Exploiting Multiple Perspectives in Tele-Immersion

Kyoung S. Park, Abhinav Kapoor, Chris Scharver, Jason Leigh
{park, akapoor, scharver, spiff}@evl.uic.edu
Electronic Visualization Laboratory
1120 SEO, MC 154, 851 S. Morgan St.
University of Illinois at Chicago
Chicago, IL 60607, USA

ABSTRACT
The work in this paper describes a preliminary observational study conducted on users of CAVE6D, a collaborative CAVE-based virtual reality tool for visualizing multivariate oceanographic data sets. CAVE6D presents the concept of multiple perspectives by allowing participants to customize their views while working collaboratively and supporting the views either privately or globally. The goal of this study is to understand how tele-immersed participants cooperate when presented with multiple perspectives and to explore ways to leverage these perspectives to allow scientists to more rapidly interpret massive multi-dimensional data sets.

Keywords
Tele-immersion, Multiple Perspectives, CSCW

INTRODUCTION

Tele-immersion [5] is defined as the integration of persistent collaborative virtual reality with audio and video conferencing in the context of data-mining and significant computation. The ultimate goal of tele-immersion is not merely to reproduce a real face-to-face meeting in detail, but to provide the next generation interface for collaborators, world-wide, to work together in a virtual environment that is seamlessly enhanced by computation and large databases. When participants are tele-immersed, they are able to see and interact with each other in a shared virtual environment.

This shared environment may be the design of a new car, a visualization of climatological data, or other three dimensional environments that do not physically exist, or cannot be physically visited. The participants may be rendered as avatar. These participants are not merely talking about a thunderstorm, but they are standing inside it; they are not looking at a scale model of a new car design, but they are standing inside the full size engine block. By transmitting gestures as well as audio and video between the collaborators, these shared virtual environments give their users a greater sense of presence in the shared space than other collaborative media. By encouraging collaboration and conversation within the data, these environments may become the preferred place to work and interact even if more traditional face-to-face meetings are possible.

Our focus is on supporting both synchronous and asynchronous collaboration over trans-oceanic distances. Our model is that of a persistent virtual world where the virtual environment is sustained by a computer simulation that is left constantly running. This world continues to exist and evolve even when there are no participants – it may autonomously control supercomputing computations, query databases, or collect the results for visualization when the participants return.

Our users are the members of CAVERN [7,8] – the CAVE™ Research Network, a collection of participating industrial and research institutions equipped with CAVE® [1] and ImmersaDesk™ VR systems, and high-performance computing resources. High-speed networks connect them to support tele-immersive engineering and design; education and training; scientific visualization; and computational steering. With over 100 CAVE and ImmersaDesk installations around the world, one of the important problems facing this growing community is how to provide a mechanism to support long term collaborative work, from a technological as well as human factors point of view.

In the real world, individuals who are trying to solve a common problem gather (in workshops, for example) in the hopes that their combined experiences and expertise will contribute new perspectives and solutions to the problem. In most tele-immersive applications to-date, the default assumption has been to display the collaborative world in the same way to all its participants. We believe that by employing multiple perspectives and in particular by encouraging role-specialization, collaborators will be able to solve problems more effectively in tele-immersion. The following two examples will help motivate this concept:

- In a tele-immersive industrial design application, an engineer may be collaborating with an artist. They may be collaborating on the design of a new office chair. The artist can manipulate the design using tools that are specifically tailored for his/her expertise, such as interactive sculpting tools. Simultaneously the engineer can evaluate the impact of such design changes in terms of material stress and strain.
- In the context of scientific visualization we envision a potential application of multiple perspectives in the visualization of multi-dimensional data-sets. Here a large number of dimensions may be partitioned across multiple viewers to assist in reducing the overall complexity of the content being visualized. The challenge is in providing the necessary interface to support this collaboration while minimizing the confusion and additional cognitive load that may result from having to coordinate the activity of multiple viewers all simultaneously viewing disparate parts of the data-set.

This paper describes our application of multiple perspectives to collaborative virtual environments with a special focus on scientific visualization. CAVE6D [14] is a collaborative CAVE-based visualization tool for exploring multivariate oceanographic data sets (Figure 1). A unique feature of CAVE6D, as compared to other tele-immersive visualization tools, is its ability to allow participants to customize their views while working collaboratively. Hence even though the participants are all viewing the same data-set they are seeing the data from decidedly different physical perspectives as well as through different filters. CAVE6D also allows participants to activate visualization filters locally (and hence affecting only one's own view privately) or globally (affecting everyone's view).

The study we conducted serves three purposes: 1. it allows us to evaluate CAVE6D's user interface for supporting collaborative visualization; 2. it provides an opportunity for testing some initial ideas of how to support multiple perspectives in a tele-immersive environment; and 3. it allows us to explore the general problem space in supporting multiple perspectives.

EVALUATION OF THE USE OF MULTIPLE PERSPEC-TIVES IN A COLLABORATIVE SESSION

This observational user study constitutes our first exploratory design study as a precursor to future and more focused experiments on multiple perspectives. As such the goals and main contributions of this study are to explore the parameter space of issues involved in employing multiple perspectives.

Specifically the goals are to:

1. provide feedback to the developers of a collaborative visualization tool called CAVE6D;
2. inform us on the logistical intricacies involved in performing a user study of tele-immersed participants;
3. allow us to gain first-hand experience of how users behave when attempting to work cooperatively in tele-immersion;
4. allow us to explore the problem space in multiple perspectives in order to stimulate new hypotheses for future experiments.

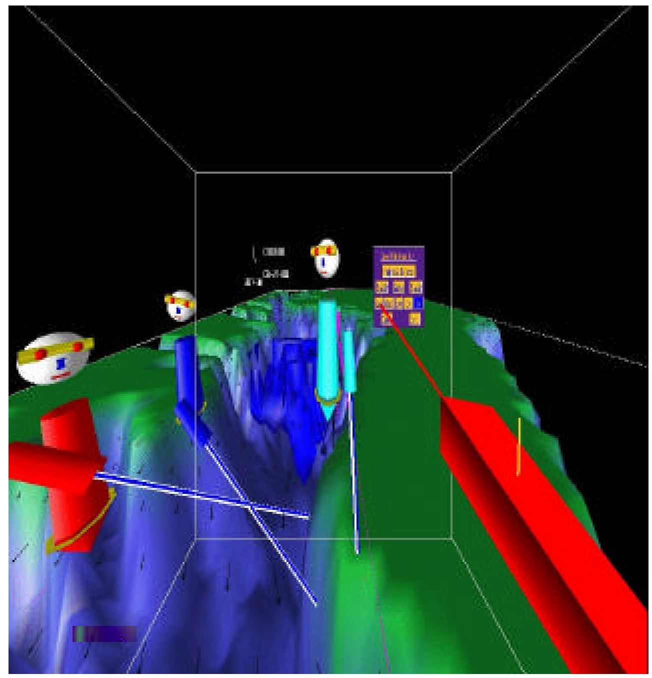

Figure 1. CAVE6D and the Chesapeake Bay data-set. Avatars that depict the users are equipped with long pointing-rays that can be used to point at features of interest in the data set.

METHOD

CAVE6D

CAVE6D, co-developed by Wheless and Lascara from the Center for Coastal and Physical Oceanography and Hibbard from the University of Wisconsin Madison, is a configurable VR application framework. CAVE5D [13] is supported by Vis5D [3], a very powerful graphics library that provides visualization techniques to display multi-dimensional numerical data from atmospheric, oceanographic, and other similar models, including iso-surfaces, contour slices, volume visualization, wind trajectory vectors, and various image projection formats.

CAVE6D is a tele-immersive extension of CAVE5D that allows multiple users of CAVE5D to jointly visualize, discuss and interact with the data-set while they converse over a telephonic conference call. Visualization parameters such as Salinity, Circulation Vectors, Temperature and Wind Velocity Slices, Larval Fish Distributions etc. that can be visualized by CAVE5D, have been extended in CAVE6D to allow participants to collectively operate them. The data-set being used in this study was a simulation of tidal patterns in the Chesapeake Bay.

CAVE6D allows participants to activate visualization filters locally (*and hence affecting only one's own view*) or globally (*propagating their changes to everyone*). This affords each participant the ability to individually customize and/or reduce the cluttered-ness of the relationship between one subset of the total dimensions of the data while another participant may be simultaneously correlating a different subset.

Time on the other-hand is globally shared and hence participants view all time-varying data synchronously. This is enforced to provide a common frame of reference for its collaborating users, to reduce confusion when coordinating between potentially divergent views.

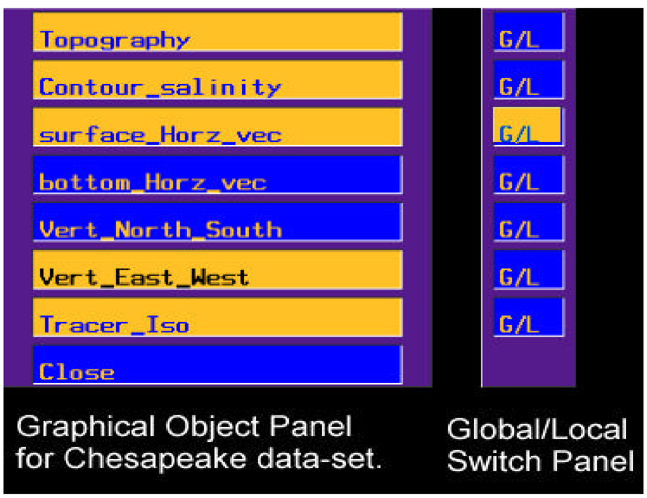

Figure 2. CAVE6D's menu interface for the Chesapeake Bay data-set.

Figure 2 shows the CAVE6D's menu interface, which has the graphical object panel and the global/local switch panel. The graphical object panel includes the available visualization tools for the Chesapeake Bay data-set. The corresponding menu on the global/local switch panel allows users to toggle between global and local settings. In this illustration, the surface horizontal vector is the only parameter that is being globally shared amongst all other users in CAVE6D session. The black text highlighting the vertical east/west vectors on the left panel indicates that the tool is interactively manipulable and it only affects to one's own view locally. Table 1 lists the visualization tools (or parameters) available in CAVE6D.

Topography	A solid or wire frame representation of the landmasses.
Contour Salinity	This visualizes salinity levels as contour lines. The magnitude of the salinity is depicted as a number next to the contour line.
Surface Horz Vec	This visualizes the velocity of the tide along the surface of the bay. This is interactively adjustable, allowing one to inspect tidal velocities at varying depths in the bay.
Bottom Horz Vec	Used in conjunction with Surface Horz Vec one can view the tidal velocities at both the deep and shallow regions of the bay.
Vert North South	This visualizes the tidal velocities cross sectional from North to South. This tool can be moved along the east/west axis.
Vert East West	This visualizes the tidal velocities cross-sectional from East to West. This tool can be moved along the north/south axis.
Tracer Iso	This visualizes salinity levels as iso-surfaces (3D contour diagrams). Red represents higher salinity levels and blue represents lower salinity levels.

Table 1. Descriptions of visualization tools in the Chesapeake Bay data-set in CAVE6D.

The Participants

Three pairs of graduate students in EVL participated as volunteers in this study. To minimize individual differences we chose to limit the study to computer science students who have already considerable experience with virtual reality. Furthermore, as this study would require students to work in pairs we tried, where possible, to pair students who already had an established working relationship with one another. Most were familiar with the concept of CVE, but none of them had prior experience with CAVE6D and the Chesapeake Bay data-set. In the future we intend to apply a refined version of this experiment to study real oceanographers or students of oceanography.

The Apparatus

As shown in Figure 3, each pair of participants collaborated with each other in their respective CAVEs, and a third participant (the instructor) would observe the proceedings on an ImmersaDesk. All of them could speak to each other via a high quality audio connection using headphones and mi-

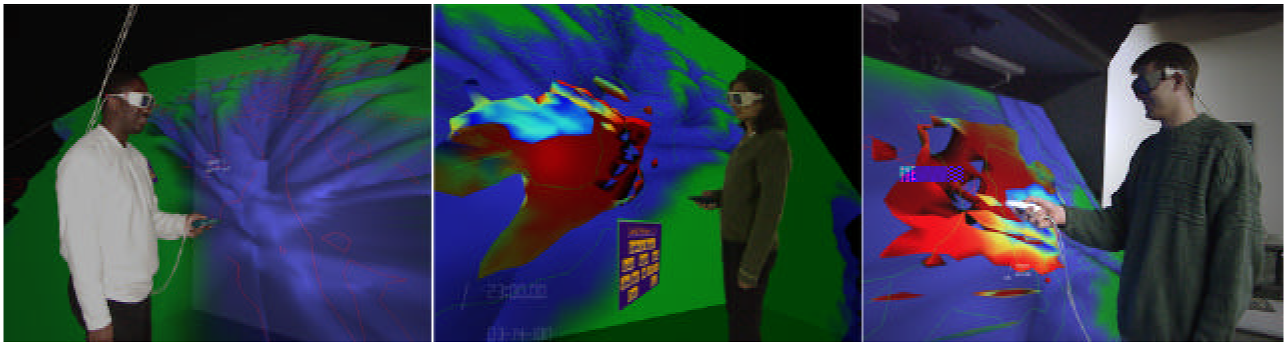

Figure 3. Two participants collaborating with each other in their respective CAVEs (on the left and middle) while the instructor, on an ImmersaDesk (on the right,) observed their activity.

crophones. An assistant was assigned to each CAVE to correct any technical problems that arose. One evaluator recorded the users' answers to a number of queries. The detail trend questions are listed in the appendix. A video camera recorded one workstation's screen and all conversations.

The Procedure

Pairs of students were organized into three treatment groups: those who were allowed to use only *local (private)* views; those who were allowed to use only *global (fully shared)* views; and those who were allowed to use *either* view. The experiments ran for three days with each group experiencing one session per day for approximately one and a half-hour. The first day consisted of a training session to introduce the groups to CAVE6D and Chesapeake Bay geography. The remaining two days consisted of guided search sessions where they were asked to search for specific trends/patterns in the data-set. On the second day, groups were assigned to one of the three treatment conditions. On the third day, they were assigned to the remainder of the two treatment conditions (Table 2). Each session was preceded by a pre-test and succeeded by a debriefing questionnaire.

	Group 1	Group 2	Group 3
Day1: Training	CAVE6D	CAVE6D	CAVE6D
Day 2: Guided Search	GL	G	L
Day 3: Guided Search	L G	L GL	G GL

Table 2. The experimental design protocol showing the task and the treatment condition that is assigned to the groups. G represents globally shared views; L represents private views; GL represents the ability to use either.

Day 1 – The Training Session

Each group experienced a 1-hour collaborative training session in which each person in the pair was placed in their own respective CAVE. The training session was given remotely by the instructor on an ImmersaDesk. Each participants was able to see each other as avatars with long pointing rays emanating from their hands so that they could point at features in the visualization. In addition high quality zero-latency audio mediated their conversations.

The training session consisted of a description of the Chesapeake Bay geography identifying the location of individual landmarks such as Baltimore, Virginia, the Atlantic Ocean, etc. Then the instructor ran through each of the CAVE6D's visualization tools, describing their functions. For example there were tools to display ocean tide vectors and tools to display iso-surface visualizations of salinity.

Finally the instructor familiarized the groups in the use of global and local views. Participants then practiced manipulating the visualization parameters until they felt comfortable with them.

Day 2, 3 – The Guided Search Sessions

The goal of the guided search tasks was to allow participants to focus on the single task of using the tools to verify the existence or absence of a specific trend in the data. The plan was to offer the more difficult, unguided free-form search opportunity on the third day, after they had gained some expertise in wielding the tools.

In the guided search task, the instructor read one trend at a time and asked the participants to use the tools to determine whether the trend existed or not. A sample question would be "How often does the cycle of the tide repeat itself?" Some of the questions were intentionally ambiguous to encourage discourse and hopefully coordination.

When the pair had agreed on an answer for the trend, the evaluator would be called in to record it. The record would note the visualization tools that were activated by each participant and the tools that were considered the most important in revealing the trend. After recording the answer, the instructor would move on to the next question.

Each group was given five trend questions for a session on the second day, and allowed no more than ten minutes for each trend even if they were not able to agree on an answer. Because of time constraints on the third day, we could only deliver three trend questions to each group rather than five questions (Appendix A).

RESULTS

The Training Session

As described earlier, two students in their respective CAVEs were taught by a remote instructor on an ImmersaDesk. They were instructed on basic Chesapeake Bay geography and how to operate CAVE6D's user interface. Although the purpose of the training session was to prepare the students for Days 2 and 3 it revealed significant aspects in CAVE6D's collaborative interface for remote training.

Avatars are useful for the most part

The students and instructor were able to see each other as avatars. The avatar was composed of a head, body and one hand. The names of each user were displayed on the jerseys of each avatar. As users are able to freely navigate around the virtual scene, they see the world from their own perspectives.

Avatars are useful in a tele-immersive environment because it allows the participant to convey one's location relative to a large space and natural gestures such as nodding ones head or waving of one's hand. The orientation of an avatar's head is very useful for determining where the avatar is looking. The pointers at the end of the avatar's hand are useful for pointing out interesting features in the visualization to one's collaborators.

In some cases however when two avatars are attempting to share the same view the presence of the avatars can occlude one's view. Furthermore it is disconcerting to users when avatars interpenetrate one another. Collision detection can be used to help mitigate this problem.

There was a mismatch between coordinate system

Since CAVE6D is a fully immersive application each student is immersed in the data. For example, when they cruise through the trenches in the Chesapeake Bay, their companions will see their avatar perform likewise. That is, that participants and their respective avatars navigate through the data in world coordinates. However the menu interface was presented to each CAVE in local coordinates. Hence, when activated, the menu would float in a constant location in the CAVE. This is the typical method of operation for most menu systems in VR. The problem however is that even though the participants are operating the menu in local coordinates, their avatars are represented to others in world coordinates. Hence the avatars appear to be giving arm gestures that may mislead the remote viewers into thinking that they were gesturing at something in the world coordinate space. The solution might be to stop world coordinate updates of the avatar when menus are being used.

Sharing intent as well as the state of the interface is important

It was observed that simply showing the state of the interface was not sufficient since it did not offer the instructor any feedback on which particular interface item the student was attempting to operate. For example CAVE6D offers a tool for displaying a cross-sectional view of the ocean vectors at any given slice in either the North/South or East/West direction. When a slice is moved, there is no indication of which user was doing the manipulation, or which slice was being manipulated. And since the arm gesture used to move the slice was represented in local coordinates, the avatar appeared as though it was raising its arm for no apparent reason. This problem was worsened by the fact that each avatar's pointing ray could not be toggled on and off hence further amplifying the unintended hand gestures.

Sharing the interface as well as the visualization

During the course of instruction it became clear that it was necessary to be able to share the view of a menu as well as the visualization. This would allow the instructor to confirm whether each student had the correct settings and was viewing the same visualization. Since CAVE6D did not support this, the instructor had to spend a significant amount of time confirming that each student had set his or her visualization parameters correctly.

The Guided Search Session

Overall we found that regardless of treatment condition, collaborators almost always worked independently except to converge on their discoveries. They also tended to work independently when they needed to search over a large body terrain, even if they were allowed to use global-only view. They would partition the terrain so that each would search for a trend over a different part of it. Hence navigation provided multiple perspectives over a single perspective. Hindmarsh et al. [4] describes such views as Fragmented Views, i.e. features of the world are fragmented to different perspectives due to the narrow field of view in desktop VR.

Usage patterns for global, local, either views

When groups used global-only views they ran into frequent problems of both users toggling the same viewing options at the same time hence accidentally deactivating the viewing option. This is a common problem in a strictly shared virtual environment – usually relying on social cues to resolve this conflict. Participants also complained about cases where one person was observing a part of the visualization and the other person would inadvertently change the visualization without forewarning them. In this case, it would be useful if the teleimmersive environment provided some feedback to indicate the actions performed by each collaborator.

In strictly localized views participants engaged in lengthier discussions to ensure that their visualizations were consistent. So while a localized view typically favored independent work, this additional discourse was required for participants to increase the coordination.

They preferred to switch from a global view to a local view when they wanted to test smaller hypotheses without disturbing the overall view, then they would use a global view to present their findings to their partner. These patterns of individual and coordinated activity are consistent with observations in the literature on workflow [11]. The use of localized views was particularly prevalent when users wanted to manipulate the visualizations, such as the depth of the tidal vectors in the bay. The global views were often used when participants agreed on one location to study observation for trends.

When groups were allowed to use either view, they tended to work independently preferring to use localized views. However, during the debriefing session, participants commented that overall they preferred a global view, even though their interaction history seemed to indicate more frequent use of localized views.

Usage of visualization tools

The strategy adopted by all groups was to turn on all tools that seemed relevant to the task and then favor tools that revealed the trend most clearly. The most favored tools may have been different depending on the viewer's point of view and their comfort with using the tool.

Individual differences

Individual differences between participants played a noticeable role in the nature of the collaboration. The more dominant participant tended to take control of the collaboration. In addition, it appeared that more pro-active participants made greater use of localized views to explore alternative solutions.

Learning of the domain material occurred

As a side effect of this study it was found that participants were able to clearly completely articulate each of the trends they found during the debriefing session. This implies that the guided search process may be useful as an instructional technique.

CONCLUSIONS AND FUTURE WORK

CAVE6D was found to be a useful vehicle for performing this initial study of the application of multiple perspectives in teleimmersive scientific visualization. At the same time the results have had immediate practical value by allowing us to offer feedback to the developers of CAVE6D.

Initially we anticipated that there would be little to no use of localized views for guided searches since the areas and parameters of interest were relatively well defined. Conversely we anticipated that localized views would be used more in unguided searches to allow the participants to independently cover as much parameter space as possible. As it turns out localized views were found to be quite useful even in guided searches.

Overall we found that collaborators mostly worked independently, even in globally shared views. They worked cooperatively to divide the collaborative task to two different roles, or to converge on their discoveries. The recurring pattern of activity was 1. problem understanding, 2. agreement on visualization filters to use, 3. independent search for a trend including some adjustments to viewing filters, 4. reporting of discoveries to their partners, 5. negotiating a conclusion based on their combined discoveries. When groups were allowed to use either views, they preferred to use local views to test out small individual hypotheses without disturbing the overall view then used global views to present their findings to their partner.

We were also able to serendipitously re-affirm that some of the findings in previous CSCW (Computer Supported Co-operative Work) works are still applicable in teleimmersion. These include:

- the need for individual pointers to allow collaborators to point at shared data items [6,12];
- that multiple pointers can however become a source of distraction and hence users should have the ability to toggle them on and off [11];
- that some cue of which region of space a user is manipulating is useful [9,10];
- that even in a fully shared environment, participants found the need to work with localized views [10];
- that there is a frequent transition between parallel/independent and coordinated activities [11];
- that in a fully shared WYSIWIS (What You See Is What I See) system frequent usage collisions will occur [11].

We will continue our study of the various aspects of multiple perspectives in the domain of scientific visualization. This initial study has generated many new ideas to explore. These include a comprehensive re-evaluation of prior findings in CSCW in the context of teleimmersion, a comparison between guided versus unguided search for trends, an evaluation with a high-dimensional multivariate data set, and an evaluation with several collaborators in various VR devices, such as CAVEs, ImmersaDesk and fish-tank VRs simultaneously.

ACKNOWLEDGMENTS

We would like to thank all the students for serving as volunteers for this pilot study. In particular we would like to thank Cathy Lascara and Glen Wheless at the Center for Coastal and Physical Oceanography for giving us comprehensive information on CAVE6D and the Chesapeake Bay data set. We would also like to thank Natt Mintrask and James Costigan for providing audio/visual support during data capture.

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible by major funding from the National Science Foundation (NSF), awards EIA-9802090, EIA-9871058, ANI-9980480, ANI-9730202, and ACI-9418068, as well as NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the National Computational Science Alliance. EVL also receives major funding from the US Department of Energy (DOE), awards 99ER25388 and 99ER25405, as well as support from the DOE's Accelerated Strategic Computing Initiative (ASCI) Data and Visualization Corridor program. In addition, EVL receives funding from Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan. The CAVE and ImmersaDesk are registered trademarks of the Board of Trustees of the University of Illinois.

APPENDIX A – THE TREND QUESTIONS

Trend Questions on Day 2

1. Tides at the bay can be diurnal (a cycle of inflow and outflow in 24 hrs), semidiurnal (cycle of 12 hrs) or neither. Which is it?
2. Are there any differences in the speed and direction of the tides at the mouth of the bay, between the north and the south end of the bay mouth?
3. How do the tide velocities vary with depth (if at all)?
4. What are the differences in the salinity levels in the sea and in the bay in general?
5. Does the simulation show any variations of salinity with depth?

Trend Questions on Day 3 First Half Session

1. There is a very noticeable trend in the direction the outflowing tide takes after it comes out of the mouth of the bay, what is it?
2. Do you see a semi diurnal or diurnal tendency of the tides in the sea?
3. There is a trend in the way saline water enters, and fresh water flows out, at the mouth of the bay. What is it and at what part of the simulation is it very conspicuous.

Trend Questions on Day 3 Second Half Session

1. Do you find any differences in the velocities of the tide, between the bay area, and the sea? If yes, what, and how is the variation from the mouth of the bay to the sea, smooth or abrupt?
2. The tide directions change with time both at the sea, and in the bay. How do they differ? Do you see any of the tide directions following a circular motion over time? Where?
3. Where is the velocity of the tides higher, at the navigational channels or at the shoals (shallow regions)

REFERENCES

1. Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R.V. and Hart, J.C. The Cave automatic virtual environment, Communications of the ACM, 35(2), 64-72, June 1992.
2. Gutwin, C. and Greenberg, S. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *Proceedings of the ACM CSCW'98 Conference on Computer Supported Cooperative Work*, pages 207–216. ACM, ACM Press, 1998.
3. Hibbard, W. and Paul, B. Vis5D, http://www.ssec.wisc.edu// 6#6billh/vis5d.html, 1998.
4. Hindmarsh, J., Fraser, M., Heath, C., Benford, S. and Greenhalgh, C. Fragmented Interaction: Establishing mutual orientation in virtual environments. In *Proceedings of the ACM CSCW'98 Conference on Computer Supported Cooperative Work*, pages 217–226. ACM, ACM Press, 1998.
5. Johnson, A. E., Leigh, J., and DeFanti T. Multi-disciplinary Experiences with CAVERNsoft Tele-Immersive Applications. In *Proc. of Fourth Interna-*

tional Conference on Virtual System and Multimedia, pages 498–503. November 1998.

6. Lauwers, J. C. and Lantz, K. A. Collaboration awareness in support of collaboration transparency: Requirements for the next generation of shared window systems, *Proceedings of the ACM CHI'90 Conference on Human Factors in Computing Systems*. ACM, 1990.

7. Leigh, J., Johnson, A.E. and DeFanti, T.A. CAVERN: a distributed architecture for supporting scalable persistence and interoperability in collaborative virtual environments. Journal of Virtual Reality Research, Development and Applications, 2(2):217-237, 1997.

8. Leigh, J., Johnson, A.E., and DeFanti, T.A. Issues in the design of a flexible distributed architecture for supporting persistence and interoperability in collaborative virtual environments. In Proceedings of Supercomputing'97, San Jose, California, Nov 1997. IEEE/ACM.

9. Patterson, J. F., Hill, R. D., Rohall, S. L. and Meeks, W. S. Rendezvous: An architecture for synchronous multi-user applications. In A. Press, editor, *Proceedings of the ACM CSCW'90 Conference on Computer Supported Cooperative Work*, pages 317–328. ACM, 1990.

10. Roseman, M. and Greenberg, S. Building real-time groupware with groupkit, a groupware toolkit. *ACM Transactions on Computer-Human Interaction*, 3(1):66–106, Mar 1996.

11. Stefik, M., Foster, G., Bobrow, D. G., Kahn, K., Lanning, S. and Suchman, L. Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. *Communications of the ACM*, 30(1):32–47, Jan 1987.

12. Tang, J. C. Findings from observational studies of collaborative work. *International J. Man-Machine Studies*, 34(2):143–160, 1991.

13. Wheless, G. H., Lascara, C. M., Valle-Levinson, A., Brutzman, D. P., Sherman, W., Hibbard, W. L. and Paul, B. Virtual chesapeake bay: Interacting with a coupled physical / biological model. *IEEE Computer Graphics and Applications*, 16(4):52–57, July 1996.

14. Wheless, G.H., Lascara, C.M., Cox, D., Patterson, R., Levy, S., Johnson, A., Leigh, J., Kapoor, A. The Use of Collaborative Virtual Environments in the Mine Countermeasures Mission, *Proceedings of SPIE's 13th annual international symposium on aerospace / defense sensing, simulation, and controls*, April 5-9, 1999.