

# MANAGING MUTUAL AWARENESS IN COLLABORATIVE VIRTUAL ENVIRONMENTS<sup>1</sup>

STEVE BENFORD

*Department of Computer Science, The University of Nottingham,  
Nottingham, NG7 2RD, UK*

JOHN BOWERS

*Department of Psychology, The University of Manchester,  
Manchester, UK*

LENNART E. FAHLÉN

*The Swedish Institute of Computer Science (SICS),  
S-16428 Kista-Stockholm, Sweden*

CHRIS GREENHALGH

*Department of Computer Science, The University of Nottingham,  
Nottingham, NG7 2RD, UK*

## ABSTRACT

This paper introduces a spatial model of interaction which aims to support groups of people in using their natural communication skills in distributed virtual environments. First, we outline our motivation for undertaking this work in terms of the social significance of space in supporting co-operative work. Next, we summarise the model's main concepts, namely aura, awareness, focus, nimbus, adapters and boundaries. Following this, we describe routes to implementing the model. The first is an approach suited to realising the model as an application of existing VR platforms. The second is a general implementation of the model, based on its own specially tailored distributed architecture. Finally, we present our initial observations from use in the laboratory setting and outline issues for future work.<sup>2</sup>

---

<sup>1</sup> To appear in Proceedings VRST'94, Singapore, August 1994, ACM Press.

<sup>2</sup> This work has been funded under the European Commission's ESPRIT programme (the COMIC project) and by the UK's EPSRC.

## 1. Introduction and motivation

This paper explores the use of distributed virtual environments to support cooperative work. At the core of our paper is the contention that the provision of shared virtual spaces can enable people to employ a rich set of social skills for managing interaction and the sharing and exchange of information. Several distributed multi-user VR systems have already been described in the literature including DIVE [Fahlén, 93], dVS [Grimsdale, 91], Aviary [West, 93], VEOS [Bricken, 93], the MR Toolkit [Green, 1994], Rubberrocks [Codella, 92], NPSNET [Zyda, 92] and the work of the ATR Lab [Takemura, 92]. The main goal of the paper is to describe realisations of a *spatial model of interaction* which supports people's social skills in crowded virtual computer spaces and which could eventually be realised in such systems.

The motivation behind the spatial model stems from key observations concerning the social importance of space arising from the social sciences. Indeed, whereas much previous work on user interface design has exploited people's cognitive spatial skills (e.g. the ability to spatially classify and navigate), our work exploits their *social* spatial skills - i.e. how they use space to manage interaction with one another. The following paragraphs summarise some of the key aspects of space in supporting social interaction.

Drawing on Giddens (e.g. [Giddens, 84]), space constitutes a key resource for establishing and enabling an activity. In particular, the concept of boundaries as a means of segmenting space into different regions is of critical importance for constraining access, marking territory and identifying locations for specific activities. Space also allows different modes of participation in, and awareness of, activity. Beyond this, space provides peripheral awareness of the presence and activity of others, allowing us to 'see at a glance' what is occurring. Such awareness has been identified as a key factor in many real world co-operative activities (e.g. Air Traffic Control [Hughes, 92] and the London Underground control room [Heath, 91]). Again, boundaries play an important role in influencing awareness as well as in controlling movement (consider the different effects of walls, doors, curtains, windows and frosted windows).

Space allows us to negotiate access to common resources. Where such resources are other people we use the term conversation management. Thus, we see people using their body positions and orientations as well as gaze direction to control turn-taking and the joining and leaving of conversational groups (the most extreme example might be the kind of "social dance" that takes place at cocktail parties). Other negotiation behaviours include queuing, jostling and even scrumming where people use space to negotiate access to valuable resources (human or otherwise). A key factor in such negotiations is the way in which continual awareness of others allows people to flexibly manage their own activity in social situations and provides a degree of predictability concerning the likely actions of others (e.g. you can easily see when someone is heading across a room to talk to you).

Our aim is therefore to support such skills and behaviours in virtual environments through the introduction of a set of generic mechanisms which collectively define a spatial model of interaction. In particular, we intend the model to provide an alternative to existing forms of turn-taking and control in co-operative systems (e.g. 'floor control' in conferencing systems). A further goal is to support the development of large-scale shared virtual spaces (i.e. spaces that contain many inhabitants).

## 2. The Spatial Model of Interaction

We now define the key concepts which constitute our spatial model of interaction (the model has been defined in detail in a previous paper [Benford, 93] - this section summarises it here for completeness). The spatial model, as its name suggests, uses the properties of space as the basis for mediating interaction. Thus, objects can navigate space in order to form dynamic sub-groups and manage conversations within these sub-groups. Next, we introduce the key abstractions of *space*, *objects*, *media*, *aura*, *awareness*, *focus*, *nimbus*, *adapters* and *boundaries* which define our model.

### 2.1. Space and objects

The most fundamental concept in our model is *space* itself. We allow virtual spaces to have any number of dimensions where each dimension allows some measure of position. Put another way, space is defined by 'spatial metrics' - well defined ways of measuring position and direction across a set of dimensions. This general definition of space is intended to allow a wide range of potential applications of the model (not only VR).

Space is inhabited by *objects* which might represent people, information or other computer artefacts. Any interaction between objects occurs through some *medium*. A medium might represent a typical communication medium (e.g. audio, visual or textual) or perhaps some other kind of object specific interface. Each object might be capable of interacting through a combination of media/interfaces and objects may negotiate compatible media whenever they meet in space.

### 2.2. Aura

The first problem in any large-scale environment is determining which objects are capable of interacting with which others at a given time. *Aura* is defined to be a sub-space which effectively bounds the presence of an object within a given medium and which acts as an enabler of potential interaction. Objects carry their auras with them when they move through space and when two auras collide, interaction between the objects in the medium becomes a possibility. It is the surrounding environment that monitors for aura collisions between objects. When such collisions occur, the environment takes the necessary steps to put the objects in contact with one another (e.g. exchange of object IDs, addresses, references or establishment of associations or connections). Thus, aura acts as a fundamental technological enabler of interaction. An aura can have any shape and size and need not be around the object whose aura it is. Nor need it be contiguous in space. Also, each object will typically possess different auras for different media (e.g. with different sizes and shapes). Thus, as I approach you across a space, you may be able to see me before you can hear me because my visual aura is larger than my audio aura.

### 2.3. Focus, Nimbus and Awareness

Once aura has been used to determine the potential for object interactions, the objects themselves are subsequently responsible for controlling these interactions. This is achieved on the basis of quantifiable levels of *awareness* between them. The measure of awareness between two objects need not be mutually symmetrical (A's awareness of B

need not equal B's awareness of A). As with aura, awareness levels are medium specific. Awareness between objects in a given medium is manipulated via *focus* and *nimbus*, further subspaces within which an object chooses to direct either its presence or its attention. More specifically, if you are an object in space:

- The more an object is within your focus, the more aware you are of it.
- The more an object is within your nimbus, the more aware it is of you.

The notion of spatial focus as a way of directing attention and hence filtering information is intuitively familiar from our everyday experience. The notion of nimbus requires a little more explanation. In general terms, a nimbus is a sub-space in which an object makes some aspect of itself available to others. This could be its presence, identity, activity or some combination of these. Nimbus allows objects to try to influence others (i.e. to project themselves or their activity to try to be heard or seen). The necessity for a concept of nimbus to complement that of focus becomes clear once we remember the notable phenomenon uncovered in Heath and Luff's work (1991) where London Underground control room operators often projected their activity so as to encourage awareness among others (see section 1). Thus, nimbus is the necessary converse of focus required to achieve a power balance in interaction.

Objects *negotiate* levels of awareness by using their foci and nimbi in order to try to make others more aware of them or to make themselves more aware of others. We deliberately use the word negotiate to convey an image of objects positioning themselves in space in much the same way as people mingle in a room or jostle to get access to some physical resource. Awareness levels are calculated from a combination of nimbus and focus. More specifically, given that interaction has first been enabled through aura collision:

*The level of awareness that object A has of object B in medium M is some function of A's focus in M and B's nimbus in M.*

The choice of function will be application specific. However, for present purposes, the important point is that the resulting quantified awareness levels between two objects can be used as the basis for managing their interaction. Exactly how this is achieved is again a matter of interpretation by a particular application. One approach might be to use awareness levels to directly control the medium (e.g. controlling the volume of an audio channel between two objects). Another might be to allow objects to react to each other's presence depending on specified awareness thresholds (e.g. I might automatically receive text messages from you once a certain threshold had been passed).

#### *2.4. Manipulation, Adapters and Boundaries*

Next we consider how aura, focus and nimbus, and hence awareness, are manipulated by objects in order to manage interactions. We envisage four primary means of manipulation:

1) Aura, focus and nimbus may most often be *implicitly* manipulated through fundamental spatial actions such as movement and orientation. Thus, as I move or turn, my aura, focus and nimbus might automatically follow me.

2) They may on occasion be *explicitly* manipulated through a few key parameters. For example, I might deliberately focus in or out (i.e. change focal length) by simply moving a mouse or joystick.

3) They may be manipulated through various *adapter objects* which modify them in some way and which might be represented in terms of natural metaphors such as picking up a tool. In essence, an adapter is an object which, when picked up, amplifies or attenuates aura, focus or nimbus. For example, a user might conceive of picking up a 'microphone'. In terms of the spatial model, a microphone adapter object would then amplify their audio aura and nimbus. As a second example, the user might sit at a virtual 'table'. Behind the scenes, an adapter object would fold their aura, foci and nimbi for several media into a common space with other people already seated at the table, thus allowing a semi-private discussion within a shared space. In effect, the introduction of adapter objects provides for a more extensible model.

4) Finally, aura, focus and nimbus may be manipulated through *boundaries* in space. Boundaries divide space into different areas and regions and provide mechanisms for marking territory, controlling movement and for influencing the interactional properties of space. More specifically, boundaries can be thought of as having four kinds of effects: effects on aura, effects on focus, effects on nimbus and effects on traversal (i.e. movement). Furthermore, these effects can be of four sorts: obstructive, non-obstructive, conditionally obstructive and transforming. These effects are also defined on a per medium basis and different boundaries may mix these effects in different ways. For example, a 'virtual door' might conditionally obstruct traversal, aura, focus and nimbus (the condition being the possession of a key) whereas a 'virtual window' might obstruct traversal but not obstruct aura, focus and nimbus. Of course, there may also be types of boundary which do not have any real-world counterpart like one way mirrors you can walk through. To generalise a little further, any object can act as a boundary and, although one tends to think of boundaries in terms of walls, doors and windows, we are all familiar with the fact that furniture and even people can make very effective boundaries (e.g. a well placed desk or a row of policemen). Thus, although it may help to explain the model if we conceive of boundaries and adapters as being different kinds of object, they are really the same thing.

Having presented the spatial model, the following sections discuss its implementation. First, section 3 presents a much simplified implementation of the model that might be realised in a less general, but still useful, manner as an application of existing VR systems (i.e. without a special purpose architecture). In section 4 we briefly describe the DIVE system and the even simpler implementation that lies behind the Mr Nimbus spatial model demonstrator. Finally in section 5 we present a general implementation of the model based on a specially tailored distributed architecture.

### **3. A simplified interpretation of awareness**

This section describes a simplified approach to realising the spatial model, based on the concept of "awareness states" which emerges when a few key simplifying assumptions are made. Although somewhat less general, this simplified approach has the advantage that it may be realisable as an *application* of existing VR platforms without re-engineering their distributed architectures. First we make the following simplifying assumptions: aura collision has already occurred and focus and nimbus are simple containment spaces such that an object is either in or out of another's focus and/or nimbus.

Now, considering all the possible arrangements of focus and nimbus for two entities, it is possible to define ten different states of mutual awareness as follows. An entity may direct its focus on the other or not. Similarly, it can project its nimbus on the other or not. This gives four possible ways in which one entity can arrange its nimbus and focus relative to another. As we are dealing with a two entity case, we have  $4^2=16$  possible arrangements of mutual foci and nimbi. However, six of these can be eliminated by symmetry, leaving us with ten unique “awareness states”.

Although much simplified, this representation of the spatial model still captures a surprisingly rich set of communication possibilities. How might we then use this notion of awareness states to implement the spatial model as a VR application? The following outlines two possibilities:-

1. Defining focus and nimbus as basic containment spaces enables them to be represented by simple graphics objects of an appropriate shape. The underlying VR system can detect collisions between these foci and nimbi objects and other objects. Thus, increasing and decreasing awareness becomes a matter of responding to collision events between application objects (collision detection being provided by any self-respecting VR system). Thus, collision with an object's focus increases its awareness of you and collision with its nimbus increases your awareness of it. A typical response in the visual medium might be to change the graphical representation of an object (i.e. showing more detail as awareness increases) resulting in a distancing effect which may be sensitive to both distance and orientation. In this way, all ten of the above awareness states can be enabled as a result of focus/nimbus collisions in space. However, to be useful in a group situation, it must be possible for an object to display itself in different ways to different observers (not something that is supported by all VR systems).

2. We might associate different objects or spaces with different awareness states. Thus, sitting at the conference table might enable full mutual awareness, entering an alcove might disallow overhearing and so on. A room in which participants can move between mutual overhearing and monitoring one of the party's activities might simulate many aspects of working spaces where mutual unobtrusive monitoring is important.

We propose that it may be possible to realise a much simplified, but still useful, version of the spatial model as application code in many existing VR systems without having to provide a specialised distributed architecture.

#### **4. Implementation of the model in DIVE**

The spatial model has been integrated into the DIVE system, the multi-user distributed VR system from the Swedish Institute of Computer Science<sup>1</sup> [Fahlén, 93]. The current demonstrator of the DIVE implementation allows a user to experience the effects of aura focus and nimbus through interaction with an imaginary creature called Mr. Nimbus. Both the user and Mr. Nimbus are equipped with simple graphical representations of aura, focus and nimbus and these are used to trigger and control the volume of an audio signal which is emitted by Mr. Nimbus. The awareness function in this demonstrator detects collisions between focus and nimbus objects, resulting in four possible awareness states (full mutual awareness, no mutual awareness and two asymmetric cases). Another

---

<sup>1</sup> Contact [dive-request@sics.se](mailto:dive-request@sics.se) for details of how to obtain DIVE.

interesting feature of Mr. Nimbus is the way in which his embodiment reflects different awareness states across both the visual and audio media. More specifically, visual cues such as flapping ears (showing his audio awareness of us) high-lighted mouth (our audio awareness of him), widening eyes (his visual awareness of us) and blushing (our visual awareness of him) are used to convey transitions between different awareness states (see figure 1).




Figure 1 Example awareness cues as demonstrated by Mr. Nimbus.

The DIVE system has some inbuilt support for the handling of the spatial model. Objects can for example be programmed to react autonomously to aura, focus and nimbus interactions. Furthermore, there is support for continuous updating of non-discrete awareness levels (not used in the present Mr Nimbus demonstrator). A major problem with a multi-user VR system is of course how to support a number of “subjective” views, one for each participant. In the Mr Nimbus case it is likely that he wants to “smile” towards one participant and flap his ears towards another. The current DIVE version has limited support for this kind of interaction in the form of a set of predefined “views” contained within each object. When rendering a scene each individual visualizer (participant) selects an appropriate “view” as derived from the current “state” for each of the objects present in that scene [Carlsson 93].

In the next section we will describe a distributed virtual reality system that is architecturally based on the spatial model and which provides a richer set of spatial interaction capabilities when compared to the DIVE implementation.

## 5. Implementation of the model in MASSIVE

MASSIVE (Model, Architecture and System for Spatial Interaction in Virtual Environments) is a prototype implementation of the spatial model. The main goals of MASSIVE are scale (i.e. supporting as many simultaneous users as possible) and heterogeneity (supporting interaction between users whose equipment has different capabilities, who employ radically different styles of user interface and who communicate over an ad-hoc mixture of media).

### 5.1. Functionality

MASSIVE supports multiple virtual worlds connected via portals. Each world may be inhabited by many concurrent users who can interact over ad-hoc combinations of graphics, audio and text interfaces. The graphics interface renders objects visible in a 3-D space and allows users to navigate this space with a full six degrees of freedom. The audio interface allows users to hear objects and supports both real-time conversation and playback of pre-programmed sounds. The text interface provides a MUD-like view of the world via a window (or map) which looks down onto a 2-D plane across which users move. Text users are embodied using a few text characters and may interact by typing messages to one another or by "emoting" (e.g. smile, grimace etc.).

These interfaces may be arbitrarily combined according to the capabilities of a user's terminal equipment. Thus, at one extreme, the user of a sophisticated graphics workstation may simultaneously run the graphics, audio and text clients (the latter providing a map facility and allowing interaction with non-audio users). At the other, the user of a dumb terminal (e.g. a VT-100) may run the text client alone. It is also possible to combine the text and audio clients without the graphics and so on. The net effect is that users of radically different equipment may interact, albeit in a limited way, within a common virtual world (e.g. text users may appear as slow-speaking, slow moving flatlanders to graphics users). One effect of this heterogeneity is to allow us to populate MASSIVE with large numbers of users at relatively low cost (a problem for any laboratory that wishes to work in the area of large-scale VR).

All of the above interfaces are driven by the spatial model. Thus, an object can't be seen until graphics auras collide and cannot be heard until audio auras collide. The effects of focus and nimbus are most pronounced in the audio and text interfaces. Audio awareness levels are mapped onto a combination of volume and low-pass filtering with the net effect that audio interaction is sensitive to both the relative distances and orientations of the objects involved (this is noticeable in general conversation). Text messages are also displayed according to mutual levels of awareness as is shown by the following table.

<b>Awareness</b>	<b>Level</b>	<b>Display in text medium</b>
0.0 - 0.2	none	
0.2 - 0.4	presence	Chris at 0,0
0.4 - 0.6	events	"Chris says something"
0.6 - 0.8	peripheral	"(Chris says hi!)"
0.8 - 1.0	full	"Chris says hi!"

Text medium awareness levels and their effects

In the current implementation, users may explicitly manipulate awareness by choosing between three settings for focus and nimbus - normal, narrow (intended for private conversations) and wide (intended for browsing). Two adapter objects are also

provided: a podium which extends the aura and nimbus of its user; and a conference table which replaces its user's normal auras, foci and nimbi with a new set which span the table.

Each user may specify their own graphics embodiment via a configuration file. In addition, we provide some default embodiments which are intended to convey the communication capabilities of the users they represent (an important issue in a heterogeneous environment). Thus, an audio user has ears, a non-immersive (and hence monoscopic) user has a single eye and a text user has the letter "T" embossed on their head. The aim of such embodiments is to provide other users with the necessary basic communication cues to decide how to address them. The basic shape of embodiments is also intended to convey general orientation in a simple and efficient manner. Text embodiments consist of a single character (the first letter of the person's chosen name) along with a short line which indicates the direction in which this person is currently facing.




Figure 2 A typical view of a graphics user.

Figure 2 shows a typical view of a graphics user. The figure shows five people sharing a space (we are one of them). Two graphics users with audio capability (note ears and single eye) are using the conference table adapter and two "texties" (note letter T) are engaged in a side conversation. Figure 3 shows a typical view of a text user. Note the use of simple characters to represent the conference table and walls and also the display of all currently known objects along with mutual awareness levels in the right hand column.



Figure 4 shows an example of spatial trading in operation. We see two client objects which support interfaces for media (M, N) and (M) respectively (i.e. they only speak medium M in common). On entering a world the left hand client object contacts the master trader for this world and declares the interfaces (M, N) as service offers. The master trader then returns interface references (addresses) of the two aura collision managers which are responsible for interfaces M and N. If necessary, the master trader spawns new aura collision managers (i.e. the first time that a given type of interface appears in this world). The second object subsequently entering the world will go through the same process of declaring its interfaces (M) to the master trader and being put in contact with the relevant collision managers. Each of the two aura collision managers monitors the auras of its known objects using the aura protocol. Upon detecting an aura collision, the aura collision manager passes out mutual interface references to the objects involved, enabling them to establish a peer connection. In this case, the aura collision can only be in medium M and so the two client objects establish a peer connection in medium M. Notice how the concept of spatial trading meets our goals of heterogeneity and scalability. Heterogeneity is realised through the master trader effectively registering all interface types currently available in a given world and maintaining a corresponding set of aura collision managers. This enables MASSIVE to dynamically broker hitherto unseen types of interface. Scalability is supported by distributing the responsibility for detecting aura collisions between multiple aura collision managers.

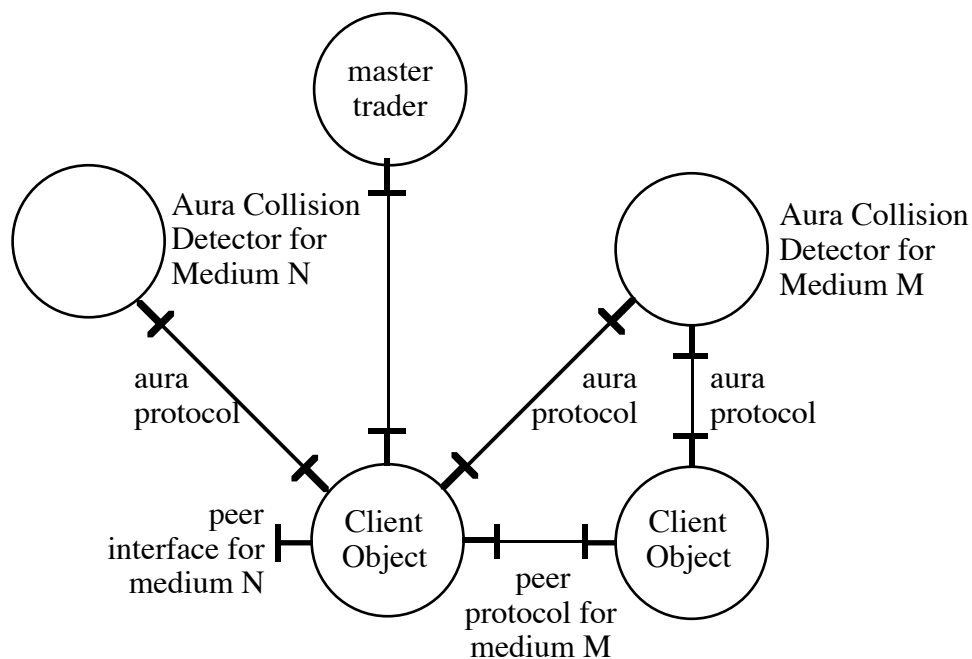


Figure 4 Spatial trader architecture

Once connected, the calculation of mutual awareness levels using focus and nimbus is the responsibility of the peer objects. In their most general sense, focus and nimbus are described by mathematical functions which map attributes of the communicating objects into focus and nimbus values. These are then combined into mutual awareness levels through a further awareness function. In the current implementation, objects are described

solely by a point location in space, and focus and nimbus are described by scalar fields over space, which are sampled at the location of a peer object to generate focus and nimbus values. Our current awareness function is multiplicative (i.e. focus and nimbus are multiplied to give awareness). This gives equal control to the observer and the observed, and is subtractive in nature - i.e. either party can force zero (no) awareness, but neither party can force awareness against the others wishes. Current focus and nimbus functions describe conical volumes projecting forwards from an object with some additional transition volume where focus or nimbus drops off to a specified background level. The operation of focus and nimbus may therefore be controlled by a few key parameters such as conical angle, maximum radius, background level and transition rate, allowing variation from generally spherical to tightly focused regions.

The implementation of adapters has proved to be a difficult problem and MASSIVE currently offers only a limited solution. There are two issues to be dealt with when implementing adapters: how to trigger the use of an adapter and how to realise its effect on focus, nimbus and awareness. Both of these are addressed though the introduction of a separate adapter medium. Adapters exist in their own medium, complete with aura, focus and nimbus. Any object wishing to use an adapter must therefore support an interface for this medium. As an object moves about it may connect to an adapter as a result of an aura collision in the adapter medium. When the object's awareness of the adapter subsequently crosses some threshold the adapter is triggered. The implementation of adapters also exploits the ability to parameterise focus and nimbus functions. When triggered, an adapter passes some new parameters back to the object across the adapter medium. These new parameters replace the object's current parameters. Thus, an adapter may extend the range of focus and nimbus, may change their shape (i.e. conical angle) or may alter the way in which they fade to a background level. When the object moves away from the adapter the object restores its normal parameters. This is handled automatically by the spatial model library. So as objects move about they automatically experience the effects of local adapter objects.

We are currently running MASSIVE on a network of SUN and SGI workstations connected to a standard Ethernet running TCP. Up to now we have managed sessions of up to six simultaneous users interacting over a mixture of graphics, audio and text (this is by no means large-scale, although it is on the high-side for current networked VR systems).

## **6. Initial observations and future work**

We conclude with some initial observations on the spatial model and its implementations gleaned from early use in the laboratory setting. We divide our observations into two categories: usability and conceptual issues.

Our first observation is a positive one - the combination of real-time graphics and audio is very natural. Indeed, it actually becomes useful and enjoyable. Our main conclusion here is don't attempt serious VR cooperation without decent audio! On the other hand, our current limited embodiments have caused problems. Although we can distinguish between different participants and can even guess at their capabilities from their bodies, determining their identity has been problematic.

In the MASSIVE case, a more surprising observation concerns inter-working between 3-D graphics users and 2-D text users. Although they are mutually visible within a

common space, their conception of that space seems quite different. In particular, the 'texties' seem to lack any notion of personal space and tend to stand directly in front of others or even walk straight through them. In contrast, graphics users tend to maintain a reasonable distance. The problem may be that the graphics field of view is much more limited than the textual one (360 degrees) and so the graphics users are forced to stand back in order to obtain a decent view. On the other hand, it may be that the graphics view is sufficiently rich that people can more easily associate the embodiments they see with other people and so feel the urge to behave in a socially polite manner. Either way, this observation would appear to point towards some deeper issues concerning interaction between users with radically different interfaces in a common environment.

From a conceptual point of view, work needs to be done on combining adapters and on representing adapters which are not obviously connected to specific objects (e.g. boundaries). We also need to explore the relationship between the users' foci/nimbi and those of adapters (e.g. merging, over-riding etc.). It is hard to create foci and nimbi which give power to the observer and the observed and which respond well to adapters (e.g. having a quiet chat but still being aware of the podium person when they step onto the podium). There is a conflict between allowing items different size auras (e.g. making them more or less demanding in terms of communications and processing) and how the world is expected to be (e.g. consistent results for the same position/distance).

## **7. Summary**

We have presented a spatial model of interaction which aims to support people's natural social skills in co-operative virtual environments. The key concepts defining the model include aura, awareness, focus, nimbus, adapters and boundaries and these may be applied independently to different media within a virtual world. Simplifications to the model, making it easier to implement, have been discussed. We then described a full realisation of the spatial model on a specialised distributed architecture and also a more lightweight approach to its implementation as an application of other VR platforms. Finally, we have presented some initial observations from use in the laboratory setting. We believe support for co-operative working to be an important issue for VR and hope that our model can help provide more flexible co-operative working environments.

## References

Benford, S.D, and Fahlén, L.E., [1993], "A Spatial Model of Interaction in Large Virtual Environments", *Proc. Third European Conference on CSCW (ECSCW'93)*, Milano, Italy, Kluwer.

Carlsson, C. and Olof Hagsand, O. [1993] "DIVE — a Multi-User Virtual Reality System" *In Proceedings of VRAIS 1993*, Seattle, Washington.

Codella, C., Lawrence R., Koved, J., Lewis J.B., Ling, D.T., Lipscomb, J.S., Rabenhorst, F.A., Wang, C.P., Norton, N., Sweeney, P., and Turk, G. [1992] "Interactive Simulation in a Multi-Person Virtual World", *Proc. CHI'92*, ACM Press.

Fahlén, L. E., Brown C. G., Ståhl, O., Carlsson, C., [1993], "A Space Based Model for User Interaction in Shared Synthetic Environments", *Proc. InterCHI'93*, Amsterdam, ACM Press.

Giddens, A., [1984], *The constitution of society*. Cambridge: Polity.

Grimsdale, C., [1991], "dVS - Distributed Virtual Environment System", *In Proc. Computer Graphics '91*, London, ISBN 0 86353 282 9.

Heath, C. and Luff, P., [1991], "Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms", In L. Bannon, M. Robinson and K. Schmidt (eds), *Proc. ECSCW'91*, Kluwer.

Hughes, J.A., Randall, D., and Shapiro, D., [1992], "Faltering from Ethnography to Design", *Proc. CSCW'92*, ACM Press, Toronto, Canada.

Takemura, H., and Kishino, F., [1992], "Cooperative Work Environment Using Virtual Workspace", *In Proc. CSCW'92*, Toronto, ACM Press.

West, A.J., Howard, T.L.J., Hubbard, R.J., Murta, A.D., Snowdon, D. N. and Butler, A. D., [1993], "AVIARY - A Generic Virtual Reality Interface for Real Applications", in *Virtual Reality Systems*, eds. E. A. Earnshaw, M.A Gigante and H. Jones, Academic Press Ltd. London.

Green, M., and White, L., [1994], "Minimal Reality Toolkit Version 1.3 Programers Manual", Department of Computer Science, University of Alberta, Edmonton, Alberta.

Zyda, M.J., Pratt, D.R., Monahan, J.G., and Wilson, K.P., [1992], NPSNET: constructing a 3D virtual world, *Computer Graphics*, 3, 147.

Bricken, W., and Coco, G., [1993], "The VEOS Project", technical report from the Human Interface Technology Laboratory, University of Washington FJ-15, Seattle 98195.