

Traversable Interfaces Between Real and Virtual Worlds

Boriana Koleva, Holger Schnädelbach, Steve Benford and Chris Greenhalgh

School of Computer Science & Information Technology

The University of Nottingham

Nottingham NG8 1BB, UK

+44 115 951 4203

{bnk, hms, sdb, cmg}@cs.nott.ac.uk

ABSTRACT

Traversable interfaces establish the illusion that virtual and physical worlds are joined together and that users can physically cross from one to the other. Our design for a traversable interface combines work on tele-embodiment, mixed reality boundaries and virtual environments. It also exploits non-solid projection surfaces, of which we describe four examples. Our design accommodates the perspectives of users who traverse the interface and also observers who are present in the connected physical and virtual worlds, an important consideration for performance and entertainment applications. A demonstrator supports encounters between members of our laboratory and remote visitors.

Keywords

Mixed Reality, Virtual Environments, Augmented Reality, Tele-presence, Tele-embodiment

INTRODUCTION

Various technologies have been developed to allow people to experience remote environments. These might be virtual environments that are experienced through virtual reality technologies or physical environments that are experienced through tele-embodiment and tele-presence technologies. A thread running through this research is the idea of using immersive technologies to establish the illusion of entering the remote environment, resulting in a sense of presence.

A major weakness in this illusion is that users clearly do not leave their physical environment behind them when they enter a remote environment. They remain firmly and visibly present within their local physical space. This is a problem for two reasons. First, their own illusion of remote presence may be destroyed by distractions from the local physical space. Examples can be found in previous experiments with virtual reality. In studies of presence in single user virtual environments, users reported that 'breaks in presence' were caused by background noise and interference from hardware such as cables [11]. Bowers *et al.* note how conduct in a collaborative virtual environment was disrupted by events in the physical environments of the participants [4]. Second,

observers of the interaction can clearly see that participating users have not gone anywhere. This is a particular problem if the interaction is being staged at least in part for the benefit of these observers, for example as part of an entertainment or performance application. It might also be a problem if these observers may themselves become participants at a later date. For example, if they are waiting their turn in an entertainment application or in a shared working environment due to the limited availability of equipment.

Our response to these problems is the concept of traversable interfaces. These enhance the illusion of immersion by making it appear that participants leave their local physical environment in order to enter into a new remote environment. They aim to do this in a way that makes sense to the participants who are entering the remote environment, to observers who are already in the remote environment, and to observers who remain behind in the local physical environment. Our discussion will focus on traversal between physical and virtual environments. However, a traversable interface could also be used in a tele-presence application to link a local physical environment to a remote physical environment.

Further motivation for traversable interfaces is provided by recent work on mixed reality. Paul Milgram has classified mixed reality technologies according to a 'virtuality continuum' [9]. At one extreme of this continuum we find purely physical environments and at the other purely virtual environments. In between, we find augmented reality where physical environments are enhanced with digital information, and augmented virtuality, where virtual environments are enhanced with physical information.

Traversable interfaces provide a mechanism for people to dynamically relocate themselves along this continuum. At one moment they may be primarily located in augmented reality, with a view into an adjoining virtual environment. They may then traverse the interface and find themselves primarily located within an augmented virtuality, with a view back into a physical environment. Traversal allows people to move back and forward between primarily real and primarily virtual environments, repositioning themselves along the virtuality continuum, according to their interest and whether they want the physical or virtual to be their primary focus.

To be published in the proceedings of CHI2000

www.acm.org/chi2000

TRAVERSABLE INTERFACES

We begin with a general design for a traversable interface. Figure 1 summarises the illusion that we wish to create. On the left we see a physical environment that is connected to the virtual environment on the right. Our design needs to consider the perspectives of the four classes of participant, A, B, C and D. A is an observer in the physical environment. B is an observer in the virtual environment. C is crossing from physical to virtual, and D is crossing from virtual to physical.




Figure 1: the illusion that we wish to create

An important point is that the illusion should potentially work for all of these classes of participants, although some applications may give priority to one class over another. For example, a performance might require that the audience believe the illusion, while the performers could be aware of the mechanisms involved. This observation challenges traditional approaches to interface design that have focussed on the experience of the direct participant, but have tended to neglect the experience of observers. We suggest that this is an important consideration for any application where an interface is deployed in a shared or “public” environment, including office environments as well as performance and entertainment applications.

Two other general points should be noted. First, objects as well as participants might traverse the interface. Second, partial traversal might be possible, for example pushing a limb through the interface. However, in this paper we restrict our consideration to complete traversal by humans.

Our general design for a traversable interface integrates a number of techniques:

- mixed reality boundaries [2] for creating windows between physical and virtual environments.
- tele-embodiment for allowing remote virtual participants to enter a physical environment [8,10].
- immersive interfaces for accessing virtual environments, including head-mounted displays (HMD) and projected displays ranging from single screens up to multi-surface CAVEs [5].
- non-solid projection surfaces to allow participants to seemingly pass through a projected image, moving from a public to a more private physical space.

The following sections describe how these are integrated into an overall design, beginning with the idea of mixed reality boundaries.

Mixed reality boundaries

Mixed reality boundaries represent a specific approach to mixed reality that involves creating transparent windows between physical and virtual environments so that occupants of each can communicate with the other [2]. In contrast to other approaches that focus on superimposing the two environments on top of one another (e.g., augmented reality typically overlays a virtual environment on top of a physical environment), the spaces on either side of the boundary are adjacent, but remain distinct. A feature of this approach is that multiple boundaries might be used to join together many different physical and virtual environments into a larger mixed reality structure.

Figure 2 shows how a simple mixed reality boundary can be created. On the left is a physical environment and on the right a virtual environment. An image of the virtual environment is projected into the physical environment and an image of the physical environment captured from a video camera is displayed as a live video texture within the virtual environment. The physical and virtual cameras and projections are aligned so that the images appear to be the reverse sides of a common boundary.




Figure 2: a simple mixed reality boundary (from [2])

A variety of mixed reality boundaries might be created with different properties in terms of their ‘permeability’, the extent to which they allow information and objects to pass across them; ‘situation’, their spatial relationship to the connected spaces; ‘dynamics’, their temporal properties; and ‘symmetry’ [7]. Permeability properties are particularly interesting here because they include the sub-property of ‘solidity’, the extent to which a boundary allows objects and participants to pass through it. This can be broken down into two issues, how to allow participants and objects to enter the remote environment and how to create the illusion that they have left their current environment when doing so.

Entering the remote environment

Entering a remote physical environment can be achieved by taking control of a remote physical proxy such as a robot. The field of tele-robotics is well established, particularly in areas such as working in hazardous environments such as outer space and the deep ocean. Of more direct relevance here is recent work on tele-embodiment in collaborative

settings, where participants take control of a physical proxy or surrogate [8]. In one recent example, participants control a tele-embodiment called a Personal Roving Presence (PRoP) that is armed with a video camera, microphones and speakers, and steer it round a remote environment in order to meet and converse with others [10]. Designs for early PRoPs include ‘space browsers’, helium filled blimps that act as airborne tele-robots and ground based platforms called ‘surface cruisers’. By placing a PRoP on the physical side of a mixed reality boundary and integrating the controls for this PRoP and the video and audio from it within the virtual environment, participants on the virtual side could enter the physical.

An alternative approach towards introducing remote virtual participants into a physical environment would be to use shared augmented reality technology such as [3]. See-through HMDs could display avatars superimposed onto the physical scene. In fact, this could be combined with the use of PRoPs. The position of the PRoP could be tracked and the image of the avatar superimposed upon it.

Techniques that allow a user in a physical environment to enter a remote virtual environment are well known and include a range of immersive displays including HMDs and different tracking and interaction mechanisms for interacting with a projected image of a virtual environment.

Leaving the current environment

The illusion of traversal requires that a user is seen to leave their current local environment when they enter the remote one. We propose that this may be achieved by using non-solid projection surfaces so that the user can appear to directly step into and through the image of the remote environment.

This is straightforward in the virtual environment. The image of the remote physical environment is displayed as a video texture attached to a graphical object. This can be non-solid, enabling avatars to pass through it.

It is more difficult in the physical environment. Later on, we shall describe four different approaches that we have implemented involving projection onto non-solid materials such as water, the use of fabric curtains as well as mechanical devices such as sliding doors and movable screens. For the remainder of this section we shall assume the existence of such technologies.

It should be noted that in all cases, what actually happens is that the user passes from a public space through the image, into a more private space beyond. From the physical environment they move to a physical antechamber beyond the screen where they find the immersive technology required to enter the virtual environment. From the virtual environment, their avatar moves to a virtual antechamber beyond the screen where they may find the controls to access a PRoP. The physical antechamber may take on a variety of forms. In a performance, the public space will be the key focus of activity, with the antechamber being ‘the

wings’ or behind the scenes. Conversely, the antechamber might be the main focus of the activity, for example it might be a CAVE installation [5], with the traversable interface providing an entry point to and from the outside world.

An integrated design

Figure 3 shows how the above techniques for entering and leaving physical and virtual environments can be integrated into a general traversable mixed reality boundary.




Figure 3: creating a traversable boundary

On the left is a physical environment containing a non-solid projection surface onto which is projected a view of the remote virtual environment. Behind this is an antechamber containing the immersive technology required to become embodied within the virtual environment. On the right is a virtual environment containing the video view into the physical environment. Behind this is a virtual antechamber that contains controls for a remote PRoP and that also contains a second video texture showing the view from this PRoP as it moves around the physical environment.

We can now consider how the four participants A, B, C and D from figure 1 will experience this design. Participant A is the observer in the physical environment. They will see participant B through the mixed reality boundary. They will see participant C step through the physical projection screen, apparently into the virtual world. At the same time, they will see C’s virtual avatar, C^V emerge into the virtual world. They will see participant D’s avatar approach the projection screen and then disappear from view. D’s ProP, D^P, will then emerge through the physical screen.

Participant B is the observer in the virtual environment. They will see A through the boundary and will see C approach them in the video view, disappear and then replaced by C’s avatar, C^V, appearing through the video texture. They will see D’s avatar approach the video texture, pass into it and then see D’s physical proxy, D^P, appear in the video image.

Participant C traverses from the physical to the virtual. They will step through the physical projection screen, entering the physical antechamber. There they will find the technology required to independently access the virtual environment. This might be a headmounted display, desktop computer, CAVE, specialised vehicle (for example, a ‘pod’ in a simulation ride) or further projected display. Their avatar will initially appear in the virtual antechamber and they will

then steer it through the video texture into the public virtual environment.

Participant D traverses from the virtual to the physical. They will steer their avatar through the video image of the remote physical environment, entering the virtual antechamber. Here they will find the virtual controls for the remote PROP, D^P, as well as a further video texture showing the view from its onboard camera. They will then be able to steer the PROP from the physical antechamber, through the physical projection surface into the public physical environment.

Design considerations

This design for a traversable interface is a general one. A particular realisation will have to make a number of specific design choices in order to meet the two goals of traversable interfaces as outlined in the introduction.

The first goal was to minimize distractions for participants who wish to become present in a remote environment. This is achieved by locating the VR equipment required to access this environment in a private antechamber. This can be designed to provide an optimal operating environment for this equipment, for example, being painted and lit to support video tracking, being free of other equipment that might interfere with electromagnetic tracking, and generally being free of clutter on which the user might snag themselves.

The second goal was to create the illusion of physically leaving the current environment in order to enter a new remote environment. Successfully meeting this goal will require considering the following design issues.

The physical and virtual antechambers can be decorated to support the transition to the new environment. For example, in a theme park ride, the physical antechamber might be modeled to match the virtual world. If the user thinks that they were going to pass into a virtual cave, then this antechamber should look like that cave. The physical and graphical design of PROPs and avatars can also support the illusion of traversal. In a theme-park ride, the PROP might be a sophisticated animatronic figure (such figures are already used in theme-parks). Likewise, the positions of physical bodies, PROPs and avatars at the key transition points will be important. With careful design, it may be possible to make them appear to directly replace one another, to be overlaid on one another, or to time the sequence of appearances and disappearances to reinforce the illusion of traversal.

Traditional theatrical techniques may be used to enhance or alter the illusion of traversal, including changes in lighting, the use of smoke and sound effects. Another key effect is the use of shadows. Several of the non-solid projection surfaces that we introduce below can be configured to show the physical user beyond the screen as shadow. In some cases it will be important to avoid shadows so as to maximise the illusion of traversal. In others, the silhouette of a participant's body seen against the image of the virtual environment may be used for its artistic effect (see figure 8)

or as one way of overlaying participants' physical and virtual bodies as noted above.

DESIGNS FOR NON-SOLID PROJECTION SURFACES

The use of non-solid projection surfaces is an essential part of our design. It has also been the most challenging part to realise. This section describes four attempts to construct such surfaces: fabric curtains, water curtains, a sliding door, and a flip-up screen. Figure 4 summarises the four designs and shows examples of each.

Fabric curtain


Curtains are familiar devices for partitioning physical space. Curtains can provide privacy and can be readily traversed, introduced and removed. They have been extensively used in theatre to hide and reveal actors and objects and to give the illusion of transitions between scenes. There are a wide variety of familiar designs of curtains; they can be pulled back, raised, vertically slit and be formed into blinds.

Curtains can be made of materials that can hold a projected image and so represent a natural choice for creating non-solid projection surfaces. Our initial design as shown in figure 4 (a) is based around a number of vertical segments of projection screen fabric, weighted at the bottom to hold their shape. A user can easily push through these and the curtain settles down to its regular shape within a few seconds. We back project the image onto the curtain by bouncing it off of a mirror on the ceiling. This creates an area in the antechamber where a participant may stand or sit without casting a shadow onto the screen. Conversely, they may be deliberately positioned so as to create a shadow for artistic effect as noted above. Figure 4 (b) shows a participant emerging through the curtain.

Water curtain

We have also experimented with a second curtain – a curtain of water. In 1998 we began collaborating with the performing arts company Blast Theory who were already experimenting with projecting images and video into a vertical curtain of water. Projection into water has also been explored in other contexts. For example, Disney-MGM studios projected film clips into fountains and a water screen as part of a dream sequence in their "Fantasmic" show in their October 1998 program.


The overall design of the water curtain is shown in figure 4 (c). The curtain is produced by several fine spray nozzles (originally designed for spraying pesticide) attached to a metal pipe that is suspended roughly two meters above a trough on the ground. Water is pumped through the pipe, descends as a fine spray about half a meter thick and is collected from the trough and recycled. Figure 4 (d) shows this physical infrastructure. The water curtain holds a back-projected image surprisingly well, although early experimentation showed that the projector needs to point straight at the curtain, making shadows unavoidable as participants pass through it.



(a) Fabric curtain design




(b) Emerging through the fabric curtain



(c) Rain curtain design




(d) Rain curtain infrastructure (e) Rain curtain in use



(f) Sliding door design



(g) Opening the sliding door



(h) Flip-up screen design



(i) Raised as ambient display (j) Raising the screen

Figure 4: four designs for non-solid projection surfaces

Being completely fluid, a person or object can pass through the water curtain much more seamlessly than they can with a fabric curtain (so long as they are prepared to get wet!). It is also transparent when viewed from behind, allowing for easy observation of its users (e.g., by performers who can then time their emergence through the curtain to match the user's actions). Like a fabric curtain, the water curtain can be readily introduced and removed by switching the pump on and off. Holes can be dynamically punched through it by using solid objects to interrupt the flow of the water. Finally, it has a powerful aesthetic, in terms of the continually shifting quality of the visual image, the sound of the water and its physical feel.

In January 1999 we staged a public demonstration of using a water curtain as an interface to a virtual environment. Participants undertook a journey through a virtual world, during which they were interrupted by a performer emerging through the curtain – an event that had a significant theatrical impact. Figure 4 (e) shows the performer emerging through the water curtain. We are currently developing a full-scale public performance that will involve the use of six rain curtains to allow an audience to experience a shared virtual environment.

Sliding door

Unlike a curtain, a door is a solid projection surface that is traversed by physically moving a large section of it. As with conventional doors, there are many potential designs including hinged, sliding and rotating. Our first design has been a sliding door made from perspex as shown in figure 4 (f). Figure 4 (g) shows a participant opening the door in order to step through it.

The sliding door has several interesting properties. Being solid, it can more easily be locked than a curtain, allowing participants to minimise possible interruptions. Its solidity also favours applications where it is part of a more permanent architectural framework. Early tests suggest that our sliding door can simultaneously hold two different images, one on each side, provided that the images that have similar contrasts (otherwise there is too much visible interference between the two). This potentially saves space, as it only requires one projection surface to display both public observers' and immersed participants' views of the virtual environment. The properties of solidity and holding multiple images could usefully be combined in using a sliding door as the entrance to a CAVE. One surface of the CAVE could be slid open to allow participants to enter. Visitors remaining on the outside could see a specially tailored (e.g., without head-tracking) public view of the activity in the CAVE on the outside of the door.

Flip-up screen

Our final example is a flip-up screen as shown in figure 4 (h). This is a screen that can be moved from a vertical to a horizontal position at ceiling height, allowing people to pass underneath it. Figure 4 (i) shows a participant raising the screen. The flip-up screen is essentially a specialised form of

door. However, it has the additional property of being able to act as an ambient display surface when in the raised position, reflecting the idea of ambient display media proposed in [6]. This is possible because the projected image is bounced off of the mirror on the ceiling and hits the screen when it is in both its vertical and horizontal positions.

This property suggests an alternative mode of use to the previous examples. Instead of stepping through the projected image, the user may remain in one physical location, but choose to lower or raise the flip-up screen according to whether their interaction is primarily focussed in the physical or the virtual environment. To focus on the physical environment, the user raises the screen, opening up their physical space to the public space beyond and displaying a peripheral image of the virtual environment on the ceiling. Figure 4 (j) shows a participant who is focussed on a task in the physical world and so has set the flip-up screen to its ambient position. To focus on the virtual environment, they place the screen in its vertical position, shielding their local physical environment from the public space beyond, and providing users in this public space with an image of their avatar in the virtual environment instead of their physical self using the immersive technology. In this way participants can reposition themselves along Milgram's virtuality continuum as noted in the introduction.

An extension to this approach would be to use the physical raising and lowering of the screen to drive a switch to automatically configure a user's local environment according to whether they were currently in the physical or virtual environment. The switch might configure lighting and tracking technologies and might minimise distractions, for example by routing the user's phone to their voice mailbox when they were immersed in the virtual environment. This reflects previous work on using physical doors to manage electronic privacy in an office environment, using a so-called "doormouse" [12].

In summary, we have realised four different kinds of non-solid projection surface that might be used in traversable interfaces. These can be broadly grouped into the two categories of curtains (fabric and water) and doors (sliding and flip-up). The curtains potentially offer the most seamless illusion of traversal and could be especially suited to performance, art and entertainment. The doors provide a less fluid illusion of traversal, but may offer some practical advantages for use in everyday environments such as offices and the home. Of course, there are many other possibilities. Perhaps we can use other materials such as smoke to create highly fluid projection surfaces, and no doubt there are other possible mechanical designs based on doors and curtains.

DEMONSTRATION

We have developed a demonstration of a traversable interface in order to show that our design is technically implementable. It should be noted that we do not claim that this is yet a real or effective application, although our future plans involve developing and evaluating such an application.

Our demonstration has been constructed in our laboratory. Its aim is to provide a social space where lab members can meet with visitors who “drop in” over the Internet. A mixed reality boundary allows lab members and visitors to see and talk to one another. Both can also traverse this boundary. A single visitor at a time can take charge of a simple P_{RoP} and use it to explore an area of the laboratory. A single lab member at a time can step into the virtual world to become part of a virtual meeting. Figure 5 shows the collaborative virtual environment that we are using in our demonstrator. This has been realised using the MASSIVE-2 system [1]. The image shows the video texture that forms half of the mixed reality boundary with the physical environment.



Figure 5: the virtual environment with video texture



Figure 6: controlling the P_{RoP} from MASSIVE-2.

Avatars can step through this boundary to enter a small virtual antechamber where they find the interface to control our remote P_{RoP}. This consists of a second video texture that shows the view from the P_{RoP}'s on board camera as well as six buttons, four to move the P_{RoP} forwards and backwards and to rotate it left and right, and two to tilt the camera up and down vertically. Figure 6 shows the view over a remote user's shoulder when they have just entered this antechamber. Inset is a close up of the virtual controls for the P_{RoP}.

The P_{RoP} itself is a small wireless robot that has been constructed using a LEGO Mindstorm kit (see figure 7).

This platform can be moved around the floor, includes a raisable arm for the camera and can be controlled over an infrared link. A small wireless video camera and microphone have been mounted on the platform along with a pen-torch to illuminate nearby objects. The wireless connections currently have a limited range and there are as yet no on-board speakers (so the P_{RoP} can see and hear, but not talk). The P_{RoP} is also rather small, standing at approximately one foot tall. However, it does provide an inexpensive workable solution for initial demonstrations and application development.




Figure 7: the P_{RoP}



Figure 8: immersed in the virtual environment

The physical side of the boundary can utilise the fabric curtain, sliding door or flip-up screen designs. The images in figures 4 (b), 4 (g) and 4 (j) all show examples of the view looking into our virtual environment, as if from out of the video texture. In each case a video camera is mounted on the top of the frame of the boundary to provide the video view shown in figure 5. This positioning is less than ideal as the two sides of the boundary are not strictly spatially aligned and a solution that allows a small camera to be located in the centre of the screen is ideal. Mounting the camera in the centre of the flip-up screen would also allow it to provide a peripheral view from above the user's workspace when in the raised (ambient) position. Having traversed the physical screen, the user enters the physical antechamber and finds

equipment to access the virtual world. Figure 8 shows an example where the user has stepped through the screen and has donned a head-mounted display. In this case, they have been deliberately positioned so that we see their shadow.

SUMMARY

This paper has developed the idea of traversable interfaces that give the illusion that participants in a local physical environment can completely cross into a remote virtual (or indeed physical) environment and vice versa. The key innovation in the paper is the extension of the now familiar illusion of entering a remote environment to include appearing to leave one's current environment. We have argued that this is particularly important when the interaction may be observed by people in the two environments as well as experienced directly by the participants. This will be the case in many performance and entertainment applications, but will also be relevant whenever virtual environments and tele-presence technologies are deployed in shared environments, be they public, working or domestic.

We have presented a general design for a traversable interface between a physical and a virtual environment that combines three key components. The first is the use of Physical Roving Proxies (PRoPs) to allow a virtual participant to enter a physical environment. The second is the use of VR technologies to allow a physical participant to enter a virtual environment. The third is the use of non-solid projection surfaces to allow a participant to seemingly step into a projected image of a remote environment. We have presented four early designs for non-solid projection surfaces, a fabric curtain, a water curtain, a sliding door and a flip-up screen. Finally, we have described a demonstrator that shows one possible realisation of our design.

Among the most obvious applications of traversable interfaces are entertainment applications where it may be important to establish a strong illusion of entering a virtual environment. VR-based theme park rides that wish to smooth the transition between watching the ride while waiting for a turn and entering the ride as a participant are a particularly strong candidate, especially as such rides already use animatronic figures and participants occasionally get wet! We also anticipate that our design might be incorporated into more general immersive interfaces. For example, a traversable interface based on our sliding door design might form one side of a CAVE facility, allowing passage to and from the CAVE and providing an external public view of the activity inside.

Our future plans involve developing and evaluating real applications of traversable interfaces. Evaluation will employ ethnographic techniques of the kind that have been previously used to study social interaction in collaborative virtual environments (e.g., [4]).

We would like to finish by reinforcing two points that have more general relevance to human-computer interaction.

First, is the idea that shared and public interfaces need to be designed with third party observers in mind as well as direct participants. Second, is the observation that virtual reality and telepresence technologies have always been concerned with creating an illusion – the illusion of entering a new and remote environment. This paper has explored how more traditional theatrical effects, such as moving screens and curtains, and changes in lighting might enhance this illusion, an approach that might be applied to the design of a wide range of human-computer interfaces.

ACKNOWLEDGMENTS

We thank the ESPRIT IV I³ programme for supporting this work through the eRENA project and the EPSRC for their support through PhD studentship awards. We are indebted to Blast Theory for their collaboration on the rain curtain.

REFERENCES

1. Benford, S. D., Greenhalgh, C. M. and Lloyd, D., Crowded Collaborative Virtual Environments, *Proc. CHI'97*, pp. 59-66, 1997
2. Benford, S., Greenhalgh, C., Reynard, G., Brown, C., and Koleva, B., Understanding and Constructing Shared Spaces with Mixed Reality Boundaries, *ACM ToCHI*, 5 (3), 185-223, ACM Press.
3. Billinghurst, M., Weghorst, S., Furness III, T., Shared Space: An Augmented Reality Interface for Computer Supported Collaborative Work, *Proc. CVE'96*, Nottingham University, UK, Sept 1996.
4. Bowers, J., Pycocock, J. and O'Brien, J., Talk and Embodiment in Collaborative Virtual Environments, *Proc. CHI'96*, 1996.
5. Cruz-Neira, C., Sandin, D. J., DeFant, T. A., Kenyon, R. V. and Hart, J. C., The Cave - Audio Visual Experience Virtual Environment, *CACM*, 35 (6), pp 65-72, 1992.
6. Ishii, H. and Ullmer, B., Tangible bits: Towards seamless interfaces between People, Bits and Atoms, *Proc. CHI'97*, pp 234-241, ACM Press, 1997.
7. Koleva, B., Benford, S. and Greenhalgh, C., The Properties of Mixed Reality Boundaries, *Proc. ECSCW'99*, Copenhagen, September 1999, Kluwer.
8. Kuzuoka, H., Ishimoda, G. and Nishimura, T., Can the GestureCam be a surrogate?, *Proc. ECSCW'95*, Stockholm, Sweden, September 1995, Kluwer.
9. Milgram, P. and Kishino, F., A Taxonomy of Mixed Reality Visual Displays, *IEICE Transactions on Information Systems*, Vol E77-D (12), Dec. 1994.
10. Paulos, E. and Canny, J., PRop: Personal Roving Presence, *Proc. CHI'98*, pp 296-303, 1998.
11. Usuh, M., Arthur, K., Whitton, M., Bastos, R., Steed, A., Slater, M. and Brooks, F., Walking > Walking-in-Place > Flying, in *Virtual Environments, Proc. SIGGRAPH'99*, pp359-364.
12. Buxton, B., Living in Augmented Reality: Ubiquitous Media and Reactive Environments, *Proc. Imagina'95*. Monte Carlo, February 1995.

