

Collaboration in Tele-Immersive Environments

J. Mortensen, V. Vinayagamoorthy, M. Slater and A. Steed

Department of Computer Science, University College London, London WC1E 6BT, UK

B. Lok and M. C. Whitton

Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill NC 27599-3175, USA

Abstract

This paper describes a study of remote collaboration between people in a shared virtual environment. Seventeen subjects were recruited at University College London, who worked with a confederate at University of North Carolina Chapel Hill. Each pair was required to negotiate the task of handling an object together, and moving a few metres into a building. The DIVE system was used throughout, and the network support was Internet-2. This was an observational study to examine the extent to which such collaboration was possible, to explore the limitations of DIVE within this context, and to examine the relationship between several variables such as co-presence and task performance. The results suggest that although the task is possible under this framework, it could only be achieved by various software tricks within the DIVE framework. A new Virtual Environment system is required that has better knowledge of network performance, and that supports shared object manipulation across a network. The participant-study suggests that co-presence, the sense of being together with another person, was significantly and positively correlated with task performance.

Keywords: Collaborative Virtual Environments, Internet-2, Virtual Reality, Presence.

1. Introduction

This paper describes a user study that investigated the extent to which people in physically remote locations can collaborate together within a shared virtual environment (VE) in order to carry out a joint task. In particular we were interested in what happens when two strangers, physically thousands of miles apart, meet together in a shared VE and have to negotiate and execute an object manipulation task. The pairs of subjects for each trial of the study were located at UCL in the UK, and at UNC-CH in the US.

In the study to be described, two people, one at UCL and the other at UNC-CH, met together in a shared VE and had to lift an object together and move it to another place. At UCL the subject was in a system similar to a CAVETM*8 called a ReaCTor. At UNC-CH the participant used a head-tracked head-mounted display. The subjects

were represented to one another by simple block-like avatars. They could talk to one another. The set up is described in detail below.

A further goal of this research was to examine the extent to which the DIVE system¹ could be used within the context provided by Internet-2 for this purpose. Although DIVE has been developed over several years, and used successfully for remote collaboration, in this work the notion of collaboration was significantly extended. In every almost collaborative VE experiment in which we were involved prior to this, the form of collaboration was limited to people interacting with one another through “seeing” and talking^{2,9}. The fact, that they were in the same virtual place at the same time, implied that they could see the same environment, and talk to one another about its features. Moreover, they could organise their spatial locations within the space in order to optimise carrying out the task together. Very sophisticated behaviour could be achieved within this paradigm to the extent that such a system was successfully used for acting rehearsal, where real actors in remotely located places

* CAVE is a registered trademark of the University of Illinois Board of Trustees.

were able to rehearse a short play, and then later meet for real and carry out the live performance in front of an audience with very little physical face-to-face rehearsal³.

Object manipulation in VR and the effects of tracking lag and frame rates on performance in single-user experiments is described in work by Ware et al.^{11,12}. In this project, however, our goal was to examine whether people could interact more directly, through physically manipulating virtual objects together. In every-day-life we take such actions for granted - such as lifting furniture together. We had never tried this in virtual reality across a network. In an experiment carried out at the MIT Touch Lab⁴ we had explored the extent to which haptic feedback adds to task performance and co-presence when two people in remote places manipulate an object together. Co-presence refers to the extent to which people have a sense of being together, in the same space, rather than their interaction being mediated by a computer interface. In that experiment though, we avoided the question of Internet time delays by running the process on a single computer with two remotely located monitors. With the addition of haptic feedback, both the task performance and their sense of co-presence were significantly enhanced.

This earlier work therefore led to our requirements for the Internet-2 study described in this paper. First, irrespective of haptic feedback, and given the time delays of the network, to what extent was it possible for people to carry out a relatively complex joint manipulation task together? In particular, they would carry a virtual object from one place to another in a large-scale virtual environment (i.e. large virtual area). Hence, this would be manipulating large objects (essentially life sized) together across a distance of several metres.

In Section 2 we describe how the DIVE* system was extended to deal with Internet-2 and the problems encountered in this endeavour. In Section 3 we provide some data on the overall performance of the network. In Section 4, we describe the study, the virtual scenario and the results. The conclusions are presented in Section 5.

2. Extending DIVE for Internet 2

The version of DIVE originally available only supported hardware directly connected to the computer running DIVE, i.e. trackers and head-mounted displays. This was sufficient for the equipment at UCL. At UNC-CH, hardware interfacing with the computers was made possible through using the Virtual Reality Peripheral Network (VRPN) library. At UNC-CH, the peripherals are connected to separate systems than the one running

DIVE. In order to resolve this, a plugin was written for DIVE extending it so that it could interface with these hardware devices using VRPN. This required writing the software plugin itself and providing a C++ interface to the core DIVE C libraries.

Secondly, the subjects at UNC-CH could not navigate the environment merely by physically walking around in their wide-tracked environment due to the large scale of the virtual model and the physical space constraints. Extra buttons on their handheld button device added functionality to allow the subject at UNC-CH to move forwards and backwards using their tracked joystick. The locomotion was implemented using velocity. When the user pressed the forward button an event was sent to the avatar to start moving forwards at a set velocity. Using velocity yielded a realistic (smooth) animation of the avatar. However, this placed a relatively high importance on the individual packets containing these events. During the trials, the user physically walked around for fine grain movement and control while using the buttons to traverse large distances.

When packets containing stop/start locomotion events were lost on the network, this resulted in a considerable loss of synchronization between the local copies of the environment. Since locomotion was velocity based, loss of a locomotion stop packet, for instance, would result in situations where the avatar would be moving in one copy of the environment but be stationary in the other. This was resolved by sending extra copies of the packet when a state transition occurred (start/stop). This does not address the fundamental problems inherent in this event model, which are addressed in another paper⁷.

Thirdly, the object to be manipulated by the subjects was described to them as a 'stretcher'. The implementation of this was based on sharing a single unique object. This involved manipulating the local copy of the object and letting the DIVE system propagate translational and rotational changes to other local copies of that object on the network thereby creating a sense of shared ownership of the entity in question. This mode of manipulation would only guarantee a synchronized environment as long as the changes were applied in the same order in both instantiations of the environment. In turn this would only be possible if the events were generated, sent and processed at a higher resolution than the frequency of the manipulation of the object. If not, there would be disparate states of the object in its various instantiations, each of which would then send updates of its global position resulting in significant jitter of the object. Also, it would continually swap between the local perceived state and the one received in the packets from other instantiations of the VE. Before any given frame is rendered the state of the

* <http://www.sics.se/dive>

object would be determined either by the local or the remote state due to processing of remote packets.

Until pilot experiments were run, the experiment was carried locally and as the LAN provided less than 10ms turn around times the system did not present any problems. As soon as a link up with UNC-CH was carried out, the system faced ~80ms turn around times and caused the stretcher to jitter. This was resolved by implementing an alternative approach that employed distinct local copies of the contents of the stretcher and shared handles, each owned by a single avatar. A local TCL script then updated the distinct local object based on the state of shared global objects. The stretcher would then align locally based on the position/orientation of the handles. In this set-up direct manipulation of a shared object was avoided. The VE appeared synchronized and visually correct, even though the two instantiations would differ slightly due to the lag in updating the positions of the handles. The stretcher would align according to the position of the hand of the subject locally and the position of the rendered hand of the remote avatar. So the alignment of the stretcher was based on the information available locally at the time of rendering.

3. Network Performance

The throughput between UCL and UNC-CH was excellent. In fact the bottleneck was the 10Mbit connection to the machine running the VE at UCL. In the initial phase of this project the round-trip times measured between UCL and UNC-CH were quite unpredictable. This problem was identified as a queuing problem at the NY PoP* where the UK backbone SuperJanet4 peers with the Internet2 backbone Abilene (see Figure 1). When this problem at the NY PoP was resolved, round-trip times to UNC-CH stabilized at around ~80-90ms, which was sufficient for our experiments.




Figure 1: Logical network topology

4. The Study

4.1 Background

A study was carried out in order to assess the extent to which users can collaborate together in carrying out a joint manipulation task. There were no explicit hypotheses for this study. It was simply to observe and record the results, as input to subsequent research.

The task itself had two levels of difficulty. In real life when people carry a large object together (e.g., a bed) around corners or up stairs, it is rarely a simple feat. There are the difficulties associated with weight and with getting the object around sharp bends. Such tasks normally involve a degree of negotiation, often with one of the partners playing a leadership role.

Carrying out a similar task in a VE has several major added difficulties, with only one difficulty removed: that of the physical weight. The first issue was that the virtual depiction of the people (avatars) was simplistic in comparison to reality. Previous studies² have shown how non-verbal feedback such as body language and facial expression are crucial to successfully carrying out virtual tasks together (even tasks that don't involve collaboration). In this study the avatars were block-like structures, with only two moveable parts: head and a pointer indicating the position of the person's tracked hand.

The second predicament was that in real-life physical manipulation, there are natural constraints imposed. When two people are jointly holding an object such as a stretcher with a weight on it, they are constrained to move in

* **PoP** - Point of Presence; a point on a network where connections to the network can be made.

conjunction to each other. If one pushes or pulls, the other will feel it, or if one walks faster the other will be pulled along by this or will have to resist it. In a VE, there are no such constraints. We cannot physically prevent a participant from moving, nor can we enforce any of the constraints imposed by the physical world. We can visually simulate such phenomena but then there would be a contradiction between the visual simulation and what the participants are physically able to do. From previous studies of presence in VEs, it is known that it is necessary to avoid such contradictions. Hence the absence of real weight of the stretcher was in fact a disadvantage.

Thirdly, all the real-life feedback associated with locomotion through a physical space is missing. The feedback from moving your feet along the floor, bumping into objects, edging against a wall, feeling the contact between the physical object being carried and other objects such as edges of walls or doors are entirely lacking. In real life one can easily move in any direction without thinking about it. Within a typical VE the art of navigation becomes a skill associated with manipulating a joystick or other device. Since the carrying of the object also involves the use of such a device, the task is doubly complex - selecting and lifting an object and while holding that object moving around with it.

At UNC-CH, a large scalable wide-area ceiling tracker, the UNC Hiball Tracker, was used, leading to other problems. As the UNC participant physically walks, the virtual model had to be re-scaled in relation to the user, in order to enable them to walk to the boundaries of the model. There is current research for interaction paradigms to overcome this issue, the issue of the relationship between physical size and virtual size, but this research was not sufficiently mature to be employed in this project, and alternatives were put in place as described in Section 3.

Overall then, the task of manipulating an object together in a virtual environment, where the object is large (i.e. body size) is actually far more difficult than carrying out the same task in real life. This was therefore an exceedingly difficult test of the extent to which collaboration between people in remote locations was possible.



Figure 2 - Person in the ReaCTor



Figure 3 - Person in a Virtual Kitchen Displayed in the ReaCTor

4.2 Scenario

A person at UNC-CH Department of Computer Science entered the VE at the same time as a person at UCL Computer Science. At UCL the system used was the ReaCTor system^{*}, a set of four projection walls each with area 3*2.2 meters. The participants held the joystick with four buttons which was tracked, in addition to the head, using the Intersense tracking system. The system is shown in Figures 2 and 3. The UNC-CH system is described in ⁵. In addition, the participant carried a second joystick that was tracked with a second Hiball Tracker.

^{*} <http://www.cs.ucl.ac.uk/research/vr/Projects/Cave>

The VE depicted a large space with a flat ground with a simple building in the middle. On the ground outside the building near where both participants started, there was an object on the ground. This object was described to them as a “stretcher” but only for the sake of giving it a name. In actual case, it was just a flat block with two “handles” attached to either ends with an object on top of the flat-block. The environment and the stretcher are shown in Figures 4 and 5.




Figure 4 - The VE: a building with a blue path




Figure 5 - The object (stretcher)

The person at UNC was not actually an experimental subject, but rather a confederate. At UNC the study was conducted by three people, one of who virtually met each of the 17 subjects recruited at UCL. Altogether there were 17 people for whom results are available (the UNC participants although repeating the study only answered the questionnaire the first time).

The subjects at UCL were recruited by email advertisement from 2nd and 3rd year undergraduate Computer Science students, and MSc students from the Bartlett School of Architecture. There were 5 women and 12 men, ranging in age from 19 through 34.

On arrival for the study the UCL subjects were asked to sign a consent form, which informed them of possible

negative effects from using the system such as simulator sickness. They were told that they could withdraw from the study at any time without giving a reason, and that they agreed that they would not be driving or operating complex machinery for at least 3 hours after the conclusion of the study.

They were then given a sheet outlining the task. This sheet informed them that the task was to meet with the other person at UNC, and negotiate lifting a stretcher together. They were to take the stretcher together along a blue path that led into the building, and then put it down on a red coloured area inside the building.

There were two experimenters at UCL. One took the subject through a training task. This was to make sure that each of the subjects was familiar with how to navigate the VE, how to pick things up and hold them, and then carry them as they navigate. This was done in a different VE to the main experimental task. When the training was completed, the second experimenter started up the link to UNC, and then explained the task again.

The subjects were told to negotiate with the other person about lifting of the stretcher and taking it along a blue path. It was emphasised that the blue path was only a guide and that more important than following it, was to follow the direction indicated. Once all this was explained, the UNC-CH avatar entered the VE and the study started. The collaboration was stopped after 5 to 8 minutes, depending on the stage to which the process had reached. Each person was represented to the other using an avatar.



Figure 6 - The DIVE Avatar

Subjects were free to move around the ReaCTor space subject to the physical constraint of the walls. It was noticed in earlier pilots that subjects tend to use their bodies involuntarily – e.g. shaking or nodding heads while

talking. Although the avatars were very basic, simple movements can be viewed: e.g. when the UNC person nodded his head, bent down or held his head low, the UCL subject would see this in real-time and vice versa. We therefore wondered if in spite of such a basic avatar, to what extent could “mood” be conveyed across such a narrow bandwidth. Hence, we instructed the UNC person to behave as either “very happy” or “very depressed”. As part of the data we gathered we asked the UCL subject the extent to which they recognised the mood of the other person. Of course such mood guessing would be based on voice, but also on the disposition of the avatar body: drooping head indicating depression.

4.3 Questionnaire

A questionnaire was designed to assess the behaviour and views of the subjects. This questionnaire was administered online* to each of the UCL subjects straight after their experience. It was also given to two of the UNC helpers after their first experience with a UCL subject. The questionnaire obtained responses on each of the following:

- *Demographic information*: such as age, gender, status etc.
- *Task Performance*: assessments of self and other’s performance, the degree of harmony and cooperation between the participants.
- *Co-presence*: the sense of being together rather than interacting through a computer interface.
- *Similarity to real life*: the extent to which the experience was similar to moving an object together with someone in real life.
- *Mood assessment*: Assessment of the mood of the other person.

Each of these (apart from the demographic data) was assessed on a 1 to 7 scale, as shown in the Questionnaire (can be found at *). There was also an open-ended question where subjects could write their answer:

Please enter your comments. Things you could consider are:

Things that hindered you or the other person from carrying out the task; what you think of the person you worked with; and any other comments about the experience and your sense of *being there* with another person. What things made you “pull out” and more aware of the computer...

The purpose of this was to try to get behind the purely quantitative results to what the subjects were thinking about their experience.

4.4 Quantitative Results

These results are only for the 17 UCL subjects, since the 3 UNC helpers were confederates in the experimental design.

Co-presence was assessed from the following questions:

- To what extent, if at all, did you have a sense of being with the other person?
- To what extent were there times, if at all, during which the computer interface seemed to vanish, and you were directly working with the other person?
- When you think back about your experience, do you remember this as more like just interacting with a computer or working with another person?
- To what extent did you forget about the other person, and concentrate only on doing the task as if you were the only one involved?
- During the time of the experience, did you think to yourself that you were just manipulating some screen images with a mouse-like device, or did you have a sense of being with another person?
- Overall rate the degree to which you had a sense that there was another human being interacting with you, rather than just machine?

Each was measured on a 1-7 scale, and in the analysis the directions adjusted so that 7 always means the highest co-presence, and 1 the lowest. The overall mean co-presence was 3.8 ± 1.1 .

In order to examine the association between other variables and co-presence, we score each subject by the number of ‘high’ scores on the 6 individual questions. A ‘high’ score in answer to a question is one that is above 4 out of 7.

Hence the overall score is actually a count of the number of high scores out of the 6 questions - e.g., if the result is 4 then it means that in 4 out of the 6 questions the subject responded with a score that was 5, 6 or 7. This is a conservative way to treat the results, and has been used several times before, for example ⁶.

We do this in order to carry out a logistic regression analysis between the co-presence results and the other variables. Here we report only results that are significant at the 5% level.

In the following results the means and standard deviations for the explanatory variables are shown in brackets after the question statement. Co-presence is most significantly

*<http://www.cs.ucl.ac.uk/research/vr/Projects/Internet2/Report/>

and positively associated with two other variables in one single fitted model:

- Please give your assessment as to how well you contributed to the successful performance of the task. (4.2 ± 1.3).
- Please give your assessment as to how well the other person contributed to the successful performance of the task. (3.9 ± 1.1).

In each case the higher the score the higher the co-presence score (and these two explanatory variables are uncorrelated). In other words this demonstrates a link between co-presence and the subjectively assessed level of task performance. The coefficients for the two variables are almost equal, indicating that self- and other-performance were equally weighted.

Other variables that are *positively* associated with co-presence, but individually, not within the same overall model:

- To what extent were you and the other person in harmony during the course of the experience? (3.9 ± 1.7).
- Think about a previous time when you co-operatively worked together with another person in order to move or manipulate some real thing in the world (for example: shifting some boxes, lifting luggage, moving furniture and so on). To what extent was your experience in working with the other person on this task today like the real experience, with regard to your sense of doing something together? (2.9 ± 1.3).
- Please give your assessment of how well you and the other person together performed the task (4.3 ± 1.4).

The second question is quite important, since it gives an overall view of how “*real*” the collaboration felt. The mean response (2.9 ± 1.3), indicating overall that the degree of similarity to moving an object in real life was relatively low.

The following was *negatively* associated with co-presence:

- To what extent, if at all, did you hinder the other person from carrying out the task? (3.3 ± 1.8).

In other words the more the subject believed that they had hindered in carrying out the joint task the lower the sense of co-presence. Interestingly co-presence was not correlated with the degree to which the subject felt that the other person had hindered the carrying out of the task.

Almost all of the variation in co-presence can be explained by just two variables taken together in one model:

- Please give your assessment as to how well you contributed to the successful performance of the task. (4.2 ± 1.3).
- If you had a chance, would you like to meet the other person? (4.6 ± 1.7).

The higher the self-assessed contribution of the subject, and the more she or he wishes to actually meet the other person, the greater the degree of co-presence.

UCL subjects were able to assess the mood state of the UNC person (recall that the UNC person was acting as happy or as depressed). The mean mood score (i.e., estimated degree of happiness) amongst those who experienced the depressed acting was: 2.5 ± 0.8 , and for those who experienced the happy state: 3.9 ± 1.4 . The difference is significant at the 5% level. The relatively high variance amongst the “happy” responses probably reflects that it seemed to be much harder for the UNC subjects to convincingly act in a “happy” manner than in a depressed manner.

4.5 Written Results

The full list of written results can be found on the accompanying web pages*. The responses to the question about what hindered the subjects in performing the tasks can be classified into a number of categories:

- Problems in system behaviour
- Limitations in capabilities
- Breakdown in communications
- Lack of realism
- Problems with network performance

The following are some quotations that illustrate several of these points. (Problems with network performance were the least often mentioned).

I thought the VE wasn't very realistic. The stretcher did not look like a stretcher and my partner looked like a bizarre robot thing.

My sense of presence was not very high at all. Probably because there is no tactile feedback and the fact the graphics were chunky and basic. The sense of immersion was high.

I wasn't sure whether 'laws of physics' applied, i.e.: whether walking 'through' the stretcher would prevent the other person from completing the task.

*<http://www.cs.ucl.ac.uk/research/vr/Projects/Internet2/Report/>

I wasn't sure whether the other person was always listening. I was expecting to be told what to do by the other person, but he was not forthcoming.

The quality of the sound was distracting - leading to the 'not listening' problem. I didn't always know if the other person had heard what I said.

There was a slight transmission delay that dulled the realistic sense of the experience and the "stretcher" which did not particularly look all that realistic, which made it hard to know what to manipulate in order to get the task done.

Not being able to walk sideways meant that it was a constant struggle to continue to hold something and move in a sideways direction.

The person I was working with didn't realise if I was in difficulty - when I asked to put down the object - in order to reassess the situation, an "I'm okay" answer was received indicating a communication problem.

I felt the person was real but because of the physics of the task of moving the stretcher I got confused as visually there was no fixed distance we could be apart when we were holding the stretcher.

The weight of the glasses reminded me all the time that it was just with a computer.

Not being able to effectively walk using the controls. When carrying the 'stretcher' I had the front of it. As it was hard to perceive what was going on I tried to walk backwards with it (which I gave up on eventually, as going forwards worked) and this made it VERY hard to walk round the corners. A strafe left/right function like in Quake would be very useful!

In the beginning it felt more like a that I was in a virtual world but once we started the task and communication between us harmonized then it felt as we were doing the task in a real world. I think in a situation like this one the only thing that makes u realize that you are not in a real world is the lack of communication

5. Conclusions

An overall conclusion from this research is that the Internet-2 system will definitely support the kind of immersive interaction between people as described in the types of experiments carried out here. However, this is a highly qualified statement. What is really meant is that the transmission speed is satisfactory, and our preliminary evidence suggests that this is the case for visual and auditory communications between people. Of course the scope should be taken into account. This involves just two people interacting in a very limited task. But the evidence is promising. The real problem, the major difficulty with this work is that there is no adequate software support!

DIVE is simple to use and easy to set up, provided that all that the participants do is look at and talk to one another and do not interact synchronously on the same object. As it currently stands it can barely support the more complex interactions involved in people collaborating in a more sophisticated way, such as joint manipulation of objects. The reasons are clear: each packet is highly significant. There is basically no way that the system can recover this information in time for there not to be a break in consistency between the people involved. This even happened with navigation: when the remote collaborator released the button indicating that they had stopped moving forward, and that button press event was lost, as far as the local person was concerned the other person was zooming off to infinity, but as far as the remote person was concerned they were not doing so and were in the right place. These problems must be addressed at a fundamental network and database level and moreover an event subsystem that is decoupled from the rest of the system subparts, such as rendering, interaction, audio, must be developed.

We have used DIVE for about six years in our collaborative VE research. All this time we have found this to be an adequate system framework within which we could pursue our technical and experimental work. The reason why this has been possible is that all of our collaborative work has concerned social interaction where total consistency is not strictly required. In simultaneous object manipulation where the consistency requirements are higher we experience the problems described above. Overall the project clearly indicates the need for VE systems to improve their understanding of the network.

There are potential solutions within DIVE to all of these problems, and we will be pursuing these in future research, as well as looking for solutions beyond DIVE.

We have tried types of interaction between people that we have never carried out before. The data suggests that in order to have a sense of being with another person, it is vital that the system 'works' in the sense that people have an impression of being able to actually do what they wish to do. To enable this, the system has to be secure in transmitting vital data, which does not aggregate disparate packets and maintains consistency. Significant work is required for a robust solution, rather than a fix for a certain problem at a certain time.

A general point that this research makes is that collaboration in VEs is really hard to do when it involves manipulation of shared objects. It is hard at many levels - interface issues, tracking delays, network delays, communication (aural and visual), synchronisation and lack of important cues such as haptics. We believe that the

type of collaboration under the conditions described in this paper remains unsolved in VR.

Acknowledgements

This research is supported by UKERNA. We would like to thank David Chapman of Geomatics at UCL and João Oliveira of UCL Computer Science for helpful input to the project. We extend our gratitude to Manuel Oliveira of UCL Computer Science for comments on project reports, and to Sharif Razzaque of UNC-CH, a visitor at UCL. Thanks to David Swapp, the manager of UCL's ReaCTor system for help throughout the project. Finally, we would like to acknowledge the volunteer work put in by Paul Zimmons and Samir Naik during the running of the trials.

References

- [1] Steed, A., Mortensen, J., and Frecon E., "Spelunking: Experiences using the DIVE System on CAVE-like Platforms", Proceedings of the Joint Fifth Immersive Projection Technology Workshop & Seventh Eurographics Workshop on Virtual Environments, Stuttgart, May 2001.
- [2] Slater, M., Sadagic, A., Usoh, M. and Schroeder, R., "Small Group Behaviour in a Virtual and Real Environment: A Comparative Study", Presence: Teleoperators and Virtual Environments 9(1), 37-51, 2000.
- [3] Slater, M., Howell, J., Steed, A., Pertaub, D-P., Garau, M. and Springel, S., "Acting in Virtual Reality", ACM Collaborative Virtual Environments, CVE'2000, 103-110, 2000.
- [4] Basdogan, C., Ho, C., Srinivasan, M. A. and Slater, M., "An Experimental Study on the Role of Touch in Shared Virtual Environments", ACM Transactions on Computer Human Interaction, 7(4), 443-460, 2000.
- [5] Welch, G., Bishop, G., Vicci, L., Brumback, S. and Keller, K., "The HiBall Tracker: High-Performance Wide-Area Tracking for Virtual and Augmented Environments", VRST'99, Proceedings of the ACM Symposium on Virtual Reality Software and Technology, 1999.
- [6] Slater, M. and Steed, A., "A Virtual Presence Counter", Presence: Teleoperators and Virtual Environments 9(5), 413-434, 2000.
- [7] Mortensen J., Oliveira M., Slater M., and Steed A. "Virtual Environment Systems are Non-Network Friendly: A Case Study", (submitted 2002).
- [8] Cruz-Neira, C., Sandin, D.J. and DeFanti, T.A., "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE", Computer Graphics (SIGGRAPH) Proceedings, Annual Conference Series, 135-142, 1993.
- [9] Steed, A., Slater, M., Sadagic, A., Tromp, J. and Bullock, A., "Leadership and collaboration in virtual environments", IEEE Virtual Reality, Houston, March 1999, 112-115, IEEE Computer Society, ISBN 0-7695-0093-5
- [10] Floyd, S., Jacobson, V., McCanne, S., Liu, C-G. and Zhang, L., "A Reliable Multicast Framework for Lightweight Sessions and Application Level Framing", IEEE/ACM Transactions on Networking, November 13, 1995.
- [11] MacKenzie, I.S. and Ware, C., "Lag as a Determinant of Human Performance in Interactive Systems", INTERCHI'93 Technical Paper, 488-493, 1993.
- [12] Ware, C and Balakrishnan, R., "Object Acquisition in VR displays: Lag and Frame Rate", ACM Transactions on Computer Human Interaction, 1(4), 331-357, 1994.