all this content, it is important that users have consistent, simple, and efficient techniques for accessing the information. Three main methods are used in the Virtual Venue for information access and retrieval.

First, much information may be obtained by choosing an item on the interface tablet. This is an indirect technique, but it is an efficient one if broad information is desired on a specific topic, since the tablet is always available to the user. Abstract text information (such as a list of events held at the Center, or information on technology used there) is accessed with this method. The tablet is also used to trigger the simulated diver and the different types of experiential information, through menu items or spatial hyperlinks.

Second, text and/or audio may be retrieved by *direct selection of objects* in the VE. We wanted users to be able to select objects no matter their position in the environment, rather than being able to select only the objects that they could touch. Therefore, we implemented a ray-casting technique (Mine, 1995), in which the user points a light ray at the object of interest. In our application, the stylus is also used for this purpose. If the stylus is not touching the tablet, pressing its button causes a light ray to emanate from its tip. When the ray intersects a "hot" object, the object changes color, and the user simply releases the button to select the object. This technique is easy to use and powerful, since any object in the field of view may be selected. User studies (Bowman & Hodges, 1997) have indicated that ray casting is very efficient for selecting of objects in immersive VEs.

We used this method for the retrieval of text specifically tied to an object, such as the dimensions of the diving pool, as well as to obtain the names of objects via audio clips. This selection technique is also used to select the help icons and the scoreboard described earlier. Just as spatial hyperlinks reveal a mapping from text information to the virtual environment, direct object selection maps the environment back to abstract information (Figure 5).

Finally, some audio information is obtained via *context-sensitive, automatic retrieval*. Sounds related

to a certain area are triggered when the user enters that region of the environment. For example, when the user is taken to the swimming pool via the spatial hyperlink, he hears the splashing water as he "swims" the length of the pool. Other locations also contain environmental sound or spoken information that plays automatically.

3 Authoring Tools

In order to create complex, information-rich virtual environments such as the Virtual Venue, programming toolsets are necessary. These tools must provide abstraction away from the lower-level graphics or virtual environments software. They should allow rapid prototyping and experimentation, and they should be reusable across multiple applications. For this project, we developed two such toolsets: one for audio annotations and one for pen-based interaction. Although these toolsets were used at the programming level, they could easily be extended to use text-based input files or a graphical user interface.

3.1 Audio Annotation Tools

As we have discussed, a great deal of information in the Virtual Venue was presented in the auditory modality, including environmental sounds, place and object names, and audio help. To support the development of these parts of the system, we designed and implemented a general toolset for embedded audio within a VE application.

Other audio annotation systems have been implemented (e.g., Harmon, Patterson, Ribarsky, & Bolter, 1996; Loughlin & Hughes, 1994). Most systems have focused, however, on allowing VE users to record and place annotations within a collaborative design or visualization environment. Our toolset, on the other hand, is designed to allow developers to include audio annotations of various types within a 3-D space, and to provide techniques with which users may access these annotations.

The most important feature of our audio annotation

toolset is its support for many different annotation types. The simplest type is an audio clip that plays continuously in the environment. There are also several types that play automatically when certain conditions are met. These include when a given object comes within a certain radius of a 3-D point or of another object, when a given object enters a specified bounding box around a 3-D point or another object, or when the user looks at a certain 3-D point or object in the environment. Finally, there are object annotations that play when the user selects a specified object.

For this last type of annotation, the toolset provides a default interaction technique (ray casting) for object selection. It also allows the user to specify that a default annotation icon should be used instead of an object already in the environment.

All annotations can be selectively enabled or disabled, and all annotation attributes can be changed at any time during application execution. Since the development of the Virtual Venue, this toolset has also been used in the development of another information-rich VE application in our laboratory.

3.2 Stylus Interaction Tools

We also developed a toolset for stylus-based interaction, on which was based the handheld menu system described above. We wanted the interaction in the Virtual Venue to be simple, so we used only selectable menu items, but the framework provided by the stylus toolkit actually allows many more types of interaction.

To use the stylus library, the programmer simply designates certain objects as “widgets.” These widgets each have an associated precision (how close does the stylus have to come to the widget before it is activated) and callback routines that implement the widget’s functionality. Application routines can be called when the stylus enters or exits a widget, or when the stylus button goes up or down within a widget. Thus, we have extended the traditional pointer and widget framework from 2-D graphical user interfaces (GUIs) to 3-D virtual environments.

From this simple model, practically any 2-D user in-

terface can be constructed within a 3-D VE, including those with draggable objects. Many novel 3-D interfaces can be created as well, although these will generally lack the physical feedback afforded by the tablet in our 2-D interface.

The stylus toolset also contains two important utility routines. The calibration utility allows the user to interactively calibrate the position of widgets so that they work in the desired manner (e.g., so that the physical and virtual tablet surfaces are aligned). Each widget is highlighted in turn, and the user presses the stylus button at the desired location of that widget. This can be important when using six-DOF tracking devices, which are often imprecise or noisy. Also, a routine is included that returns the relative position of the stylus tip to the origin of a widget. This can be useful for interfaces such as image maps, where a single texture serves as an interface, and the callback function is determined by the position selected within the image. This utility may also be used in three dimensions. For example, a cube representing RGB color space could be defined as a widget, with the selected color depending on the relative location selected within the cube.

As with the audio annotation toolset, the stylus tools have been reused in other VE applications. They allow designers to focus on the visual appearance of the interface and the core functionality, rather than on the low-level processing of events.

4 Usability Study

In order to test our hypotheses about embedded information in a virtual environment (and also the interaction techniques we had designed to access that information), a usability study was performed. This was a first pass at evaluating the utility of an information-rich virtual environment, intended to show whether the application concept was compelling and enjoyable, and whether the particular set of chosen interaction techniques was usable for the task of information access. Our future plans include a formal experiment comparing an immersive VE against other methods of information access. Ten evaluators, all of whom were students, partici-

pated in the study. Both graduate and undergraduate students, and both technical and nontechnical majors were represented. None of the users had any previous experience with immersive VR. The usability study was designed to exercise all the features of the system, and to put an emphasis on ease of use, ease of learning, and efficiency.

4.1 Structure and Content

As discussed earlier, we designed two different interaction techniques for operating the handheld menu system: one using a joystick, the other a stylus. In the usability study, these techniques were compared against one another. Therefore, the ten evaluators were divided into two groups, one group for each technique. Evaluators were not told of the other menu technique until after they had completed the entire study.

Except for this distinction, the study was the same for both groups. After completing a consent form and background questionnaire, users performed some tasks in the immersive Virtual Venue system. The task section was divided into three parts. Evaluators were timed on each individual task, errors were noted, and a detailed log of user actions was kept by the system.

In order to test ease of learning, the first set of tasks was designed so that users would learn the main interactive techniques on their own. After donning the HMD, evaluators were given no instruction on the use of the menu system or object selection light ray, except that it was possible to select menu items or objects, using the buttons on the joystick and/or stylus. The tasks in this first part were simple and redundant, so that speed of learning could be measured over several attempts at the same task. Tasks involved the use of only one of the four top-level menus (travel), so that other system functionality would not be revealed until later.

Second, the users entered an exploration phase. Here, evaluators were allowed to use the system freely for ten minutes, so that they would become acquainted with the rest of the application, such as the simulated diver, the controls menu, and the use of spatial hyperlinks. Again, no instruction was given, unless the evaluator had failed

to find one or more of the major techniques at the end of the ten minutes.

Finally, a set of more complex tasks was presented. These involved multiple actions on the part of the evaluator, and often multiple interaction techniques. This section was designed to test the evaluator's overall system knowledge, including the structure of the menu system, object selection, use of spatial hyperlinks, and access paths for specific pieces of information.

Each session concluded with a feedback questionnaire and interview with the evaluator. The questionnaire (Appendix A) consisted of several five-point rating scales, for which the user circled a dot closest to their subjective rating of the item in question. These responses were translated into scores ranging from -2 to 2 . There were also several free-form response questions, allowing the user to express specific comments about the interaction techniques, information presentation, and overall experience.

4.2 Results

The usability study provided us with a great deal of useful data about both the usefulness of immersive VR in information retrieval and the types of interaction techniques that are usable and efficient in such a system. The main results are summarized below.

4.2.1 Embedded Information. As we have noted, the first major issue is whether the use of embedded symbolic information in a perceptual environment is useful or enhances the information. We did not have a nonimmersive equivalent against which to compare the Virtual Venue, so a quantitative measurement was not possible. However, comments from our evaluators generally indicated that the information access task is indeed appropriate for an immersive VE. They found the experience enjoyable, and felt immersive information gathering was a useful paradigm. Subjects reported that the largest advantage of using a VE was the sense of presence and scale it provides. For example, one user commented that the VE "gave an actual sense of being there that can't be obtained from 2-D pictures." Another noted that "it's more useful to . . . look around and see

WHERE [an object] is, rather than just seeing a picture” (subject’s emphasis). The audio and text “filled in the gaps of information that . . . VEs can’t provide naturally.” In another’s final analysis, she said, “if the amount and type of information were the same, I prefer the VE to the World Wide Web or printed text.”

Based on these comments, the VE excelled at presenting spatial information, but more abstract information types were also useful in this environment. Since this was the first VE experience for all of our evaluators, their positive reactions could be attributed in part to the novelty of the VE. However, subjects were asked in the questionnaire to compare this information-gathering experience with others using the printed page or hypermedia, and responded that the VE offers definite benefits over these other types of media.

Of the information types that were embedded in the Virtual Venue, the experiential information, the spatial hyperlinks, and the information embedded within specific objects were the most highly praised by our users. Users did not feel that the plain text information (accessed via the menu) or the image information presented in slide show format were as useful or relevant. Many subjects noted that experiential information cannot be provided in other types of information displays, and felt that this was the most enjoyable part of the VE. One user commented that information presented in its spatial context can be enhanced over simple text, as we had surmised. Regarding more-static information, another subject stated that “if the presentation of the information was more integrated into the experience, I would enjoy it more.”

This leads us to the observation that those information types that are the most enhanced for the user and enjoyed by the user are those which are *tightly coupled* to the perceptual environment (Figure 5). Experiential information allows the environment itself to be utilized as a conduit for information. Spatial hyperlinks directly relate text information to locations or actions within the environment. Finally, information embedded in objects completes the cycle, relating the 3-D world back to symbolic (text or audio) information. We feel, therefore, that the choice of access techniques can have a large effect on the perception and learning of the user. Well-

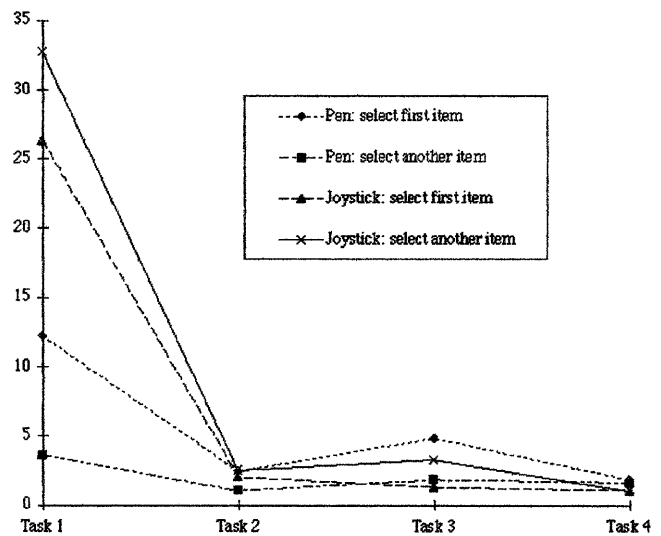


Figure 6. Comparing Ease of Learning for Two Menu Interface Techniques: Average Times for Simple Menu Tasks.

chosen access techniques, which create relational links between the environment and abstract information, will cause the information to become more relevant and learnable.

4.2.2 Handheld Menu System. The menu system is the major interaction technique used in the Virtual Venue, so the evaluators’ ability to learn and use the menu was especially important to our results. Both interfaces were well received and performed efficiently.

The pen-and-tablet interface proved especially easy to learn and use. Without instruction, most evaluators used the technique correctly the first time and times remained low for subsequent tasks (Figure 6). The joystick interface, while more indirect and less intuitive, still scored well. It took most evaluators only one or two menu tasks before they understood the mappings of the three buttons, and times then decreased to approximately the same level as those using the stylus technique.

Figure 6 breaks the tasks into those selecting the first item in a menu versus those which require selection of another item, since selecting the first item with the joystick does not require any scrolling. Times for the pen-and-tablet interface should be similar regardless of which item is selected, except that the first task in the study asked users to select the first menu item, so that times

for this task are high because learning has not yet occurred.

In a statistical comparison of these results, using a standard analysis of variance (ANOVA), we found that the pen-and-tablet interface was significantly faster for the first task in which users were asked to select the first menu item, although with only a marginal level of statistical significance ($p < 0.1$). Over all tasks in which users selected the first item in a menu, the stylus group averaged 2.38 seconds versus 7.63 seconds for the joystick group, but this did not produce statistical significance due to high variance in the times of the joystick group ($p > 0.15$). Similar results were found for the tasks in which users selected a menu item lower in the menu (stylus avg. = 1.56 secs., joystick avg. = 9.88 secs., $p > 0.2$). The large variation in task time for the joystick group was due to a few very high times, indicating that for some users the less-intuitive mapping of the joystick buttons caused problems, while others understood these mappings quickly.

Once users reached the more complex tasks in the third section of the usability study, they had mastered the use of the menu system. Five tasks in this section utilized menus heavily. For each of these tasks, the joystick group had a lower average time for completion, although these differences were not statistically significant.

Qualitatively, users preferred the stylus interface for ease of use. On the scale of -2 to 2 , with -2 representing “quite hard to use” and 2 representing “very easy to use,” stylus users rated the interface at 1.2 , while users of the joystick interface rated it at 1.0 . However, overall system satisfaction ratings were quite high for both groups (1.4 for the stylus group, 1.3 for the joystick group), even though much of the session was spent using the menu system. Overall, timings and observations showed that the stylus group tended to learn the interface more easily and make fewer errors. However, the joystick group acted more quickly and efficiently. Once the evaluators understood the button mappings, they were often able to “click ahead” of the system, producing better times on some of the more complex menu tasks.

Overall, then, both interfaces performed well. The structure of the menu was not difficult for most users to follow, and every menu task was completed by every evaluator. Use of the stylus and tablet seems to be most suited for novice users who need an intuitive and uncluttered interface. The joystick can be more efficient, however, if users are willing to spend a few minutes mastering it. Of course, these results do not show that the handheld menu system is better than other possible interfaces, but both the quantitative and qualitative measures do support our claim that such an interface is easy to learn and use.

4.2.3 Constrained Travel. In one sense, the flying-chair metaphor worked perfectly during the usability study. Not a single user became lost! Evaluators were able to remain seated for the entire session, moving only their head and hands. Although we have no proof, this may have been a significant reason why no users became dizzy or nauseated during the study, even though sessions usually lasted more than thirty minutes. In our experience with VEs where one stands and moves about, some users can become dizzy or nauseated quite quickly.

On the other hand, many evaluators commented that they felt restricted by this travel metaphor, and that they would prefer more freedom of motion to see more details of the environment. This may be partly explained by the fact that none of our evaluators had ever experienced immersive VR before, and wished to look at the 3-D environment more closely than they would had they used similar systems in the past. However, the flying chair metaphor does trade freedom for greater precision, efficiency, and user orientation. On average, users rated the flying chair at 0.6 on a scale of -2 to 2 , with -2 representing “quite ineffective” and 2 representing “very effective.”

One possible solution might be to augment the flying-chair metaphor to allow users some freedom of motion while still adhering to the main concept of predefined positions to which the user is taken. Users could be allowed to move around within a certain range of each of the main destinations, and collision-detection routines

could ensure that they would not travel through walls, floors, or ceilings.

4.2.4 Information Access Techniques. Other methods for retrieving information within the VE were also tested. Most notably, the study focused on the use of the ray-casting technique for object selection.

Selecting objects in 3-D was a foreign concept to most of our evaluators. Again, they were given no instruction in the use of the stylus or light ray and were required to discover the technique on their own. Three tasks in the first task section required users to select objects in the environment. The average times for these tasks (43.56, 6.78, and 12.56 seconds) indicate that after the initial task, users were able to use this technique fairly efficiently and accurately.

We made one important observation in the use of this interaction technique. Since most evaluators were familiar with the desktop metaphor and 2-D GUIs, they associated object selection with a “point-and-click” action. The light ray, however, requires a “press, point, and release” to work correctly. This misconception led many users to try to point the stylus in approximately the correct orientation in 3-D and then click the button.

For novice users, then, it might be beneficial for the light ray to be active at all times when the stylus is not touching the tablet. In this way, users can highlight the desired object and then click the stylus button. In other words, they would use a point-and-click metaphor.

On the whole, object selection using the light ray worked as expected. Once the technique was mastered, evaluators were easily able to select objects, even from a great distance. Because users did not have to travel to an object to select it, efficiency was maintained and large objects could be seen in context during selection.

5 Conclusions and Future Work

Through the Virtual Venue project, we have learned a great deal about the utility of virtual environments for the purpose of information retrieval and enhancement. It is clear that embedded information can be

augmented for the user if it is closely related to, and tightly coupled with, the virtual environment. We have also seen that the use of appropriate symbolic information can make the experience of immersion in a virtual environment more engaging and relevant.

In the area of interaction techniques, we have combined a useful and usable suite of tools. The handheld menu is a user interface that is approaching maturity. The concept of bringing 2-D metaphors into the virtual world no longer seems unnatural or backward. Rather, these familiar techniques translate well into 3-D, and users are often able to complete tasks more quickly and easily than by using gestures or real-world metaphors exclusively (Rygol, Dorbie, Worden, Ghee, Grimsdale, & Harvey, 1995). We have also shown spatial hyperlinks to be a technique that will prove important in relating static information back to a dynamic environment. Furthermore, the ray-casting technique for 3-D object selection works well as a parallel to traditional 2-D pointing methods. Synthesizing these techniques seamlessly into the Virtual Venue has produced a useful and usable immersive information space.

In the future, we hope to put our theories of embedded information to an even more complex test. We propose to compare the utility and learning support provided by several different information-access scenarios. These will include the printed page, hypertext documents, nonimmersive 3-D environments, and immersive VR.

We also plan to continue our development of interaction techniques for information access as well as for general user tasks in immersive virtual environments. One interesting possibility that we have already explored briefly is a general tool for database access from within a VE. If the VR application is directly linked with a database related to the environment, embedded information could be changed and updated on the fly. A tablet interface similar to the one described here can be used to specify queries to the database. Results would not only be shown in text form, but would also be linked back to the environment as spatial hyperlinks. Objects could still be selected directly, providing easy access to data related to the object.

Finally, we plan to explore issues related to the creation and authoring of these information-rich VEs. Our current toolkits for audio annotations and stylus-based interaction will be extended and integrated with other sets of tools allowing multiple information types to be embedded in a virtual world. The resulting toolset should allow developers to create rich virtual experiences with minimal effort, as default interaction and access techniques will be defined. Furthermore, the designer will be allowed to extend these techniques or create new ones, so that the coupling between the interaction metaphor and the information itself is as close as possible.

Appendix A: Feedback Questionnaire

For questions 1–6, circle the dot which is closest to your rating:

1. Rate your overall satisfaction with the system:
 • - - - • - - - • - - - • - - - •
 quite dissatisfied neither satisfied quite
 nor dissatisfied satisfied
2. Rate the ease of use of the menu system you used:
 • - - - • - - - • - - - • - - - •
 quite hard to use neither easy very easy to use
 nor hard to use
3. Rate the ease of use of the light ray for object selection:
 • - - - • - - - • - - - • - - - •
 quite hard to use neither easy very easy to use
 nor hard to use
4. Rate the effectiveness of the flying chair for travel:
 • - - - • - - - • - - - • - - - •
 quite ineffective neither effective very effective
 nor ineffective
5. Rate the effectiveness of the links found in some text information:
 • - - - • - - - • - - - • - - - •
 quite ineffective neither effective very effective
 nor ineffective
6. Rate the effectiveness of the audio annotations and help in the system:

- - - - • - - - • - - - • - - - •
- quite ineffective neither effective very effective
 nor ineffective
7. What was your favorite thing about the system?
 8. What was your least favorite thing about the system?
 9. Were there any features missing that you would have liked in the system?
 10. Do you feel that getting the information while in a 3-D virtual environment was helpful and/or useful? If so, why? If not, how would you rather obtain the information (e.g. through the World Wide Web, through a printed article, through a CD-ROM)?
 11. Please make any other comments you have about the system, the interfaces used, the information content, or the input devices used.

ACKNOWLEDGMENTS

The authors would like to acknowledge the help of our evaluators, whose time and comments are greatly appreciated. We also acknowledge the use of the physically simulated diver, developed by Wayne Wooten and Jessica Hodgins of the Gvu Animation Group. This work was supported in part by the National Science Foundation and the Hewlett-Packard Corporation.

For more information on the Virtual Venue project, see <http://www.cc.gatech.edu/gvu/virtual/Venue/>.

REFERENCES

Angus, I., & Sowizral, H. (1995). Embedding the 2-D Interaction Metaphor in a Real 3-D Virtual Environment. In *Proceedings of SPIE, Stereoscopic Displays and Virtual Reality Systems, 2409*, 282–293.

Barfield, W., Zeltzer, D., Sheridan, T., & Slater, M. (1995). Presence and Performance Within Virtual Environments. In W. Barfield & T. Furness (Eds.), *Virtual Environments and Advanced Interface Design*. Oxford University Press.

Benford, S., Snowdon, D., & Mariani, J. (1995). Populated Information Terrains: First Steps. In R. Earnshaw, J. Vince,

- & H. Jones (Eds.), *Virtual Reality Applications* (pp. 27–39). Academic Press.
- Billinghurst, M., Savage, J., Oppenheimer, P., and Edmond, C. (1996). The Expert Surgical Assistant: An Intelligent Virtual Environment with Multimodal Input. In *Proceedings of Medicine Meets Virtual Reality IV*.
- Bolter, J., Hodges, L., Meyer, T., & Nichols, A. (1995). Integrating Perceptual and Symbolic Information in VR. *IEEE Computer Graphics and Applications*, 15(4), 8–11.
- Bowman, D., & Hodges, L. (1997). An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments. In *Proceedings of the Symposium on Interactive 3-D Graphics*.
- Bowman, D., Koller, D., and Hodges, L. (1997). Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. In *Proceedings of the Virtual Reality Annual International Symposium*, 45–52.
- Bowman, D. (1996). Conceptual Design Space: Beyond Walkthrough to Design. In D. Bertol, *Designing Digital Space*. New York: John Wiley & Sons.
- Bowman, D., & Hodges, L. (1995). User Interface Constraints for Immersive Virtual Environment Applications. Graphics, Visualization, and Usability Center Technical Report GIT-GVU-95-26.
- Brooks, F. et al. (1992). Final Technical Report: Walkthrough Project. Report to National Science Foundation.
- Bryson, S., & Levit, C. (1992). The Virtual Wind Tunnel. *IEEE Computer Graphics and Applications*, 12(4), 25–34.
- Dede, C., Salzman, M., & Loftin, R. (1996). ScienceSpace: Virtual Realities for Learning Complex and Abstract Scientific Concepts. In *Proceedings of the Virtual Reality Annual International Symposium*, 246–252.
- Fairchild, K., Poltrock, S., & Furnas, G. (1988). Semnet: Three-Dimensional Graphic Representations of Large Knowledge Bases. In *Cognitive Science and its Applications for Human-Computer Interaction*. Lawrence Erlbaum Associates.
- Fairchild, K. (1993). Information Management Using Virtual Reality-Based Visualizations. In A. Wexelblat (Ed.), *Virtual Reality Applications and Explorations* (pp. 45–74). Academic Press Professional.
- Feiner, S., MacIntyre, B., & Seligmann, D. (1993). Knowledge-based Augmented Reality. *Communications of the ACM*, 36(7), 52–62.
- Ferneau, M., & Humphries, J. (1995). A Gloveless Interface for Interaction in Scientific Visualization Virtual Environments. In *Proceedings of SPIE, Stereoscopic Displays and Virtual Reality Systems*, 2409, 268–274.
- Harmon, R., Patterson, W., Ribarsky, W., & Bolter, J. (1996). The Virtual Annotation System. In *Proceedings of the Virtual Reality Annual International Symposium*.
- Hodges, L., Rothbaum, B., Kooper, R., Opdyke, D., Meyer, T., North, M., de Graff, J., & Williford, J. (1995). Virtual Environments for Treating the Fear of Heights. *IEEE Computer*, 28(7), 27–34.
- Jacoby, R., & Ellis, S. (1992). Using Virtual Menus in a Virtual Environment. In *Proceedings of SPIE, Visual Data Interpretation*, 1668, 39–48.
- Kessler, D., Kooper, R., Verlinden, J., & Hodges, L. (1994). The Simple Virtual Environment Library Version 1.4 User's Guide. Graphics, Visualization, and Usability Center Technical Report GIT-GVU-94-34.
- Loughlin, M., & Hughes, J. (1994). An Annotation System for 3-D Fluid Flow Visualization. In *Proceedings of IEEE Visualization*, 273–279.
- Mine, M. (1995). Virtual Environment Interaction Techniques. UNC Chapel Hill Computer Science Technical Report TR95-018.
- Nielsen, J. (1993). Noncommand User Interfaces. *Communications of the ACM*, 36(4), 83–99.
- Ribarsky, W., van Teylingen, R., & van der Mast, C. (1995). Virtual Data Visualizer. Graphics, Visualization, and Usability Center Technical Report GIT-GVU-95-16.
- Risch, J., May, R., Thomas, J., & Dowson, S. (1996). Interactive Information Visualization for Exploratory Intelligence Data Analysis. In *Proceedings of the Virtual Reality Annual International Symposium*, 230–238.
- Robertson, G., Card, S., & Mackinlay, J. (1993). Information Visualization Using 3-D Interactive Animation. *Communications of the ACM*, 36(4), 57–71.
- Rybol, M., Dorbie, A., Worden, A., Ghee, S., Grimsdale, C., & Harvey, J. (1995). Tools and Metaphors for User Interaction in Virtual Environments. In R. Earnshaw, J. Vince, & H. Jones (Eds.), *Virtual Reality Applications* (pp. 149–161). Academic Press.
- Song, D., & Norman, M. (1993). Cosmic explorer: A Virtual Reality Environment for Exploring Cosmic Data. In *Proceedings of the IEEE Symposium on Research Frontiers in Virtual Reality*.
- Wooten, W., & Hodgins, J. (1996). Animation of Human Diving. *Computer Graphics Forum*, 15(1), 3–13.