Evaluation of a Collaborative Volume Rendering Application in a Distributed Virtual Environment

U. Wössner, J.P. Schulze, S.P. Walz, U. Lang

High Performance Computing Center, University of Stuttgart, Germany
{woessner,schulze,walz,lang}@hlrs.de

Abstract

In this paper, we present a collaborative volume rendering application which can be used in distributed virtual environments. The application allows the users to collaboratively view volumetric data and manipulate the transfer functions. Furthermore, 3D markers can be used to support communication. The collaborative setup includes a full duplex audio channel between the virtual environments. The developed software was evaluated with external users who were asked to solve tasks in two scenarios which resembled real-world situations from the medical field: a presentation and a time-constrained search task. For the evaluation, two 4-sided CAVE-like virtual environments were linked. The collaborative application was analyzed for both technical and social aspects.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism – Virtual Reality; I.3.6 [Computer Graphics]: Methodology and Techniques – Interaction Techniques; C.2.4 [Computer-Communication Networks]: Distributed Systems – Distributed Applications

1. Introduction

Virtual environments are well suited for the exploration of large datasets, and they are especially useful to understand complex three dimensional structures like those which occur in volumetric datasets. Collaborative work with large datasets between remote locations is an increasingly important aspect in today's fast paced scientific and professional world. The software system, which we built allows to work in collaborative virtual environments (CVEs) with 3D datasets using regular Internet connections. Also, the system features direct rendering of volumetric datasets in virtual environments while keeping the visual appearance synchronized between the cooperating locations. Direct volume rendering is especially useful for medical datasets, e.g., data acquired from computed tomography scans.

Collaborative work requires a number of mechanisms for the participating parties to work together efficiently. One concept is to provide suitable coupling strategies for different use cases. In our system, three types of coupling are implemented: Master/Slave, Tight and Loose Coupling.

In order to evaluate the system's usability and the implemented coupling strategies, we developed two scenarios in which we asked volunteers to work. Both scenarios were taken from the medical world. One was the presentation of a dataset by an expert to a novice, the second scenario assumed both participants to be experts with respect to the software system, and they were asked to solve a specific task within a given time frame.

The evaluation was conducted in two CAVE-like virtual environments: the CUBE at our research lab, and the CAVEEE, which is located at the neighboring Fraunhofer Institute.

The main goal of this paper is to show that collaborative volume rendering is feasible with the software system we developed, and we want to point out remaining usability issues.

Related work will be summarized in Section 2. In Section 3, we describe the collaborative system, the volume rendering application and its improvements compared to the version used in the previous evaluation. These include usability improvements and collaborative extensions. In Section 4, we evaluate the system. Then
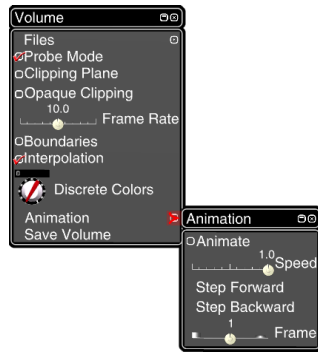
Figure 1: *Menu System*

follow the results in Section 5 and our conclusions in Section 6.

2. Related Work

There are a number of publications in the field of shared virtual environments which focus on scalability, interest management and consistency. These systems include DIVE [4], NPSNET [10], and MASSIVE [6]. Online 3D communities like SPLINE [1] or Blaxxun [7], and multiplayer-games also have greatly influenced the development of shared virtual environments.

Not much work has been done in the field of collaborative volume rendering. [11] describes a collaborative tool for virtual endoscopy, but it cannot be used in virtual environments.

Our collaborative application is based on the visualization framework COVISE that is described in [13]. Details about the integrated VR renderer COVER can be found in [12]. The volume rendering capabilities, which were needed to work with medical data, were added by integrating the library VIRVO and its 3D user interface. A description and usability evaluation of that work can be found in [14]. It was used as the basis for our new developments and for the evaluation described in this paper.

The type of usability engineering we used is mainly based on a comprehensive taxonomy of usability characteristics in virtual environments [5]. The evaluation has been inspired by suggestions in [2] and [15], where evaluations are separated into two phases: a usability inspection and a formative evaluation with task-based scenarios.

3. The Collaborative Software System

Our volume rendering application is integrated into the collaborative environment COVISE, which was developed by our group. This section describes the parts of the system needed to understand the rest of the paper.

3.1. Visualization Framework

COVISE is a modular visualization system, designed to support collaborative visualization of data in virtual environments as well as on the desktop. Included in COVISE is the virtual environment renderer COVER, which offers full VRML97 capabilities and provides a flexible plugin system. To support volume rendering, a plugin was developed which uses VIRVO to render volume data.

3.2. Collaborative GUI Library

In order to add collaborative features and fix existing issues of the volume rendering application, we decided to implement a complete GUI Library (VRUI) that solves two problems at the same time. First, this gives us much more flexibility in modifying the layout of the current user interface and second, it gives us the possibility to implement collaborative features from the beginning at a low level, so that they do not have to be implemented separately in all applications.

The API of VRUI resembles that of conventional GUI libraries like QT or Java AWT. Basic UI elements like buttons, labels and sliders (see Table 1) can be arranged manually or automatically by layout managers, which we call containers. Layout managers can be nested to allow complex dialog arrangements. Special containers provide backgrounds and frames around groups of UI elements. By combining these basic building blocks, menu classes have been developed that allow easy creation of hierarchical menu systems (see figure Figure 1). Custom popup dialogs are also easy to create as can be seen in Figure 2.

Events are processed by event listeners that can be attached to GUI elements to trigger events.

All UI elements can be derived from a class that implements support for collaborative work. An application then does not need to care about synchronizing each individual feature, as the GUI library automatically distributes events from local user interactions to all remote partners. The remote applications will then process these events locally and thus do not need to implement special code for the collaboration. Often a user interface element should not be synchronized all the time, so synchronization can be enabled and disabled, for example depending on the collaboration mode.

To allow concurrent interaction of multiple users with one UI element, the collaborative UI-class locks all corresponding remote elements during an ongoing interaction. A handshake would be required to ensure that two locking requests do not pass each other on the network. A handshake though would be much too slow in situations with high latency, like between different continents. In our approach, we do not wait for a confirmation of the locking request, but interrupt one of the interactions if a double lock was detected.

Basic UI Elements	Containers
Label	RowContainer
Button	Panel
PushButton	Frame
ToggleButton	ColoredBackground
Slider	TexturedBackground
Poti	DropdownHandle

Table 1: *User Interface Elements*

3.3. The Volume Rendering Application

Our volume rendering system allows the user to assign colors and opacities (alpha values) to scalar volume data. This assignment is represented by transfer functions. VIRVO provides a 3D GUI, the Transfer Function Editor (TFE, see Figure 2), to accomplish this.

Other functionalities of the volume rendering plugin can be selected either in a 3D Menu or by gestures. Our evaluation participants were introduced to the following features:

- Probe Mode: in this mode only a cubic sub-volume (the Probe) is visible, the surrounding dataset is transparent. The user can move around and resize the Probe with the input device.

- Clipping Plane: cuts off a part of the dataset.

- High Quality Mode: when a mouse button is clicked whilst the mouse is held above the head, the dataset is displayed in a higher image quality. Due to system performance limitations, this results in a lower frame rate.

At any given time, when working in a virtual world, cone-shaped markers (see Figure 3) can be placed in the 3D scene to point to positions of interest. They remain fixed in object space and their locations are synchronized in all participating environments, so that

Figure 2: *Transfer Function Editor*

Figure 3: *Marker*

they can be used to discuss small features of a dataset.

Since the previous evaluation, all volume rendering related interaction elements were extended for cooperative use. This includes all interactors of the TFE, the Probe Mode, and the Clipping Plane. Some of this functionality was introduced along with the new collaborative GUI library described above.

The following two sections describe improvements of the user interface, which were implemented upon user requests from the previous evaluation.

3.3.1. General Transfer Function Editor Improvements

A detailed description of the TFE's basic features can be found in [14]. These features include interactive manipulation of the opacity function by placing Pins (represented as vertical lines) in a 2D transfer function representation. Three types of Pins can be used to define the opacity function: trapezoids, ramps, and blanks. The color function can be defined by assigning specific colors to a few scalar values. The colors for the remaining scalar values are then interpolated linearly in RGB space between the neighboring color Pins.

Our last evaluation pointed out several usability issues of the TFE which were improved in the current version:

- The mouse button assignment was confusing because separate buttons were used for selection and manipulation of Pins. Now both can be done with the same mouse button.

- Imprecise Pin positioning: users weren't aware of the exact Pin locations. Now the current scalar value is displayed at the bottom of the selected Pin.

- Users asked for a trapezoidal Pin type. We enhanced the peak Pin by a width parameter, so that it can be widened to the shape of a trapezoid.

- Previously, the color bar displayed the transfer function colors and opacities in one bar. Thus, users couldn't see the colors of transparent regions. Now the color bar is divided into two parts: one displays only the color, the other additionally represents the opacity.

- After deletion of all color pins, the volume turned black even if opaque regions existed in the transfer function. In the current version, the default color is white, to make the volume visible on black projection walls.
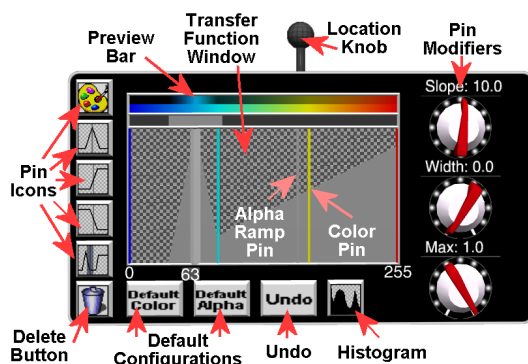
- The rotary knobs in the Transfer Function Editor were hard to use for a precise setting of values, because they transformed hand rotations 1:1 into knob rotations. Therefore, we added a fine-tuning mode, used by turning the knob with the right mouse button pressed. Now, angular changes of the hand are converted into knob rotations at 1:10. The left mouse button can still be used to do 1:1 rotations.

- Sometimes, users had requested an undo function. This we implemented to stores the last 20 actions which affected the transfer function definition. Simple clicks on the Undo button go back one step at a time.

- Since some users had asked for a histogram of the scalar values in the volume, we added a button, which toggles the display of a 2D histogram in the function field.

- Some users were irritated by the TFE's background color, which was a pattern of black and white dots, because it made text difficult to read. For the current version, we changed the background color to solid black.

3.3.2. Other Improvements

In the previous evaluation, some people were confused by the fact that, when they placed a clipping plane somewhere in the volume, they weren't aware of the exact clipping plane position. This mostly occurred if the volume contained large semi-transparent regions whose exact positions in 3D space are hard to be located by the human eye. Therefore, we implemented an Opaque Clipping Plane mode in which only the volume data directly on the plane is displayed and the opacity function is set to the maximum for all scalar values.

3.4. Collaboration Modes

For the development of a multi-purpose collaboration system, certain use cases for such an environment have been defined. Real-life working situations which involve many people, include presentations, training, problem solving, data exploration, and combinations of these. In order to suit the needs of all of the above situations and not overwhelm the users with too many options, we implemented three basic collaboration modes: Loose Coupling, Tight Coupling, and Master/Slave Coupling. They are described in the following subsections.

3.4.1. Loose Coupling

In this mode, the users' views are not synchronized at all; partners can move through one shared world independently. Probe Mode and Clipping Planes are not synchronized, while the overall state of the virtual world, e.g., dataset, transfer function, or the position of boxes is coupled. In this mode, users are represented as avatars (see Figure 4) that directly show the locations of
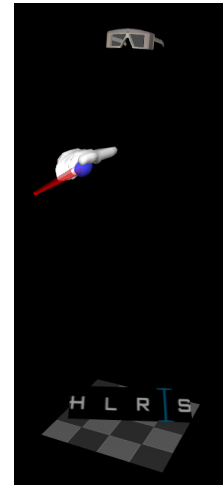

Figure 4: *Avatar*

the persons they represent. The participants may scale their world independently. The scaling is visualized by the size of the avatar. This means that people looking at a world in miniature format appear to others as large avatars standing outside of the object, looking down on it. The person standing outside will see small avatars moving through the miniature world.

3.4.2. Tight Coupling

Here, viewpoints of all partners are synchronized. If one of the participants moves or scales the world, all other views are updated accordingly. In detail, synchronized views mean that the positions of the projection environments are synchronized. The users are still able to physically move around in their projection environment and thus have a slightly different view on the data object due to head tracking. All elements of the user interface, Probe Mode and Clipping Planes are also tightly coupled.

3.4.3. Master/Slave Coupling

This mode is similar to Tight Coupling Mode except that only one participant, the master, is allowed to interact with the world.

3.5. Avatars

The avatar is a minimalistic representation of the user: it consists only of a pair of glasses, one hand, plus a ground level (gray chess board). The shape of the avatar is a compromise between full person-like bodies and small iconic representations such as simple name plates. Solid bodies have the advantage of good visibility but they hide a lot of objects when they are scaled up. It is furthermore impossible to correctly animate a full body avatar with the limited information acquired by the tracking system (head and hand position and their
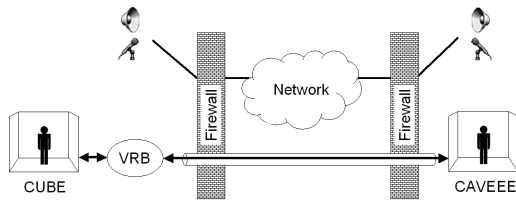
Figure 5: *System Setup*

orientation). We decided to display only the tracked features (glasses and pointer) and the feet, and use the users' imagination to fill the gaps in-between. The feet are represented by a gray chess board at the height of the floor, straightly underneath the head position. A name plate is attached to it to distinguish avatars in collaborations with more than one partner.

As [9] and our own experience showed, this limited set of features gives enough visual cues to work together efficiently. The avatar can transmit a rich set of gestures to gain your attention or to show what people are doing. You can see where a person is standing and what he or she is looking or pointing at, and you can see if people are paying attention to you or looking around.

4. Evaluation

After we had implemented the collaborative enhancements and GUI improvements, the system was evaluated in a similar fashion as in our previous work [14]. The major difference is that this time we had to use two virtual environments in order to test the collaborative features.

This section first gives a short overview of the last evaluation, it describes the hardware setup used, and it describes the evaluation process itself.

4.1. Previous Evaluation

The evaluation of the non-collaborative volume rendering system was derived from the ideas in [15] and
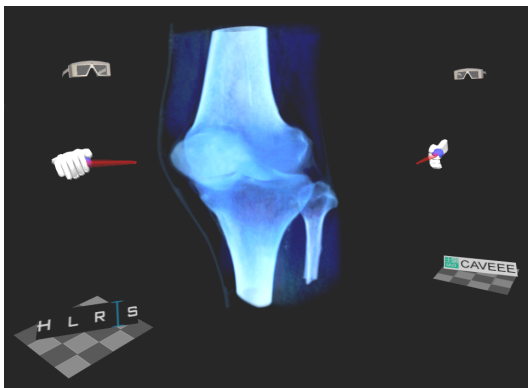

Figure 6: *Visible Human Knee With Avatars*

consisted of two phases: a usability inspection and a scenario-based evaluation. Twelve participants first had to fill out a self-rating questionnaire, then they were introduced to the software system. After they had spent a total time of 30 minutes in the CUBE, they were again asked to fill out a questionnaire, and finally we interviewed them about their experiences. Most of the participants had worked in CAVE-like environments before, and some had previous experience with volume rendering.

Most of the users were satisfied with the overall usability of the system. Among the usability issues found were some minor flaws with the Transfer Function Editor. Most of them were addressed in Section 3.3.1, since they had been eliminated in the meantime.

4.2. The Collaborative Hardware Setup

The hardware setup used in the evaluation consisted of two 4-sided CAVE-like [3] virtual environments (the CUBE at HLRS and the CAVEEE at Fraunhofer IAO), which were connected by both a direct Ethernet connection to synchronize the virtual environments and a regular multicast capable WAN for the audio conferencing. The participants were given wireless microphones, and loudspeakers were placed next to the VEs.

4.3. Scenarios

We designed two task-oriented social situations to represent typical collaborative scenarios. The first situation is a presentation. It occurs whenever there is an expert who wants to instruct a novice, as in cases of training or in the field of education. Another example would be surgical planning where a remotely located doctor who did the diagnosis, instructs a surgeon. The second situation we considered is the joint work of two people who want to benefit from their combined expert knowledge in order to solve a difficult problem.

4.3.1. Presentation

In the presentation scenario, the experts present a volume rendered computed tomography skull dataset (see color section), which was extracted from the National Library of Medicine's Visible Human Project data. Before the actual work in the CVEs took place, the experts were given detailed information about the dataset's features, which they were then to explain to the novices.

4.3.2. Joint Work

In the second scenario, a knee joint, which was also extracted from the Visible Human dataset, was loaded into the volume rendering system (see Figure 6). The dataset was altered to contain a needle at a location

unknown to the participants (in a bone as shown in Figure 7). The needle was represented by a simple straight line of volume elements with the same scalar values as bone, which could only be seen after precise modifications of the transfer functions. The participants had to collaborate to find the needle as fast as possible, and they should choose whatever collaboration mode seemed appropriate.

4.4. Evaluation Process

The evaluation was based on pairs of participants working together in two networked CVEs. The actual evaluation took place on three consecutive days and had a mean total evaluation time of three hours per pair.

4.4.1. Participants

We invited five pairs of people who never have met before. Most of the 10 participants were members of neighboring research institutes. Of each pair, one participant was assigned the role of the expert, the other one was going to be the novice. The experts were required to have experience in the use of virtual environments, in volume rendering, or in both. Figure 8 shows the participants' experience as it was rated by themselves.

4.4.2. Task Preparation

Participants were initially taken to two separate rooms for briefings. The briefings consisted of slide shows explaining the collaborative software system and the tasks which had to be solved in the CVEs. The expert was given additional knowledge about the dataset used in the presentation while the novice was not.

After the slide shows, the participants filled out a self rating questionnaire asking for their previous experience with virtual environments and related knowledge. Furthermore, they were asked about their age, eyesight restrictions (glasses, contact lenses etc.) and their handedness.
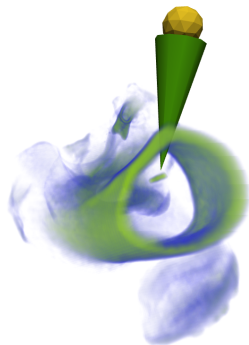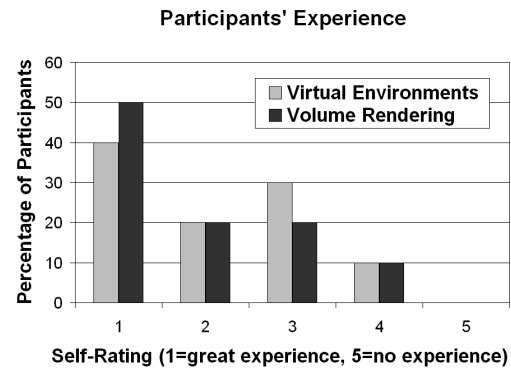

Figure 7: *Marker Pointing to Needle*

Figure 8: *Participants' Experience*

Then they were guided to the CUBE and to the CAVEEE, still not having met. Both of them were given short introductions to the GUI and to the collaboration modes. Unlike the novice, the expert was given the opportunity to get to know the dataset used in the presentation, and he was given suggestions about what to explain to the novice in the first scenario. This part of the evaluation took about 15 minutes.

4.4.3. Participant Observation and Video Recording

When both of the participants felt comfortable with the user interface, they were hooked up to the audio system and from then on could communicate with each other. Also, video recording of the session was started. Two tasks had to be solved, each within a time limit of 15 minutes. During all of this time, the participants could ask the observers questions about the user interface and the collaboration modes. If after about 10 minutes, there was no progress with the search task, the observers began to give hints about the location of the needle. During the time in the VEs, the observers made notes about the participants' questions and their performances.

4.4.4. Questionnaires

After the tasks had been solved, the participants went back to the two briefing rooms. There they filled out questionnaires asking for different aspects of social collaboration and technical issues. The questionnaires consisted of two parts: multiple choice questions and questions for free answers.

The multiple choice questions were grouped into the following sections: general impression of the collaborative system, details about virtual worlds and avatars, interaction with the system, cooperation modes and interaction with the collaborator, and technical questions about the transfer function editor and the other interaction elements.

The free questions asked for general ideas for improved usability of the cooperation modes, what other cooperation modes could be offered, and which

information the participants would have needed about each other to solve their tasks faster. Finally, general comments were noted. This part of the evaluation lasted about 20 minutes.

4.4.5. Interviews

After the questionnaires were filled out, the participants met in one of the briefing rooms. There they were questioned about their experiences in the CVEs. The interviews followed pre-designed guidelines and lasted up to one hour.

4.5. Analysis of the Evaluation Data

The analysis of the collaboration was done at two levels: one level covered technical aspects and issues the participants were faced with. We were interested in how the users interacted with the system and how they wanted to reach the technical goals of the tasks.

The other level of analysis dealt with social aspects of the collaboration, how the participants interacted with each other. For this analysis we used methods from psychological discourse analysis and sequential film analysis. Discourse analysis extracts hidden conceptions that are not easily measurable, like subtle expressions in body language. Sequential film analysis delivers micro-views of key scenes from the evaluation footage [8].

The evaluation data was analyzed in the following order:

- The interviews were fully transcribed.

- Empirical data from the questionnaires was charted and statistically processed.

- The two video recordings were assembled on a split-screen montage (see color section). They were then inspected: key scenes were captured, sequenced, and fully transcribed according to findings in the interviews and questionnaires.

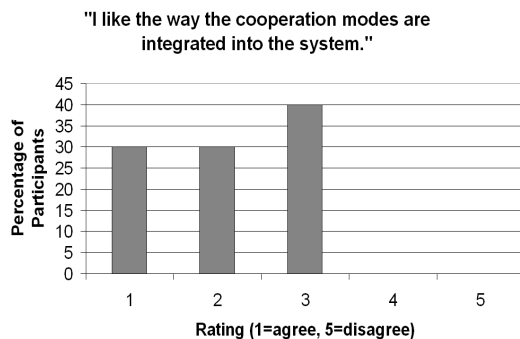- Notes from the participant observation were counterchecked with the above findings.
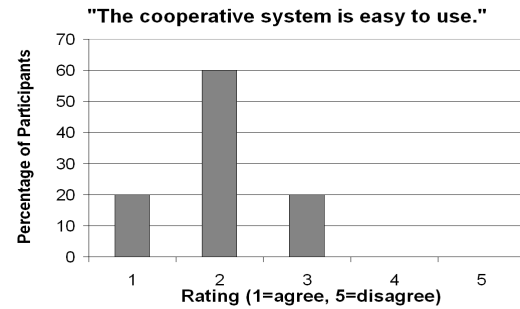

Figure 10: *System Integration*

Figure 9: *Usability*

5. Results

The analysis of all the data gathered in the evaluation led to a number of usability issues of the developed collaborative visualization system. These issues can be viewed from multiple perspectives: system related, GUI related, and social issues.

5.1. System

In general, the users thought that the cooperative system was easy to use (see Figure 9). They also liked the way the cooperation modes were integrated into the system (see Figure 10). The selection of collaboration modes was not directly linked to the current working scenario, it was more dependent on the users' motivation or method to tackle a problem. Most of the users stated in the questionnaire that they preferred the Loose Coupling mode, but they all used different modes in different situations. Some of the pairs used the Loose Coupling mode for most of the time in both scenarios while others used the Tight Coupling mode also for the second scenario.

One important user request was for a No Coupling mode. In Loose Coupling mode, the users have independent viewpoints, but changes of the volume dataset are visible to all users. Sometimes the evaluation participants wanted to work with the dataset independently of their partners so as to not interfere with or be affected by their partners' actions. Sophisticated ways of comparing and merging the two different application states would be needed when switching back to a tighter coupling mode after a period of independent work.

Another frequent problem was that the users were not entirely aware of which parts of the system were synchronized in the different collaboration modes. During the briefing they were shown a table with exactly this information, but it turned out to be too difficult to memorize all possible combinations.

All but one user were happy with the markers as they were implemented, and they used them frequently to discuss specific parts of the datasets. Some users
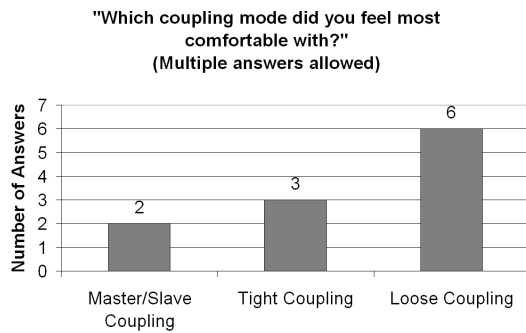
"Which coupling mode did you feel most comfortable with?"
(Multiple answers allowed)

Figure 11: *Coupling Modes*

suggested the use of different colors or attach labels to the markers if more than one was present, so that during the collaboration they could be distinguished more easily.

Half the users were satisfied with the avatar representation. The others did not have any complaints about the functionality of the avatar but they did not like its design. By analyzing the video footage, we could verify how participants explicitly used the avatar to direct the partner or correct him. One participant had problems locating an avatar once but it is unclear if a more solid avatar would have solved the problem.

5.2. GUI

Several usability issues from the previous evaluation have been subject to improvement (see Section 3.3.1). For the latest evaluation, the 10 users were asked how they liked the improved features:

- Contrary to the last evaluation, the participants were not confused by the mouse button assignment for the Pin usage. Previously, two mouse buttons had to be used for selection and manipulation, now one button can do both. Only 2 out of 10 users said they had problems with the new usage scheme, while previously 7 out of 12 were confused with the button assignment.

- Previously the Pin positioning was considered to be imprecise because there was no feedback for the exact Pin position. After the display of the currently set scalar value, only 3 users desired more information.

- The new trapezoidal Pin was accepted quickly, 8 users found it easy to use.

- The color bar division into one with and one without opacity representation, compared to only one bar with opacity, was considered helpful for all participants.

- 8 users liked the new fine tuning mode for rotary knobs, activated when the knobs were turned with the right mouse button pressed.

- 9 users were perfectly happy with the TFE's new solid black background.

- The new undo button was found to be useful by all but one user, who said it did not work as expected.

- 8 users thought the histogram display was useful.

5.3. Social

Since the software system we evaluated can be used collaboratively, it is important for us to not only address technical issues but also social issues and observations.

Figure 11 shows that several participants were uncomfortable with Tight Coupling and Master/Slave Coupling. When asked, most of them stated that they did not want to disturb their partners by moving around their datasets, and thus they preferred Loose Coupling. Only 2 out of the 10 participants felt most comfortable with Master/Slave Coupling.

A flawless audio connection proved to be crucial for the collaboration during the solving of the tasks. Once in a while the battery of a wireless microphone went flat. Since there was no warning before the connection terminated, there was a sudden moment of silence at the other end before the battery could be replaced. Whenever this occurred the workflow broke down completely and the effect of immersion in the virtual world was gone. It took a few minutes before the previously reached level was re-established.

After having reviewed the experiments, we were surprised at the following unexpected results:

- Even though the experimenters decided who were going to be experts and novices, the participants were always happy with the decisions. We had expected some people to not feel comfortable with their roles and to attempt a swap.

- Before the experiments, the participants only received weak guidelines on how to reach their goals in the two tasks. Nevertheless, all the pairs arranged their strategies that they were going to follow very quickly in the two scenarios. This happened although the participants were unknown to each other.

- Not only was the task solving well planned by the participants, it was also highly interactive: for example, during the search for the needle, most pairs showed each other views of the dataset in which they suspected the needle to be located. Then they discussed the views and alternately tried to make the needle more clearly visible by adjusting the transfer functions.

- During the presentation, we would have expected that, due to the spatial distance, it would have been more like a lecture and not very interactive. But in our experiments, the participants frequently switched back and forth between the collaboration

modes. The experts allowed the novices to practice the things they explained and to explore the dataset by themselves, accompanied by comments from the experts.

6. Conclusions

We presented a collaborative volume rendering system which can be used in distributed virtual environments. The system allows the manipulation of the transfer functions and arbitrary viewpoints with respect to the volume data. Three collaboration modes provide suitable interaction strategies for the typical contexts in which a collaborative system can be used. Avatars give the users the necessary visual cues for an increased awareness of the collaborators' presence.

The collaborative system was evaluated by a number of people from different professional fields. The evaluation consisted of two task based scenarios which represented typical real-world tasks. The participants were generally satisfied with the usability of the system.

Even though or especially because there was no physical presence of the partner, the collaboration worked very well. Due to the fact that the users did not have to share one physical device, they worked in parallel, much more than they would have in a standard face-to-face collaboration.

Together with reliable networks and working audio conferencing infrastructure, collaborative virtual environments can be a very useful tool for many scientific applications.

7. Future Work

Although our virtual environment proved to be generally usable, several issues remain. A non-coupled collaboration mode should be implemented, which would allow the users to work entirely by themselves. This mode could be extended by a feature allowing the users to temporarily switch to other collaborators' views in order to know what exactly they are doing.

Sometimes the evaluation participants were not sure about mouse button assignments or they wanted to know details about the coupling modes. In these cases, an online help system could have done what the evaluation mediators had to do.

A session management would improve the awareness of the collaborators' presence in the virtual environment. This would be especially helpful in sessions with more than two users. A simple list of active participants would already help, sophisticated graphical cues like virtual radar or a world-in-hand type display of the remote collaborators would probably be even better.

For a professionally and regularly used collaboration system, it would be important to provide different subsets of the functionalities for each profession and hide unnecessary features so as not to distract the users from their tasks. This requirement could be solved by a configurable user interface.

8. Acknowledgments

9. References

[1] Barrus, J.W., Waters, R.C., Anderson, D.B., *Locales: Supporting large multiuser virtual environments*, IEEE Computer Graphics Applications 16.6, Nov. 1997, 1997, 50-57.

[2] Bowman, D.A., *Interaction Techniques for Immersive Virtual Environments: Design, Evaluation, and Application*, Human-Computer Interaction Consortium Conference '98 (HCIC), 1998

[3] Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., *Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE*, ACM SIGGRAPH '93 Proceedings, 1993, 135-142.

[4] Frécon, E., Stenius, M., *DIVE: A Scalable Network Architecture for Distributed Virtual Environments*, Distributed Systems Engineering Journal 5, 1998, 91-100.

[5] Gabbard, J.L., Hix, D., *A Taxonomy of Usability Characteristics in Virtual Environments*, Final Report to the Office of Naval Research , 1997

[6] Greenhalgh, C., Purbrick, J., Snowdown, D., *Inside MASSIVE-3: Flexible Support for Data Consistency and World Structuring*, Proceedings of Collaborative Virtual Environments, Sept. 2000, San Francisco, 2000, 119-127.

[7] Interactive AG, B., *Web Site: http://www.blaxxun.de*, Munich, Germany, 2001

[8] Korte, H., *Einführung in die Systematische Filmanalyse*, Ein Arbeitsbuch, 2nd edition, 245 pages, Erich Schmidt Verlag, Berlin, ISBN 3-503-06115-0, 2001

[9] Leigh, J., Park, K., Kenyon, R., Wong, H., *Preliminary STAR TAP Tele-Immersion Experiments Between Chicago and Singapore*, Proceedings of HPCAsia '98, Singapore, 1998, 687-693.

[10] Macedonia, M.R., Zyda, M.J., Pratt, D.R., Brutzman, R., Donald, P., Barham, P.T.,

Exploiting Reality With Multicast Groups: A Network Architecture for Large-Scale Virtual Environments, Proceedings of IEEE Virtual Reality Annual International Symposium (VRAIS '95), North Carolina, IEEE Computer Society Press, 1995, 2-10.

[11] Niemeier, R., Fleiter, T., Benölken, P., Lang, U., Sokiranski, R., *An Advanced Collaborative Tool for Virtual Endoscopy*, Proceedings of the 12th International Symposium and Exhibition for Computer Assisted Radiology (CAR) '98, Tokyo, 1998, 118-123.

[12] Rantzau, D., Frank, K., Lang, U., Rainer, D., Wössner, U., *COVISE in the CUBE: An Environment for Analyzing Large and Complex Simulation Data*, Proc. 2nd Workshop on Immersive Projection Technology (IPTW '98), Ames, Iowa, 1998

[13] Rantzau, D., Lang, U., Rühle, R., *Collaborative and Interactive Visualization in a Distributed High Performance Software Environment*, Proceedings of the International Workshop on High Performance Computing for Graphics and Visualization, Swansea, Wales, '96, 1996

[14] Schulze-Döbold, J., Wössner, U., Walz, S.P., Lang, U., *Volume Rendering in a Virtual Environment*, Proceedings of 5th IPTW and Eurographics Virtual Environments, Springer Verlag, 2001, 187-198.

[15] Swartz, K., Thakkar, U., Hix, D., Brady, R., *Evaluating the Usability of Crumbs: a Case Study of VE Usability Engineering*, Proceedings of the 3rd International Immersive Projection Technologies Workshop, May '99, Springer-Verlag, 1999, 243-252.
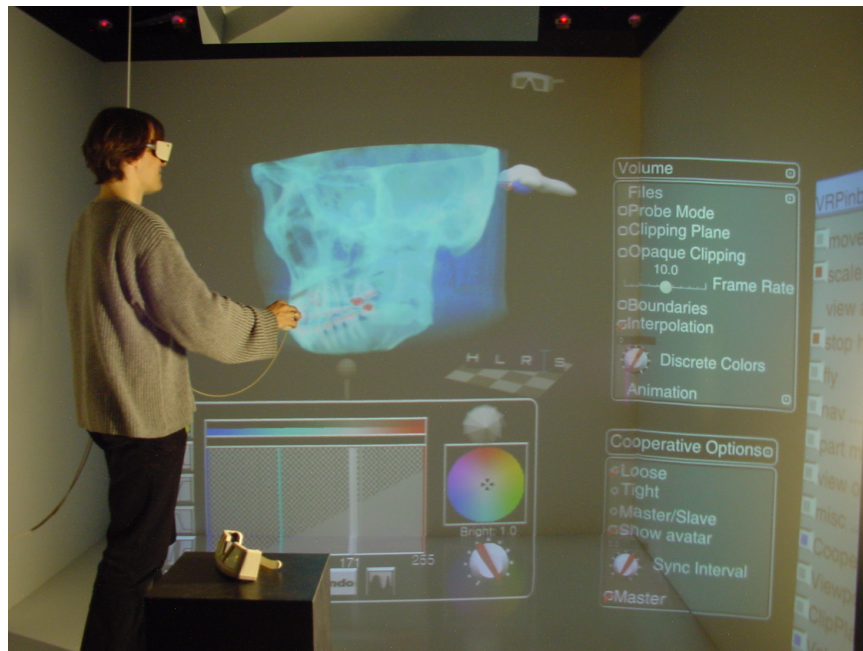
Visible Human Skull
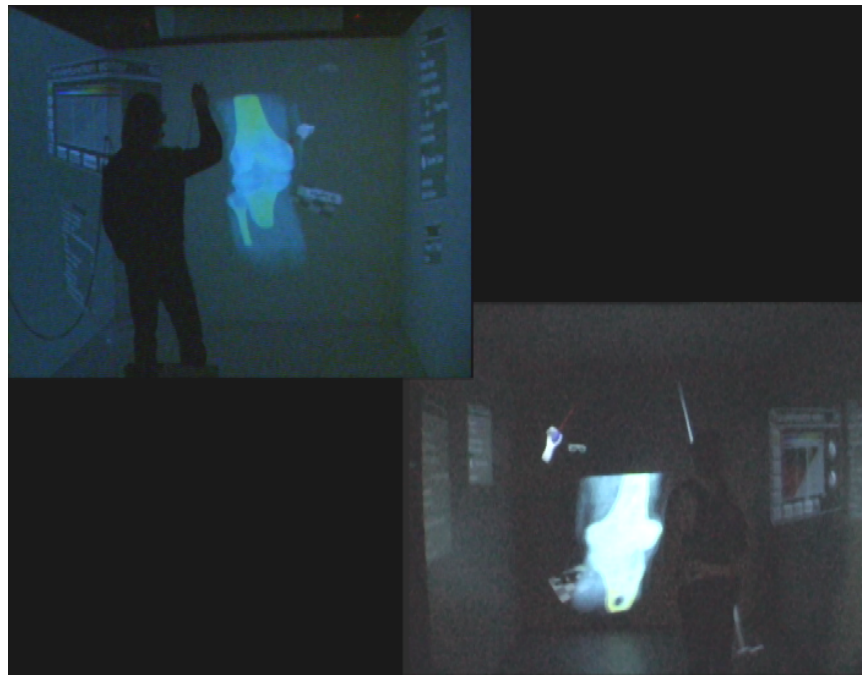
Split Screen Montage of Collaborative Video Footage