Position Paper

Experiences on the limits and differences of user behavior on Virtual Environments

Fernando Das Neves
Dept. of Computer Science
Virginia Tech
Blacksburg, VA 24061
E-mail: fdasneve@vt.edu

Edward A. Fox
Dept. of Computer Science
Virginia Tech
Blacksburg, VA 24061
E-mail: fox@vt.edu

Ronald Kriz
Dept. of Engineering Science and Mechanics
Virginia Tech
Blacksburg, VA 24061
E-mail: kriz@vt.edu

Abstract

In this paper we argue that the technology to create Virtual Environment (VE) is not yet to the point where we can ignore the effect on the user of the limits of the quality of the simulation. User cannot interact with the virtual world using the same techniques they use in the real world. We firmly believe there is a need to understand the limits and differences of the translation of real-world experience to VE. Furthermore we believe in creating a range of simulations with different quality adapted to heterogeneous hardware environments, and to study how to help users in those environments to interact. This paper presents the experience gathered while evaluating and using two VE: CAVE-ETD (a library simulation), and the CAVE Collaborative Console, a shared virtual space in which remote users can interact. While developed for different purposes, both environments provide insights about the way users modify their behavior to adapt to the quality and fidelity of the simulation.

Introduction

Within the last years the cost of hardware have decreased to the point where it is possible to create convincing virtual environments outside of the research labs; some of them provide a very high quality shared experience for specific tasks, like the last generation of multi-user on-line games (Quake III or Unreal Tournament). While the purpose and utility of the tasks to carry out in these games is debatable, they prove that consumer hardware in terms of graphics and input devices can support high quality virtual environments. If VEs are going to become widespread, it is necessary to learn how to bridge the gap between fully immersive environments that provide accurate spatial perception, and environments running on desktop machines. Even in the high-end research hardware limitations in resolution, contrast and physical space of the representation device, and in accurate sampling of body position and gestures in input devices are going to exist for a while. We used CAVE-ETD and the Cave Collaborative Console (CCC), two virtual environments developed at Virginia Tech to explore these issues.

CAVE-ETD: A library simulation

CAVE-ETD is a single-user VE that runs in a CAVE, developed to study how limitations in technology affected the task of searching and locating items in a collection that are relevant to an information need. The spatial design of CAVE-ETD followed the one of a real world library. A main foyer leads to different rooms, each room standing for a subset of the collection. Within each room, bookcases are arranged in rows (Figure 1). Books on the bookshelves stand for collection items, each book representing one item. The width of the book spine is directly related to a measure of size of the corresponding record, either to the number of pages if there is a printed version, or to some measure of the storage capacity used if the collection is only available electronically. Books are arranged as in a real-world library, going from left to right and from top to bottom, continuing with the next bookcase or with the leftmost bookcase in the next row once a bookcase is full.

All interaction with CAVE-ETD is performed using the wand, a three-dimensional pointing device. The system is always in one of two interaction modes: navigation or browsing. In navigation, the user uses the wand to move in the virtual library. The angle of the wand with the floor indicates the speed and direction of movement. Pointing the wand to the floor would make the world to move toward the user, as if walking forward, pointing the wand up would move the world away from the user, as if walking backwards. In both cases the higher the angle the faster the movement.

This movement can be combined with rotation by turning the wand to the side the user wanted to turn; the smaller the angle between the user and the wand, the faster the turn. Users found this movement technique simple and easy to learn, although it had drawbacks as will be described below.

By pressing a button in the wand, the interaction mode changed from navigation to browsing. In this mode the position of the user's body in the virtual world remains fixed, although he or she can still move inside the CAVE and freely move the wand. A hand appears indicating the book currently pointed to, and the spine of that book is shown highlighted (Figure 3). The user could browse the books by selecting them; each time a book is selected, the book title and authors appear on top of the bookcase, due to the limited resolution of the CAVE to display text. Pressing a second button displays and hides a card with the first 200 words of the abstract of the currently selected book (Figure 3).

To study how users changed the way they perform a search within the virtual environment we recruited 10 users, and asked each of them to perform 6 tasks consisting of retrieving the best match they could find for a certain topic within the items (books) available in a room of the library. All tasks were performed in the same room, populated with 384 books from Virginia Tech's Electronic Thesis and Dissertation. This room was filled with thesis from the colleges of Business and Human Resources and Education, and none of the participants had problems understanding the subjects or content of the abstracts. It is important to note that CAVE-ETD was not developed as a new, innovative interface for information retrieval, but as a way to test differences in behavior between the real and the simulated world.

Every participant performed all 6 tasks, although the order was randomized. For every task the user was presented with a topic he or she would pretend to be interested. In order to facilitate the retrieval of relevant books, the user would perform a search whose formulation and results were given by us. A 5-7 minute training session familiarized the user with the space, the change between navigation and browsing mode and the book selection. After that we handed a card to the user with the subject to search for. Users were asked to think aloud and elicit their actions and criteria used in deciding how to pick the next book, and the terms used in the query to identify potentially relevant books..

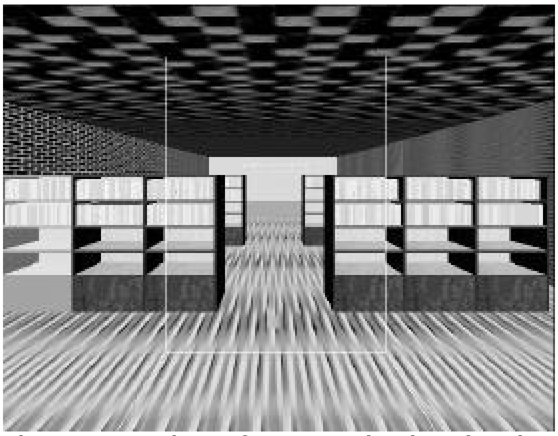

Figure 1. A view of a room in the virtual library taken from the entrance door.

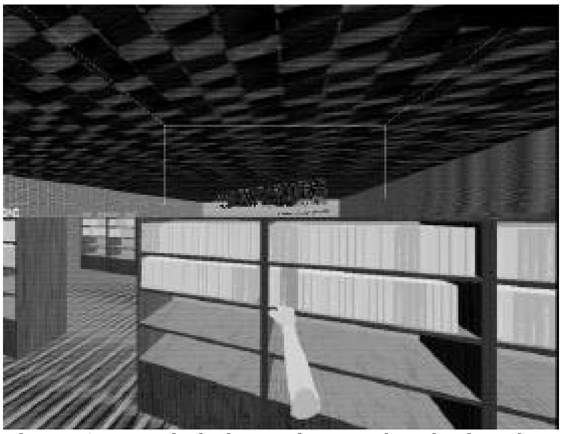

Figure 2. User is in browsing mode. The hand points to the currently selected book with is highlighted.

Figure 3. The user recalled the abstract of the currently selected book.

The setting of the experiment involved clustering techniques (collection-dependent vs. DDC) and different highlighting methods (showing only the best match, showing only potential matches and showing matches along with all the other books). Testing with different combinations helped us to identify the following results relevant to HCI design:

- Need to coordinate the design of spatial arrangement of objects and user movement. CAVE-ETD did not provide side-stepping as movement in the virtual world, assuming that users would side-step in the real world space inside the CAVE when needed to browse a shelf. This did not happen, and instead users relied almost exclusively on movements in the virtual world to access books. This seems to indicate that users have problems coordinating the actions between real and virtual world that are affect both real and virtual space.
- Difference in browsing pattern, and influence of angular vs. linear distance. The partial match between spatial arrangement and movements along with the changes in the way titles were visible on top of the bookcases modified the way the search was performed, going from the linear scan expected in a library to a circular pattern centered on potential relevant books. Having a fixed position while browsing also favored this approach, since the angular distance between books (the minimum angle to change in order to select the next or previous books) got smaller as the book was farther from the user, which made them more difficult to select.

The CAVE Collaborative Console - A Shared Workspace

This project aimed at providing a shared space where users in different sites can interact while being aware of each other's relative position and actions. The CCC was developed on top of EVL's Limbo from the University of Chicago. While Limbo provided a shared space where every user is represented by an avatar, there was no provision to make the users aware of each other's position when they are out of sight, and did not have any capabilities to coordinate actions among avatars.

The need to know one's position in the world and relative to others became evident after a demo session is Supercomputing'98, where researchers at different sites met in a shared, virtual space. Due to the extension of that space, it was often the case that a newcomer couldn't see anyone either because people where beyond the horizon, or occluded by walls and constructions. As a result of that and because of the lack of other tools, people had to rely on verbal communication to discover common

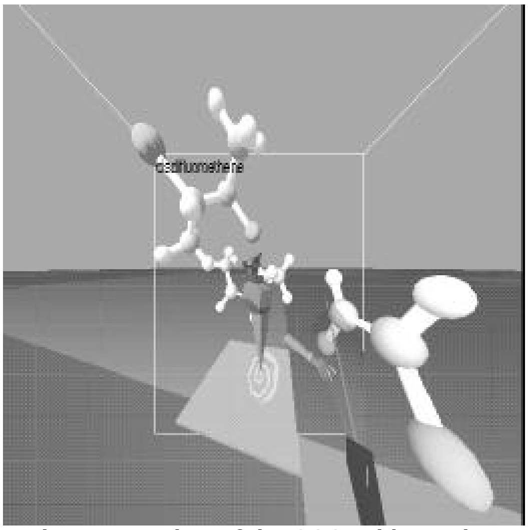

Figure 4. A View of the CCC, with another user's avatar at the center, and the user selecting a molecule.

Figure 5. Participant List and the 2D Radar View. An avatar is visible from a distance.

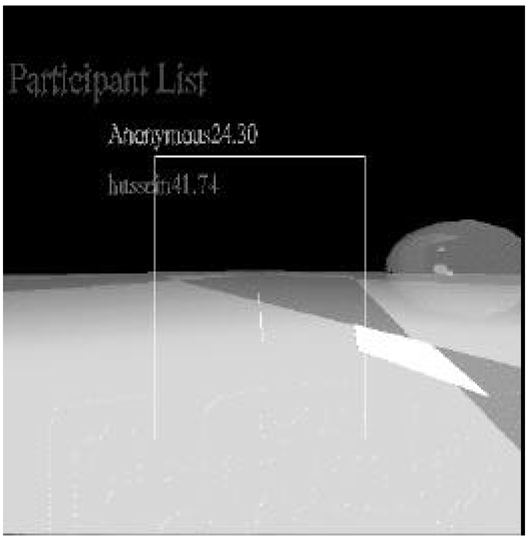

Figure 6. Participant List and 3D Radar View.

landmarks and describe their relative position to each other. Needless to say, this procedure was time-consuming and frustrating.

- CCC adds three tools to support awareness: The Participant List, the 3D radar and the 2D radar. The Participant List displays a list of the name of the users that are present in the world, plus the distance to each avatar from one's current position. Every user name is colored differently, and the same color is used in the radars to represent the same user.
- The 3D radar is an egocentric representation of the avatars in the world, where the user is at the radar center, and each blip represent a user currently in the world. Color of each blip is the same as the color of the user's name in the Participant List. Position of the avatar representations in the radar is continuously updated as avatars or the user moves around the world in any direction.
- The 2D radar is a flat representation of the 3D radar, corresponding to a projection on the $z=0$ plane of all the avatar positions. We added the 2D radar because many people found it easier to understand than the 3D radar, especially because being a flat view, perspective does not distort the representation of distance as avatars move farther away from the user.

CCC capabilities evolved beyond the awareness tools to support avatar coordination. We allow a user to

- Jump next to another user, so a group with a common activity can quickly gather around the activity organizer
- Tether to another user, so a more knowledgeable user can lead a tour through the world and show the most prominent landmarks. While a user is tethered to a leader, he or she will follow the leader wherever he goes, although it is free to detach at any point in time or to look around while the being toured.
- See through other user's eyes. We discovered the need for this capability while testing a distributed class, where the teacher was describing a feature in a small spot and it was difficult for many avatars to gather around the teacher without interfering with each other. Instead, we allow the student to see though the teacher's eyes. When a student request to see through the teacher's eyes, he or she sees at the eye level whatever the teacher is looking at. He or she cannot walk or grab objects in this mode to avoid conflicts, although he can look around from the position of the head of the user he is connected to. All avatars whose users are now seeing though somebody else's eyes remain in the place they were before the user started to share a view and until the user returns to the avatar. A sign on the avatar indicates the user is not currently available for interaction.

Although the CCC has not yet formally been tested as CAVEE-ETD was, it had an iterative development cycle where new features where evaluated by two groups of users: architects and Virginia magnet school teachers. The latter are particularly involved in the design since the CCC will be used and evaluated this summer for a distributed chemistry class with K12 students.

Both groups expressed the need to be able to collaborate with the CAVE using desktop machines. We supported this by designing CCC decoupling input from input interpretation. While in the CAVE, input is performed by using the wand and voice commands, while in desktop machines input is performed with the mouse, keyboard and a floating menu bar. All options available via voice and menu, and are coherent in naming and results.

Having these two separate input modes working coherently while being tested for the particular environment they were developed for proved to be very successful, requiring very little training for a teacher used to the desktop interface to move to the CAVE. As a result, we currently have a community of CCC users with different settings and quality of environments, including CAVE, stereo displays and standard monitors.

Conclusions

We firmly believe that CAVE-ETD and CCC prove that is necessary and possible to identify and adapt the Virtual Environment to the limitations and differences imposed by technology, and that doing that is needed in order to improve de usability and support widespread acceptance of virtual environments.We plan on continuing evaluating the usage of both environments in order to more precisely identify how to support different settings without affecting the usability of the virtual environments.