

Interaction between Users of Immersion Projection Technology Systems

A. Steed¹, D. Roberts², R. Schroeder³, Ilona Heldal⁴

¹Department of Computer Science, University College London, Gower St, London, WC1E 6BT, UK

²Centre for Virtual Environments, Business House, University of Salford, Salford, Manchester, M5 4WT, UK

³Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford, OX1 3JS, UK

⁴Department of Technology and Society, Chalmers University, 412 96 Gothenburg

A.Steed@cs.ucl.ac.uk, D.J.Roberts@salford.ac.uk, Ralph.Schroeder@oii.ox.ac.uk, ilohel@mot.chalmers.se

Abstract

Participants using immersive projection technology (IPT) systems tend to behave as if they were really in the space depicted by the displays. This behavioral presence leads them to adopt responses and strategies to events and situations that are similar to what they would adopt in analogous real-world situations. In this paper we discuss what happens when IPT systems are used for tele-collaboration, and each participant sees a virtual environment (VE) that includes representation of another (an avatar). In these situations we see the participants respond to that avatar as if it were the other person.

We analyze several recent experiments using IPTs where participants represented to each other using avatars complete tasks together. On one hand, because these are IPT systems, a participant can not see their own avatar, and doesn't know what, if any, of their appearance and behavior is conveyed. On the other hand, they can see the appearance and behaviors of other participants' avatars. We explore participants' awareness of each other through their avatars, and we demonstrate how participants successfully and unsuccessfully use their own avatar to communicate. We then analyze some of the more complex collaborative situations that are being supported in current applications and show what requirements there are for future research on avatar representations and interaction.

1 Introduction

Participants using immersive virtual environments (IVEs) sometimes behave as if the space depicted by the displays were real. This 'behavioural presence' leads them to adopt responses to events and situations that are similar to what they would adopt in analogous real-world situations (Slater & Wilbur, 1997). There is no doubt that behavioural presence can be seen in certain situations; one of the most compelling is stress response to heights (Meehan, Insko, Whitton & Brooks Jr, 2002). However does this extend to other, more subtle forms of behaviour? Specifically, how might they respond to representations of other participants in the environment?

In this paper we examine evidence of behavioural presence when participants are tele-collaborating between CAVETM-like (Cruz-Neira et al., 1993) immersive projection technology (IPT) systems. In these situations each participant sees a world that includes representation of another participant. We might expect the participant to respond to the avatar as if it were the other person, respecting well-known proxemic rules. However, we find that behavioural presence response is mixed: sometimes we see behaviours that can only be explained as a behaviour response following a social rule; at other times we see such rules being ignored.

We analyze several recent and ongoing experiments involving remote collaboration using immersive projection technology (IPT) systems. The tasks we have used include conversational tasks (Garau et al., 2003), focused object manipulation tasks (Axelsson et al., 2001, Roberts, Wolff, Otto & Steed, 2003), and more open-ended trials over longer periods of time (Steed, Spante, Schroeder, Heldal & Axelsson, 2003).

The main conclusion will be that despite many real-world behaviours being supported through the IPT system interface, we see varying levels of the participants' using these behaviours. For example, it is natural to face someone whilst speaking to them and although we do observe this, it is not universal. We will analyse this and similar situations to show where system design limitations, system failures or simple expediency have meant that participants have adapted their behaviour. Despite the observed adaptations of behaviours we do see some remarkably complex interactions emerging and being supported.

2 Related Work

So far, most studies of collaboration in virtual environments have dealt with desktop systems (Churchill, Snowdon and Munro 2002). Further, the focus has typically been on the way in which the individual interacts with system in order to collaborate rather than on the collaboration itself. This overlooks the complex interplay between the interactions between the avatars inside the virtual environment, though some recent work has examined how avatars interact with each other in terms of the social dynamic (Schroeder, 2002, Schroeder & Axelsson, 2005).

Hindmarsh et al. (2000) showed that remote collaboration on desktop systems has severe limitations due to the limited field of view and difficulties in referencing parts of the world. The study also shows that participants have problems in being able to take their partner's point view inside the environment. In immersive systems, many of these problems can be overcome because of the better capture of participant behaviour through tracking and the wide field of view of the displays. This means that participants are much more peripherally aware of their collaborator. Peripheral awareness supports communication about the task at hand but it also supports the maintenance of the collaboration itself since the participant rarely loses track of their collaborator.

A few recent studies have started to investigate how collaboration is affected by the use of various combinations of display system. A number of studies have shown that immersed participants naturally adopt dominant roles versus desktop system participants (Slater, Sadagic, Usoh, and Schroeder, 2000, Heldal et al 2005). Recent studies by Schroeder et al. (2001) and Roberts et al. (2003) have investigated the effect of display type on collaboration of a distributed team. Schroeder et al. showed that doing a spatial task together, in this case a Rubik's cube type spatial puzzle, can be practically as good as doing the same task face-to-face, whereas the same task takes considerably longer on desktop systems. Roberts et al. have shown that it is possible to successfully do a construction task (building a Gazebo) in networked IPT systems, a task which requires that partners work closely together and in a highly interdependent way.

From these studies, it has emerged that different types of networked VE systems support collaboration and social interaction in different ways. In this paper, we will examine more specifically the interplay between object-focused interactions and interpersonal interactions. Sometimes participants behave very much like they would in the real world, and at other times real world rules seem not to be applied at all.

3 Implementation

3.1 Tasks

Seven tasks are discussed in the examples in Section 4. Each is discussed in more detail in related papers, and in this paper we will draw out common types examples from these trials. Figure 1 gives a snapshot from each task. The seven tasks are:

Gazebo. Pairs of participants build a gazebo. The task involved close collaboration to fit planks and supports together. Certain tasks require both participants such as carrying objects and fixing parts together.

Rubik's. The task was to do a small-scale version of the popular Rubik's cube puzzle. Eight blocks with different colours on each side had to be assembled so that each side would have a single colour.

Landscape. The world in this case was a small townscape with surrounding countryside ringed by mountains. Subjects were instructed to familiarize themselves with this landscape and count the number of buildings. They were also told that they would be asked to draw a map of the environment at the end of the task

Whodo. This task was based on a popular game, in this case the 'who-dunnit' board game, Cluedo. The subjects were asked to find five murder weapons and five suspects in a building with nine rooms. They needed to locate the murder victim's body and find and eliminate weapons and suspects.

Poster. This environment consisted of a room with ten posters stuck on the walls. The posters each contained a list of six sentence fragments. When all the fragments were put in the right order, they would make a popular saying.

Modelling. This environment contained 96 shapes (square blocks, cones, etc.) in six different colours. The subjects were told to make a building or model of a building.

Conversation. In this environment two participants simply meet and have a conversation. Each is given a role to play and the focus of the task is on improvising to fit within the conversation.

These tasks were used in two different types of study: focussed case-based experiments and longer-term ethnographic analyses. In this paper we will use fragments of conversation captured from video, audiotapes and observers' notes from both types of study to frame our discussion. In other papers more detail may be found about the tasks, the data captured and the analyses done. The Gazebo trial is discussed in (Roberts et al., 2003). The Rubik, Landscape, Whodo, Poster and Modelling worlds are discussed collectively in (Steed et al. 2003). In that paper experiment participants experienced all five worlds over the course of a day. In the rest of this paper this experiment is referred as the "long-term" experiment. In addition the Rubik trials is discussed in (Axelsson et al. 2001, Heldal et al., 2005, Schroeder et al., 2001). The Conversation task is discussed in (Garau et al., 2003).

3.2 Hardware and Software

All of the tasks we will present were implemented using the DIVE (Distributed Interactive Virtual Environment) systems. DIVE is an Internet-based multi-user virtual reality system in which participants can navigate in a shared 3D space and interact with each other. DIVE was originally designed for desktop environments, but has been extended to support IVEs such as HMDs and IPT systems (Steed, Mortensen & Frécon, 2001). The immersive version of DIVE hides all the complexity of the immersion support and allows an immersive client to interact with a collection of other desktop and immersive clients connected over the Internet.

DIVE is a peer-to-peer system and clients that join a multi-participant session fetch the current state from one of the peer group. Some of the sessions have been run over local area networks, but others have been over wide area networks. Sites involved include: University College London, UK; Chalmers University of Technology, Sweden; University of Reading, UK and University of Salford, UK. In order to support multicast between these sites, we have used the DIVEBONE facility (Frécon, Greenhalgh & Stenius 1998). This is an application-level multicast bridge that allows us to connect together local area networks so that they can participant in shared DIVE sessions.

Several varieties of immersive technology were employed. The trials under consideration always used at least one IPT system. The Conversation task used an IPT and a HMD. The Gazebo task was occasionally run with three IPTs. The IPTs were all four-sided (three walls and a floor) with the exception of the Chalmers TAN VR-Cube which was five-sided. All used some type of SGI Onyx computer to drive the displays. A variety of tracking technologies were used, though for all systems the participant's head and one of their hands were tracked. The HMD used in the Conversation task was a Virtual Research V8. In all trials, the participants could talk to each other. This was usually achieved using the Robust Audio Tool (RAT). Participants were represented to each other using avatars. These avatars had freely moving hands and head. The body would rotate in azimuth only under the head, and the right or left arm would bend to fit between the tracked hand and the relevant shoulder. The participant would not see their own avatar.




Figure 1: Snapshots of the worlds used in the seven applications discussed in this paper. From top to bottom and left to right: Gazebo, Rubik's, Landscape, Whodo, Poster, Modelling and Conversation.

4 Interaction Examples

We present commentary on examples of two participants interacting through the IVE. The main purpose is to draw out situations in which participants followed or didn't follow expected social and environmental rules.

4.1 Referencing Objects

As mentioned in Section 2, Hindmarsh et al. (2000) have shown that it is difficult to reference objects when working with desktop systems. In the Hindmarsh et al. study, participants given the task to relocate furniture in a virtual room spent more time identifying objects than deciding where to put them. The main reason for this was the limited field of view. For a desktop system this problem can be reduced, but by no means eliminated, by placing the viewpoint behind the avatar. To a large extent this problem does not happen on IPT systems because of the wider field of view.

In the Gazebo task, the environment is more cluttered than the room of furniture in the Hindmarsh study. Yet in collaborating on the Gazebo task, participants routinely and easily referred to the objects and to the place in the environment where they were located. The IPT tracking systems and use of avatars means that pointing gestures are conveyed accurately. Participants could locate themselves in relation to the objects, in a wide swath of directions limited only by the number of walls in the IPT. The wide field of view, direct control of gaze and accurate representation of pointing supported efficient referencing of objects. A typical conversation between two collaborators - Bob and Lara in this case - engaged with moving a heavy beam using carry tools, can illustrate this:

Bob: *Hey, let's move this beam over there.* [Points with his hand to the beam in front of him and to the left]

Lara: [Rotating to see Bob and then follows his hand movements]

Lara: *Ok, I will take this carry tool here.* [Points to the tool and moves to pick it up]

Bob: *I'll take the other tool then. Where is that?*

Lara: *Just next to you.*

Bob: *Ah, ok.* [Rotates and picks up the tool]

Lara: *I have my end of the beam now.*

Bob: *Yup, I am right with you. Ok now let's move it over there.*

Lara: *I am following your direction.*

Verbal and non-verbal communication can be used in a nuanced way while referring to objects and places within the environment. For example, a participant might point to an object and say "let's pick that up" and then rapidly turn and point to a place in the environment saying "and take it over there". Another example is a participant simply taking an object and telling the other participant to do the same. These gestures happen very fluidly and there is an implicit assumption on the part of the participant performing the gesture that their collaborator will see and hear the gesture immediately.

Similar behaviour can be seen frequently in the Rubik's task which is also highly object-focussed and where success in task depends on the activity of the other. Indeed we often see participants making quite complex gestures towards objects such as pointing to parts of objects they cannot see because they are occluded from their current position, and gesturing towards objects that aren't directly in front of them. These both require the participant to have good awareness of their own position and posture as well as faith that the other person will also see this position and posture.

4.2 Handling Objects and the Virtual Space

Participants treat objects in the space in naturalistic and un-naturalistic ways. For example, in the Whodo task one participant at one point walked straight through several pieces of furniture, yet only seconds later bent down physically to look under a chair in order to see if something was hidden there. Or again, in the Poster task, one person stood in the middle of a shelf and walked through a wall without showing awareness of this, and yet shortly thereafter commented verbally on a similar situation:




Figure 2: Example of a closely coupled interaction involving a hand over. The five frames cover 13 seconds of interaction. The local participant places a block on his right, then the remote participant hands him another block which he stacks on top.

Person 1: *Are you outside [the room with the Posters] as well?*

Person 2: *No, I'm inside.*

A further example can be taken from the Rubik's cube puzzle. In the Rubik's cube the objects stayed suspended in the air without simulated gravity. Still, participants working closely together in handling the cubes easily coped with this in handing each other the cubes, even if they were initially surprised that they could hand each other objects:

Person 1: *What is happening if I try to take this from you?*

Person 2: *I don't know. Oops. [surprised as she hands over a cube and the other person takes it]...do you have a red one [one with the red side showing which has just been handed over] at your place?*

Figure 2 shows a sequence where the two participants collaborate very closely to start building the cube in the Rubik's task.

4.3 Interactions between Avatar Embodiments

When participants encountered each other briefly and only for a single task, as in the Gazebo or Conversation tasks, they did not typically comment or inquire about their partner's avatar or circumstances. On the other hand, when participants collaborated for a longer period, as in the long-term experiment, they would typically comment on each other's avatar appearance and ask about their own appearance at the outset of the session. They would also inquire about where the other was physically located, about their occupation, about the weather and the like. They would also wave goodbye to each other or try to shake each other's hand upon successful completion of a task.

In the experiment involving the Conversation task one participant made the following comment about the avatar in a post-experience interview:

"Even if it is not a very realistic avatar, it helps a little. It gives you something to focus on. Although you do not think of it as a person, strangely it does stop you turning away or doing anything inappropriate. Also your mind does not wander as much as it might on the telephone. "



Figure 3: Examples of close, face-to-face interaction. Left: Two participants meeting in the Gazebo application. Right Two participants discussing the task in the Poster application.

There is a difference between how participants treat each other's avatar depending on the stage of the task and their focus: at the beginning and end of each of the five tasks, avatars would generally stand facing each other and treating each other as in conversation with real people. Figure 3 shows two examples of this from different tasks. In the Conversation task, without out exception the two participants would stand directly facing the other's avatar for the complete conversation. In longer trials or in trials involving object interaction there was a definite middle phase when they would only react to collisions with or proximity to partner's avatar when necessary, and otherwise would seem not to consider it important. For example, in the course of one session, one person stood inside another's avatar for several seconds. On other occasions, a person would go through the other's avatar several times within a short time. Avatar embodiments would thus be noticed if they were related to the task, for example when one person was blocking the other person's view in the Poster task:

Person 1: *Move on! You are in my way.*

Person 2: *OK.*

In the Gazebo task, or during the long-term trial, the users would frequently glimpse at each other to ascertain each other's position or to give an indication to their partners that they were aware of what they were doing. They also communicated verbally to reassure each other or to keep the conversation or the flow of the activity going. Short phrases such as 'all well?' or 'this is fine' or 'I am over here' played an important role, even if they were not always explicitly acknowledged.

4.4 Working Separately and Together

For parts of some of the tasks the participants would need to temporarily split up to undertake separate but related jobs, such as gathering materials in the Gazebo or Modelling tasks. In this phase, they could still keep an eye on each other in case help was needed. During such periods, communication was generally more limited as they concentrated on their own task. At such times, small talk replaces work related verbal communication, and this seems to maintain the social feeling and increases the feeling of co-presence.

Changes in activity or level of communication of the other person can signal a need to bring collaboration closer together again and communicate directly. In the Gazebo and Cubes tasks, we often observed participants looking over to check how their partner was getting on and offering assistance when necessary by, for example, fetching a tool or object for the other person to use. Because of the wide field of view of the IPTs, the other participant was usually in view. Furthermore because of the tracking small but significant changes in behaviour were easily noticed.

It was quite complex to collaboratively manipulating an object in the Gazebo task. Jointly carrying an object means that they must negotiate who will take the lead and then coordinate where they are going. It is also important that the participants can see which end of the object is to be picked up. This is well supported in the IPT because the avatars and their orientation often mean that there is an obvious solution to the problem where each participant takes the end

that is closest to them. It is also easy to infer which object, or part of an object the other participant is talking about through their body language.

A good illustration of this is in a particularly complex phase of building the gazebo where the task requires that different objects are connected by jointly holding them together while drilling a hole through both and inserting and tightening a bolt. Simulated gravity requires that one participant must hold the object in place while the other fixes it, see the left frame of Figure 4. This means that verbal communication may be required to move the task along. For example, one participant will position an object while the other is told to gather the necessary tools to fix it. Or there can be verbal instructions about how to fix the object or about where a tool can be found (“look behind you”).

A complex set of gestures can be found in this collaborative assembly. For example, nodding with the head towards the partner indicate agreement, whilst nodding towards an object indicates that object or a direction. Participants also frequently walk around the object of interest. This is accomplished much more easily with an IPT system than with a desktop system it can be achieved by physically walking around inside the IPT.

4.5 Planning and Instruction

For most of the tasks some form of planning or instruction is necessary to negotiate or convey a strategy. For example, in the Gazebo task, one participant would demonstrate how to undertake a given operation, such as using a tool to fix two construction objects together. In the Cube, Whodo, Poster, Landscape and Modelling tasks, initial discussion is necessary to understand the task, form a strategy and partition the work. Verbal communication is essential at this stage to describe the upcoming tasks and to agree on locations and coming steps, but other cues, such as gestures, are also widely used to point out directions and underline the verbal communication.

In the planning phase, observations show that participants mainly stand around or use body-centric gestures such as facing each other while they discuss their further action, see for example Figure 2. The use of IPTs supports this kind of behaviour by allowing the participant to turn and move naturally within the spatial context. Planning can also involve the use of objects. For example, when explaining how or where to use a tool, it is often easier to take the object and to demonstrate it rather than give verbal instructions.

The planning stage of the Gazebo task includes a complex combination of speech and gestures. This includes nodding and shaking of the head and pointing, not only with the hand but also the head and body. A participant may, for example, face his or her partner, point with the hand to the object and move the head between the two. Once the object has been lifted the participants keep each other in sight, either by walking side by side or one person in front walking backwards. Once the hand is used to carry an object, however, it can no longer be used to gesture and the body and head may be employed to do this.

Another feature of the planning and instruction process is that participants divide the roles between themselves. For example, collaborating partners will adopt the role of instructor and student. In the middle frame of Figure 4 the avatar on the left, the instructor, is explaining and the avatar on the right is testing a tool. The traces in the figure show that the instructor is more animated.

4.6 Interaction with the Technology

Occasionally it is obvious that the technology itself is interfering with the experience. One example is the use of the non-tracked hand. This phenomenon occurred frequently and still occurred when participants had had plenty of experience with the system. Mostly this went unnoticed, though on occasion the mistake was explicitly noted, as in the following exchange during the Rubik’s cube task:

Person 1: *Here, here...against me...Do you see my...Of course, you can’t see my hand. Damn, I’m waving with my real hand, hmmm...*

Person 2: *Sure, yes...If you point with your joystick hand I’ll see it.*

Another example is interacting with the physical walls of the IPT. Participants coped with this in different ways: sometimes they bumped into the walls and commented on this, at other times they worked quite close to the physical walls while avoiding touching them in a skilful way. The participant in the right frame of Figure 4 was using her non-tracked hand to avoid colliding with the wall while at the same time using her tracked hand to continue working with the objects. In other words, participants were able to cope with treating the physical space and the virtual objects as such simultaneously.



Figure 4: Left: Two participants fix a bracket to an upright in the Gazebo task. The tracks indicate recent head and hand moves Middle: When planning a task the instructor (on the left) moves much more frequently than the student (on the right). Right: A participant manipulates an object whilst making sure that they do not touch the physical wall of the IPT.

4.7 Finding Your Way Together and Orienting Yourself

Moving and navigation are essential to all of the tasks, but we find that participants had some trouble with these. Unlike the naturalness with which they were able to manipulate objects and talk about their surroundings, locomotion over short and long distances seems to remain a little problematic and orientation is occasionally tricky.

Being able to orient oneself requires that the environment have a reasonable number of landmarks. For most tasks the environment is small and distinctive enough that orientation is not a problem. However in the Landscape task, where participants had to find different landmarks in an environment, all of the pairs of participants became lost.

Moving around together can also be problematic. This seems to be because of participants having different skill levels with the joystick controls. For example, in the Landscape, participants would often find that one person was racing ahead or lagging behind. In the Landscape task they would also run ‘through’ each other’s avatars occasionally, though whether they noticed this or not depended on the stage in the task. At the beginning, or when they were not preoccupied with other things, they might comment on this. For example one participant said ‘Sorry, I’m running through you again’ when he collided with his collaborator. However a short while later when this same pair was preoccupied with finding their way around, they ran through each other several times without mentioning it at all.

In large open spaces like the Landscape, and also in confined spaces such as the rooms in the Whodo and Poster tasks, we observed participants moving backwards and forwards in a zigzag manner. These patterns were quite unlike any that would be encountered in the real world. For example, when they needed to read the sayings on the posters, they would zigzag backwards and forwards instead of stepping or turning side wards as one might do in a real world art gallery. For example, in the 4th frame of Figure 3 the remote participant steps through the cube he has just put down whilst trying to turn to the left. In the Modelling World, often when participants had to fetch objects from the periphery of the world, when coming back to the centre they would often use the joystick to reverse backwards instead of turning around and going forwards as one would expect in real world. Such movements would often cause them to step through objects or the other avatar.

Overall we see that the participants usually try to avoid walking through each other. However because of the difficulty of moving with the joystick, there are occasional collisions with objects and their partner. These are often quickly rectified, by reversing or stepping out of the way. In the Landscape the inability to move and orient accurately was a definite cause of frustration. Fortunately, for most of the other tasks, a relatively a small proportion of the time was spent moving about the space and participants did not get frustrated.

5 Conclusion

In desktop distributed VEs, participants follow certain norms; for example in relation to greeting and acknowledging the other, maintaining a sense of personal space, not passing through each other and positioning oneself in relation to one's interaction partner (Becker and Mark, 2002). The same applies to immersive VEs, though here the conventions apply to life-size avatars, their movements and forms of non-verbal communication. In both types of environments, this following of real world norms appear to be applied selectively: at different stages in the interaction, the norms can be contravened. We have seen that this selectivity applies to every major facet of interaction including handling objects, navigating, communicating non-verbally, maintaining awareness, and engaging with the virtual landscape or with virtual spaces such as walls. The selectivity appears to depend on whether the objects or avatars are the focus of attention. If an object isn't the focus of attention, the norms might be broken accidentally or for reasons of expediency.

It is worth pointing out that the tendency of participants to follow the conventions of interpersonal behaviour did not depend on whether the person one is interacting with is a stranger or someone whom one knows well. Familiarity does not seem to change how participants interact with each other (Steed et al. 2003). It is also noteworthy that what we have observed applies not only to short-term tasks but also to longer (three hour plus) long sessions in networked IPTs. In other words, it is not a question of 'getting used to' the system.

The overall conclusion is that although IVEs have great promise for tele-collaboration, we have not yet reached a point at which communication can be completely natural. It is worth stressing that we have seen participants use quite subtle forms of real-world behaviour to support virtual interaction: small bodily gestures and verbal communication to signal to the other person. The form of IVE displays seems to overcome certain restrictions of other types of display such limited peripheral awareness and difficulty in making fluid gestures with accompanying dialogue. However movement about the world still seems to be a problem.

There are two main suggestions for further work. The first is that to have collaborative interactions that are more similar to those in the real world, much more of the non-verbal behaviour needs to be captured or simulated. The second is that interaction design needs to focus not only on efficiency and reliability of single person interactions, but on similar considerations for multi-person interactions.

Acknowledgements

Thanks to all of those who took part in running the studies that we have drawn upon. At UCL these include Vinoba Vinayagamorthy, Andrea Brogni, Maia Garau, Mel Slater and David Swapp. At Salford these include Robin Wolff and Oliver Otto. At Chalmers these include Iona Heldal, Ann-Sofie Axelsson and Maria Spante.

References

- Axelsson, A., Abelin, A., Heldal, I., Schroeder, R. & Widstrom, J. (2001) Cubes in the Cube: A Comparison of a Puzzle-Solving Task in a Virtual and a Real Environment. *Cyberpsychology & Behavior*, 4(1), 279-287.
- Becker, B. & Mark, G. (2002). Social Conventions in Computer-mediated Communication: A Comparison of Three Online Shared Virtual Environments. In R.Schroeder (Ed.). *The Social Life of Avatars: Presence and Interaction in Shared Virtual Environments*. London: Springer, 19-39.
- Churchill, E., Snowdon, D. & Munro, A. (Eds.) (2002), *Collaborative Virtual Environments: Digital Places and Spaces for Interaction*. London: Springer.
- Cruz-Neira, C., Sandin, D.J. & DeFanti, T.A. (1993) Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. *Proceedings of Computer Graphics (SIGGRAPH)*, pp. 135-142.
- Frécon, E., Stenius, M. (1998) DIVE: A scaleable network architecture for distributed virtual environments. *Distributed Systems Engineering Journal*, 5, 91-100.
- Frécon, E., Greenhalgh, C. & Stenius, M. (1999) The DiveBone - An Application-Level Network Architecture for Internet-Based CVEs. *ACM Symposium on Virtual Reality Software and Technology*, University College London, UK, December 20 – 22, pp. 58-65.
- Garau, M., Slater, M., Vinayagamoorthy, V., Brogni, A., Steed, A. & Sasse, M.A. (2003). The Impact of Avatar Realism and Eye Gaze Control on Perceived Quality of Communication in a Shared Immersive Virtual Environment. *Proceedings of the SIG-CHI conference on Human factors in computing systems*, April 5-10, 2003, Fort Lauderdale, FL, USA.

- Greenhalgh, C. (2001) Understanding the network requirements of Collaborative Virtual Environments. In Collaborative Virtual Environments: Digital Places and Spaces for Interaction, D. Snowdon, E. F. Churchill, and A. J. Munro, Eds. London, UK: Springer Verlag, pp. 56-76.
- Heldal, I., Schroeder, R., Steed, A., Axelsson, A.-S., Spante, M. & Widestrom, J. (2005). Immersiveness and Symmetry in Copresent Scenarios. Proceedings of IEEE VR, Bonn, Germany.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S. & Greenhalgh, C. (2000) "Object-Focused Interaction in Collaborative Virtual Environments", ACM Transactions on Computer-Human Interaction (ToCHI), vol. 7, pp. 477-509.
- Meehan, M., B. Insko, M.C. Whitton & F.P. Brooks Jr. (2002) Physiological Measures of Presence in Stressful Virtual Environments. ACM Transactions on Graphics, 21(3), 645-652. (Proc. of ACM SIGGRAPH'02, San Antonio, TX).
- Pinho, M. S., Bowman, D. A. & Freitas, C. M. D. S. (2002) Cooperative Object Manipulation in Immersive Virtual Environments: Framework and Techniques. ACM Symposium on Virtual Reality Software and Technology (VRST), Hong Kong, China.
- Roberts, D., Wolff, R., Otto, O., Steed, A. (2003) Constructing a Gazebo: Supporting Team Work in a Tightly Coupled, Distributed Task in Virtual Reality. Presence: Teleoperators and Virtual Environments, 16(6), 644 - 657.
- Ruddle, R. A., Savage, J. C. & Jones, D. M. (2002) Symmetric and asymmetric action integration during cooperative object manipulation in virtual environments, ACM Transactions on Computer-Human Interaction (ToCHI), vol. 9, pp. 285-308.
- Ruddle, R. A., Savage, J. C. & Jones, D. M. (2002) Verbal communication during cooperative object manipulation. Proceedings of the ACM Conference on Collaborative Virtual Environments (CVE'02), pp. 120-127.
- Schroeder, R., Steed, A., Axelsson, A.-S., Heldal, I., Abelin, Å., Wideström, J., Nilsson, A. & Slater, M. (2001) Collaborating in networked immersive spaces: as good as being there together? Computers and Graphics, 25, 781-788.
- Schroeder, R. (Ed.) (2002). The Social Life of Avatars: Presence and Interaction in Shared Virtual Environments. London: Springer.
- Schroeder R. & Axelsson, A.-S. (Eds.). (2005). Avatars at Work and Play: Collaboration and Interaction in Shared Virtual Environments. London: Springer.
- Slater, M. & Wilbur, S. (1997) A framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual environments, Presence: Teleoperators and Virtual Environments, 6(6), 603 - 616.
- Slater, M., Sadagic, A., Usoh, M. & Schroeder, R. (200) Small Group Behaviour in a Virtual and Real Environment: A Comparative Study. Presence: Teleoperators and Virtual Environments, 9, 37-51.
- Steed, A., Mortensen, J., Frécon, E. (2001) Spelunking: Experiences using the DIVE System on CAVE-like Platforms. In B. Frohlicj, J. Deisinger, and H-J. Bullinger (eds) Immersive Projection Technologies and Virtual Environments 2001, 153-164. Springer-Verlag/Wien.
- Steed A., Spante, M., Schroeder R., Heldal I., Axelsson A-S (2003) Strangers and Friends in Caves: An Exploratory Study of Collaboration in Networked IPT Systems for Extended Periods of Time. ACM SIGGRAPH 2003 Symposium on Interactive 3D Graphics, 27 - 30 April 2003 Monterey, California.