

A Spatial Model of Interaction in Large Virtual Environments

Steve Benford

The University of Nottingham, UK.

Lennart Fahlén

The Swedish Institute of Computer Science (SICS), Sweden.

Abstract: We present a spatial model of group interaction in virtual environments. The model aims to provide flexible and natural support for managing conversations among large groups gathered in virtual space. However, it can also be used to control more general interactions among other kinds of objects inhabiting such spaces. The model defines the key abstractions of object aura, nimbus, focus and adapters to control mutual levels of awareness. Furthermore, these are defined in a sufficiently general way so as to apply to any CSCW system where a spatial metric can be identified - i.e. a way of measuring position and direction. Several examples are discussed, including virtual reality and text conferencing applications. Finally, the paper provides a more formal computational architecture for the spatial model by relating it to the object oriented modelling approach for distributed systems.

1. Introduction

Our paper presents a model for supporting group interaction in large-scale virtual worlds¹. The model provides generic techniques for managing interactions between various objects in such environments including humans and computer artefacts. Furthermore, the model is intended to be sufficiently flexible to apply to any system

¹This work is part of the COMIC project, a European ESPRIT Basic Research Action to develop theories and techniques to support the development of future large scale CSCW systems.

where a spatial metric can be identified (i.e. a way of measuring distance and orientation). Such applications might range from the obvious example of multi-user virtual reality through conferencing systems, collaborative hypermedia and even databases and information spaces. Where the interacting objects are humans, the model provides mechanisms for conversation management. These contrast with existing floor control and workflow modelling techniques by adopting a "spatial" approach where people employ the affordances of virtual computer space as a means of control. In so doing, our underlying philosophy has been to encourage individual autonomy of action, freedom to communicate and minimal hard-wired computer constraints. Where the interacting objects are artefacts, the model provides mechanisms for constructing highly reactive environments where objects dynamically react to the presence of others (e.g. you may activate a tool simply by approaching it).

2. Rooms and virtual spaces

We have chosen to base our work around the metaphor of interaction within virtual worlds. Under this metaphor, a computer system can be viewed as a set of spaces through which people move, interacting with each other and with various objects which they find there. The use of such spatial metaphors to structure work environments is not particularly new, having previously been explored in areas such as user interface design, virtual meeting rooms, media spaces, CSCW environments and virtual reality. Xerox used a rooms metaphor to structure graphical interfaces (Henderson 85, Clarkson 91) and this was later followed up with VROOMS (Borning 91). Audio Windows applied a spatial metaphor to audio interfaces (Cohen 91). Multi-media virtual meeting rooms have been demonstrated in a variety of projects (Leevers 92, Cooke 91). The CRUISER system explored social browsing in larger scale virtual environments (Root 88) and multi-user recreational environments have been available for some time (e.g. MUD (Smith 92) and Lucasfilm's HABITAT (Morningstar 91)). Spatial metaphors also feature heavily in discussions of Virtual Reality (VR) (Benedikt 91) including early collaborative VR systems (Codella 92, Takemura 92, Fahlén 92). In contrast to virtual reality, media-spaces explore the role of space in providing more embedded support for cooperative work (Gaver 92a, Gaver 92b). Finally, spatial metaphors have been adopted as an integrating theme for large scale CSCW environments (Michelitsch 91, Navarro 92).

In short, spatial approaches to collaborative systems have become increasingly popular. One reason for this is their strong relation to physical reality and therefore their highly intuitive nature. However, from a more abstract standpoint, space affords a number of important facilities for collaboration including awareness at a

glance; support for ad-hoc as well as planned interaction; use of body language and other social conventions in conversation management; flexible negotiation of access to resources (e.g. queuing, scrumming and hovering), and structuring, navigation, exploration and mapping of large-scale work environments.

However, we believe that current spatially-oriented systems will not effectively scale to heavily populated spaces. More specifically, as the number of occupants in a virtual space increases beyond a few people, the need to effectively manage interactions will become critical. One example, is the need for conversation management. As a starting point, we might consider borrowing the conversation management and coordination mechanisms developed in other areas of CSCW. Previous conferencing systems have introduced a range of floor control mechanisms such as chairpeople, reservations and token-passing (Crowley 90, Sarin 91, Cook 92). Alternatively, the work-flow and process oriented techniques from asynchronous systems also represent a form of conversation management (e.g. THE COORDINATOR (Winograd 86), DOMINO (Victor 91), CHAOS (Bignoli 91), COSMOS (Bowers 88) and AMIGO (Pankoke 89)). However, we believe that these approaches are generally too rigid and unnatural to be applied to spatial settings. As an example, a real-world implementation of explicit floor control would be tantamount to gagging everyone at a meeting and then allowing them to speak by removing the gags at specific times. New techniques are needed which support natural social conventions for managing interactions. One approach might be to take advantage of the highly fluid and dynamic nature of space. The following section introduces a *spatial model* of interaction which aims to meet these goals. Furthermore, although we base our discussion on a consideration of three dimensional space, the model is intended to be sufficiently generic to apply to any system where a spatial metric can be identified, including possible higher dimensional *information terrains*.

3. The spatial model

Virtual spaces can be created in any system in which position and direction, and hence distance, can be measured. Virtual spaces might have any number of dimensions. For the purposes of discussion we will consider three. The objects inhabiting virtual spaces might represent people and also other artefacts (e.g. tools and documents). Our model has been driven by a number of objectives including ensuring individual autonomy; maintaining a power balance between “speakers” and “listeners” in any conversation; minimising hard-wired constraints and replacing them with a model of increasing effort; and starting with support for free mingling and only adding more formal mechanisms later if needed.

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 **MEDIUM**, **AURA**, **AWARENESS**, **FOCUS**, **NIMBUS** and **ADAPTERS** which define our model.

Any interaction between objects occurs through some **medium**. A medium might represent a typical communication medium (e.g. audio, visual or text) 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.

The first problem in any large-scale environment is determining which objects are capable of interacting with which others at a given time (simultaneous interaction between all objects is not computationally scaleable). **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 (Fahlén 92). 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. Note that an object typically has different auras (e.g. size and shape) for different media. For example, 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. Also note that it is the surrounding environment that monitors for aura collisions between objects.

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 (Benford 92). The measure of awareness between two objects need not be mutually symmetrical. 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, the more an object is within your focus, the more aware you are of it and the more an object is within your nimbus, the more aware it is of you. Objects therefore 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,

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

The resulting quantified awareness levels between two objects can then be used as the basis for managing their interaction. Exactly how this is achieved depends upon the 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 allowing objects to actively 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). Notice that the use of both focus and nimbus allows both objects in an interaction to influence their awareness of each other. More specifically, they support our stated goals of autonomy and also power balance between "speakers" and "listeners".

Now we consider how much of this apparent complexity the user needs to understand. The answer is very little, because a person need not be explicitly aware that they are using aura, focus and nimbus. First, aura, focus and nimbus may often be invisible or may be implied through "natural" mechanisms such as the use of eyes to provide gaze awareness and hence convey visual focus. Second they will be manipulated in natural ways which are associated with basic human actions in space. To be more specific, we envisage three primary ways of manipulating aura, focus and nimbus and hence controlling interaction:

1. Implicitly through movement and orientation. Thus, as I move or turn, my aura, focus and nimbus will automatically follow me. A number of novel interface devices are emerging to support this kind of movement. These are generally known as six dimensional devices (three for position and three for orientation) and include space-balls, body-trackers, wands and gloves.
2. Explicitly through a few key parameters. A user interface might provide a few simple parameters to change aura, focus and nimbus. I might change the shape of a focus by focusing in or out (i.e. changing a focal length). This might be achieved by simply moving a mouse or joystick.
3. Implicitly by using various **adapter** objects which modify my aura, focus or nimbus. These can be represented in terms of natural metaphors such as picking up a tool.

Adapters support interaction styles beyond basic mingling. 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 space. In effect, the introduction of adapter objects provides for a more extensible model.

To summarise, our spatial model defines key concepts for allowing objects to establish and subsequently control interactions. Aura is used to establish the potential for interaction across a given medium. Focus and nimbus are then used to negotiate the mutual and possibly non-symmetrical levels of awareness between two objects which in turn drives the behaviour of the interactions. Finally, adapter objects can be used to further influence aura, focus and nimbus and so add a degree of extendibility to the model.

4. Applying the spatial model

The spatial model is intended to be applicable to any system where a spatial metric can be identified. We now briefly describe some example applications of the spatial model including the multi-user virtual reality and text conferencing systems currently being prototyped at SICS and Nottingham respectively.

4.1. Multi-user virtual reality - the DIVE system

Perhaps the most obvious application of the spatial model is to virtual reality systems. A prototype multi-user Virtual Reality (VR) system, DIVE (Distributed Interactive Virtual Environment) (Fahlén 91) (Carlsson 92) has been developed as part of the MultiG program (a Swedish national research effort on high speed networks and distributed applications (Pehrson 92)). DIVE is a UNIX-based, multi-platform software framework for creating multi-user, multi-application, three-dimensional distributed user environments. There is support for multiple co-existing "worlds" with gateways between them to enable inter-world movement. Users are represented by unique graphical 3D-bodies or icons whose position, orientation, movements and identity are easily visible to other participants. In this first realisation, aura is implemented as a volume or sphere around each user's icon which is usually invisible. Aura handling is achieved through a special collision manager process. When a collision between auras occurs, this manager sends a message containing information such as the id's of the objects involved, positions, angles and so on, to other processes within the DIVE environment. These processes (e.g. the owners of the objects involved) then carry out appropriate focus, nimbus and awareness computations. It is possible to have support for a multiple users, objects, media and service specific aura types with associated collision managers mapped onto separate processing nodes in a network. Focus and nimbus handling can be mapped in a similar way. Further details on the aura implementation in DIVE can be found in (Ståhl 92b). Figure 1 shows a screen dump from DIVE of an aura collision, with the auras made specially visible.




Figure 1: Body Images with Colliding Auras

A more general toolkit has been developed as a first step towards constructing a distributed collaborative environment and for experimentation with the concepts of aura, focus, nimbus and awareness. Presently it consists of four major components, the whiteboard, the document, the conference table and the podium.

The whiteboard (Stähl 92a) is a drawing tool similar in appearance to its real world counterpart. Several users can work together simultaneously around the whiteboard. There can also be groups of whiteboards, with the contents being duplicated across the group. That is, the actions performed by one user on one whiteboard are immediately replicated by the other whiteboards in the same group. The aura surrounding the whiteboard is used to enable whiteboard access and use (e.g. by automatically assigning a pen to a user when their aura collides with that of the whiteboard). The content of a whiteboard can be copied into something called a document that a user can pick up and carry away. Apart from being "single user", documents have the same functionality as a whiteboard. More specifically, when document auras intersect, their contents are copied to other users documents and onto whiteboards.

The conference table detects participants presence, and establishes communication channels (video, voice and document links) between them via aura. The auras, foci and nimbi of the conference participants around the table are then extended to cover everyone in attendance. So, by having a very long table, people can form larger collaborative groups than “direct” aura/focus/nimbus functionality makes possible. Users can come and go as they please and it is easy to obtain an overview of who is present. The conference table can also distribute documents to conference participants and to whiteboards. To do this a user simply places a document in the centre of the table and then the aura collision manager initiates the distribution. Figure 2 shows a screen dump of a meeting in Cyberspace involving the whiteboard and conference table.




Figure 2: A Conference in Cyberspace

A participant can enter a podium and is thereby allowed to be “heard” (or seen) by a dispersed group of users that “normally” (e.g. without the podium) are not within communication distance of the “speaker”. The aura and nimbus of the participant on the podium are enlarged to cover, for example, a lecture hall or town square. The podium is an example of an aura/nimbus adapter and it is asymmetric, i.e. the “listeners” can’t individually communicate back to the “speaker” without special provisions.

A teleconferencing subsystem is also under construction and will be integrated into DIVE in the near future (Eriksson 92). Apart from the CSCW toolkit, some other concept demonstrators have also been developed within the DIVE environment, including control of a real-world robot, a 3D customisable graph editor for drawing and editing graphs in 3D space, a 3D-sound renderer allowing objects or events to have sounds and for these sounds to have a position and direction and finally a computer network visualiser and surveillance tool.

4.2. Text conferencing - the CyCo system

We can also apply the spatial model to less sophisticated technology. For example, several text conferencing systems have been produced over recent years. Such systems support communication through the medium of text messages and often introduce a floor control mechanism for managing conversations in groups of more than a few people. Consider, instead, the application of the spatial model to such a system. We might define rooms to be two dimensional spaces which could be readily mapped in a window on a typical workstation screen. Aura might be circular in shape and focus and nimbus might be modelled as segments of a circle projecting from a person's current position that could be manipulated both by moving and turning. In the simplest case, these areas might provide for discrete values of focus and nimbus (i.e. if an object is inside the area, then it is in focus/nimbus; if it is outside, then it is not). Considering two people, A and B, we can now evaluate three possible levels of awareness (see figure 3) :-

- A is fully aware of B if B is within A's focus and A is within B's nimbus. In this case, A would receive text messages from B.
- A is not aware of B if B is not within A's focus and A is not within B's nimbus. In this case A sees no messages from B.
- A is semi-aware from B if either B is within A's focus or A is within B's nimbus, but not both. In this case A wouldn't receive messages from B, but would be notified that B was speaking near by.

Even with such a relatively crude application of the spatial model (i.e. using simple discrete valued foci/nimbi), some interesting and novel effects come into play. In particular, there is a semi-aware state in which I am notified that you are trying to speak (perhaps in a separate window) without hearing what you say. Notice also that there is a power balance between A and B in terms of their abilities to influence the conversation and also that their levels of awareness may be asymmetrical. A prototype application of the spatial model to a text conferencing system is being realised in the *CyCo (Cyberspace for Cooperation)* system at Nottingham University (Benford 92). CyCo provides a large environment of connected virtual rooms and is implemented on top of the ANSA Distributed

Processing Platform (ANSA 89). The current prototype supports two user interfaces, an X Windows interface using the Motif widget set and a teletype interface based on the UNIX Curses C library. CyCo can be configured to support specific world designs by creating new room descriptions and topology information and also provides inbuilt mapping facilities to aid navigation.




Figure 3 - Levels of Awareness in Text Conferencing.

4.3 Other applications

We can also envisage the application of the spatial model to a range of other CSCW systems. One interesting example might be that of collaborative hypermedia. A hypermedia document can be considered as a one dimensional space where the spatial metric is the number of links between two nodes. Simple aura, focus and nimbus might then convey a sense of awareness between people browsing through such a space. Hypermedia browsers could use measures of awareness to take actions such as notifying people of the presence of others or automatically opening up communication channels.

To go a stage further, it may be possible to spatially organise more general information domains, classification schemes and taxonomies. One approach to the

spatial visualisation of large databases is given in (Mariani 92). As a second example, work has been carried out into the spatial mapping and classification of scientific disciplines based on a statistical analysis of the co-occurrence of keywords in academic papers. More specifically, the analysis resulted in measures of *inclusion* and *proximity* between keywords and these were used to automatically draw maps of scientific areas (Callon). The spatial model could be applied to manage interactions across such a space. Similar techniques might have applications in areas such news systems, bulletin boards and shared databases. Perhaps in the future, we will see collaboration taking place across large *populated information terrains* of spatially arranged data.

5. Distributed support for the spatial model

This section outlines a more formal computational framework for the spatial model by relating it to current object-based approaches for building distributed systems. This process also highlights a number of key requirements for future distributed systems support for collaborative virtual environments.

5.1. Object-based models of distributed systems

Much effort has been invested in to the development of platforms for building large scale distributed systems including the Open Distributed Processing framework (ODP) (ISO 91a); the work of the Object Management Group (OMG) (OMG 90); OSF's Distributed Computing Environment (DCE) (OSF 92) and systems such as ISIS (Birman 91). Although not identical, these emerging platforms share much in common; particularly the use of an object-based modelling approach. The following discussion uses terminology from the ODP work. However, the underlying principles are generally applicable to other emerging platforms.

A distributed system can be modelled as a set of *objects* which interact through well defined *interfaces*. An interface groups together a set of related *operations* which are invoked by one object on another. A distributed platform provides some mechanism for establishing contact between objects, negotiating the use of interfaces and invoking operations. In the Open Distributed Processing model, this is supported through the process of *trading*, probably one of the most important concepts to emerge from distributed systems work in recent years (ISO 91b). In order to trade, a provider object exports its interfaces by registering them with a well known system object called the trader. The trader notes the *type* of each interface and also the *context* in which it is provided (effectively the name of the service provided). A consumer object that wishes to use an interface queries the trader, supplying both the desired interface type and also target contexts. The trader

looks for a match and, if one exists, returns an *interface reference* to the consumer. This interface reference can then be used by the consumer to invoke operations on the provider. Notice that, in current trading models, it is the consumer who decides when to request an interface reference from the trader and that, in effect, the trader is a passive service. The main advantage of trading is that it provides a high degree of transparency for object interactions. The concepts of objects, interface, operations and trading are summarised by figure 4. Other distributed platforms define similar mechanisms (e.g. the Common Object Request Broker (CORBA) in the OMG work).

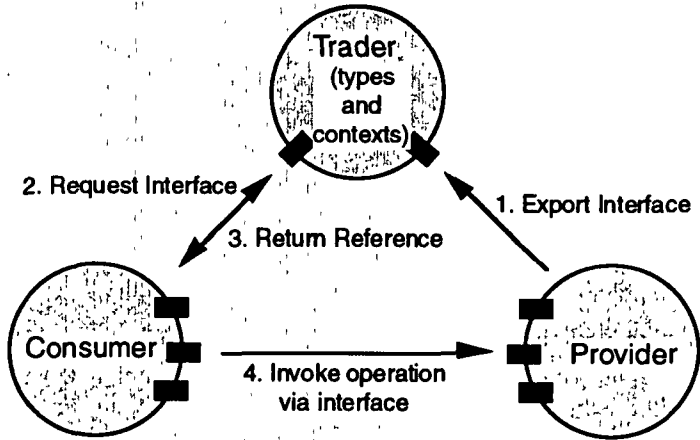


Figure 4: Trading

5.2 Requirements of trading in virtual environments

We expect that collaborative virtual environments will be characterised by a number of features which will impact on the nature of object interactions and on fundamental ideas such as trading. First, they will include objects which represent human beings. Human beings are intelligent and autonomous, often liking to explore their environments. Interaction between objects will therefore often be ad-hoc and opportunistic. Objects will not always know in advance which interfaces they require and so the passive trading model will not be sufficient. Instead, objects will require the trader to actively inform them of new services that become available as they move about (i.e. services that come into range). Second, in addition to interface type and context, trading will be based on the spatial proximity of objects. In other words, as objects get closer to each other, they will become more aware of each other and will be able to invoke new operations on each other. In this way the environment becomes more reactive (i.e. objects react to each other's presence). A

good example might be moving up to a bulletin board. At a great distance you don't see it. Closer to, you see it is there. Even closer and you can read messages. Even closer and you can write on the board. In summary, trading in collaborative virtual worlds will be active as well as passive and will be based on a notion of spatial proximity, and hence awareness, between objects. Finally, large distributed systems will contain many traders, each of which is responsible for a specific set of objects. In this case, we say that each trader manages a local *trading domain*. Furthermore, traders may federate together in order to exchange information about trading domains and so achieve a distributed trading service.

5.3 Extending object interfaces and trading

Next, we outline key extensions to the distributed object model to support the spatial model. At the same time, this provides a more computational and general definition of the spatial model itself. First we consider a general mapping of terms. People and artefacts are represented as objects. Communication media are mapped onto different interfaces (e.g. "audio" or "text") allowing interaction between these objects. A single virtual space containing many objects maps onto a trading domain managed by a given trader. Now we can introduce the idea of managing object interactions through inter-object awareness. We can associate an aura with each interface. When two auras collide, the relevant interfaces are enabled - in other words, the objects mutually acquire interface references. It is the role of the trader to detect aura collisions and to actively pass out interface references. Next, we associate focus and nimbus with an interface. This time it is the objects themselves, not the trader, that negotiate awareness levels. These levels can then be used in two ways. Operations within an interface can be associated with an awareness threshold at which they become available to other objects. Also, objects can decide to invoke operations on others once certain thresholds are passed. This ability for objects to determine levels of mutual awareness requires support from standard operations to return values of focus and nimbus from a given interface. Notice that, in terms of where computation takes place, the trader is concerned with supporting aura whereas the objects themselves deal with the use of focus and nimbus. These key extensions of aura, focus and nimbus in object interfaces are shown in figure 5. We also need to consider how aura, focus and nimbus are formally represented and computed. Given that we require a quantitative measure of awareness, we can model them as mathematical functions which map from spatial properties of objects such as position and orientation into real number values. This is similar to the way in which functions can be used to describe properties of surfaces in surface modelling. We then combine values of aura, focus and nimbus through a separate awareness function. A more detailed mathematical treatment of focus and nimbus is given in (Benford 92).

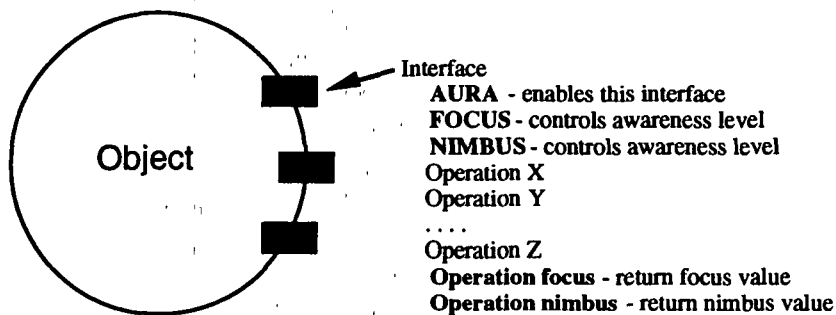


Figure 5: Spatial model extensions to object interfaces

As a final comment, by considering general object interfaces, we need not only use awareness to control conversation across communication media; it can also be used to govern any kind of interaction between objects in distributed systems. Thus, the spatial model might eventually provide a more generic platform for building a variety of virtual environments.

6. Summary

Our paper has described a spatial model of group interaction in large-scale virtual environments. The model provides mechanisms for managing conversations between people, as well as interactions with other kinds of objects, in spatial settings. The notion of awareness is used as the basis for controlling interaction and the model provides mechanisms for calculating awareness levels from the spatial properties of objects (e.g. position and orientation). This allows objects to manage interactions through natural mechanisms such as movement and orientation in space. The model defines the key concepts of aura, focus, nimbus and adapter objects all of which contribute to awareness. Furthermore, these concepts are defined in a sufficiently general way so as to apply to any system where a spatial metric can be identified. The paper then considered several example applications including virtual reality and text conferencing, both of which are currently being prototyped. Finally, we outlined a more computational definition of the spatial model by relating it to recent work on distributed systems; in particular, to the notions of objects, interfaces, operations and trading. Much work remains to be done. The current prototypes require extension and eventually proper evaluation. Additional applications also need to be modelled and demonstrated. However, at this stage, we are optimistic that spatial models of interaction such as the one described in this paper, will form an important aspect of support for CSCW, particularly as new technologies such as virtual reality become more widespread in the next few years.

References

- (ANSA 89) *The ANSA Reference Manual*, Architecture Projects Management Ltd, Poseiden House, Castle Park, Cambridge CB3 PRD, UK, 1989.
- (Benedikt 91) Michael Benedikt, *Cyberspace: Some Proposals*, in *Cyberspace: First Steps*, Michael Benedikt (ed), MIT Press, 1991, pp 273-302.
- (Benford 92) Steve Benford, Adrian Bullock, Neil Cook, Paul Harvey, Rob Ingram and Ok Ki Lee, *From Rooms to Cyberspace: Models of Interaction in Large Virtual Computer Spaces*, The University of Nottingham, Nottingham, UK (to appear in the Butterworth-Heinmann journal, *Interacting With Computers* in 1993).
- (Benford 92) Steve Benford, Wolfgang Prinz, John Mariani, Tom Rodden, Leandro Navarro, Elsa Bignoli, Charles Grant Brown and Torbjørn Naslund, *MOCCA - A Distributed Environment For Collaboration*, Available from the MOCCA Working Group of Co-Tech.
- (Borning 91) A. Borning and M. Travers, *Two Approaches to Casual Interaction over Computer and Video Networks*, In Proc. CHI'91, New Orleans, April 27-May 2, 1991, pp13-19.
- (Bignoli 91) C. Bignoli and C. Simone, *AI Techniques to Support Human to Human Communication in CHAOS*, in *Studies in Computer Supported Cooperative Work: Theory, Practice and Design*, J. Bowers and S. Benford (eds), Elsevier Science Publishers, 1991.
- (Bowers 88) J. Bowers and J Churcher, *Local and Global Structuring of Computer Mediated Communication: Developing Linguistic Perspectives on CSCW in COSMOS*, In Proc. CSCW'88, Portland, Oregon, September 26-28, 1988.
- (Callon) Michel Callon, John Law and Arie Rip (eds), *Mapping the Dynamics of Science and Technology*, Macmillan press, ISBN 0-333-37223-9.
- (Carlsson 92) Carlsson, C. and Hagsand, O. *The MultiG Distributed Interactive Virtual Environment*, In Proc. 5th MultiG Workshop, Stockholm, December 18, 1992.
- (Clarkson 91) Mark A. Clarkson, *An Easier Interface*, BYTE, February 1991, pp227-282.
- (Codella 92) Christopher Codella, Reza Jalili, Lawrence Koved, J. Bryan Lewis, Daniel T. Ling, James S. Lipscomb, Favid A. Rabenhorst, Chu P. Wang, Alan Norton, Paula Sweeney, Greg Turk, *Interactive Simulation in a Multi-Person Virtual World*, In Proc. CHI'92, ACM, 1992.
- (Cohen 91) Michael Cohen and Lester F. Ludwig, *Multidimensional Audio Window Management*, in *Computer Supported Cooperative Work and Groupware*, Saul Greenberg (ed), Harcourt Brace Jovanovich, 1991, ISBN 0-12-299220-2.
- (Cook 91) Cook, S., Birch, G., Murphy, A., and Woolsey, J., *Modelling Groupware in the Electronic Office*, in *Computer-supported Cooperative Work and Groupware*, Saul Greenberg (ed), Harcourt Brace Jovanovich, 1991, ISBN 0-12-299220-2.
- (Cook 92) Neil Cook and Graeme Lunt, *XT-Confer: Dynamic Desktop Conferencing*, In Proc. European X User Group Fourth Annual Conference, September, 1992
- (Crowley 90) Crowley, T., Milazzo, P., Baker, E., Forsdick, H. and Tomlinson, R., *MM Conf: An Infrastructure for Building Shared Multi-media Applications*, In Proc. CSCW'90, October 1990, ACM Press.
- (Eriksson 92) Eriksson H., Frecon E., Carlsson C., *Audio and Video Communication in Distributed Virtual Environments*, In Proc. 5th MultiG Workshop, Stockholm, Dec. 1992.
- (Fahlén 91) Fahlén, L. E. *The MultiG TelePresence System*, In Proc. 3rd MultiG Workshop, Stockholm, December 1991, pp. 33-57.
- (Fahlén 92) Fahlén, L. E. and Brown, C.G., *The Use of a 3D Aura Metaphor for Computer Based Conferencing and Teleworking*, In Proc. 4th Multi-G Workshop, Stockholm-Kista, May 1992, pp 69-74.
- (Fahlén 93) Lennart E. Fahlén, Charles Grant Brown, Olov Stahl, Christer Carlsson, *A Space Based Model for User Interaction in Shared Synthetic Environments*, The Swedish Institute of Computer Science (SICS) (to appear in Proc. InterCHI'93).
- (Gaver 92a) Gaver W., *The Affordances of Media Spaces for Collaboration*, In Proc. CSCW'92, Toronto, November 1992, ACM Press.

- (Gaver 92b) Gaver, W., Moran, T., MacLean, A., Lovstrand, L., Dourish, P., Carter, K. and Buxton W., *Realising a Video Environment: EuroPARC's RAVE System*, In Proc. CHI '92 Human Factors in Computing Systems, Monterey, Ca., USA, May 1992.
- (Henderson 85) Henderson and Card, *Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention*, ACM Transactions on Graphics, Vol. 5, No. 3, July 1985.
- (ISO 91a) ISO/IEC, *Basic Reference Model of Open Distributed Processing, Working Document RM-ODP - Part 1: Overview, December 1991*, Available through national standards bodies.
- (ISO 91b) ISO/IEC, *Basic Reference Model of Open Distributed Processing, Working Document on topic 9.1 - ODP Trader, December 1991*, Available through national standards bodies.
- (Leevers 92) David Leevers, *Prototyping Multimedia Communications for Manufacturing SME's*, Presented at the CI Europe Seminar, 2-3rd July 1992, and also an internal report of BICC Central Development, Maylands Avenue, Hemel Hempstead, UK.
- (Mariani 92) J. Mariani, Lougher, *TripleSpace: an Experiment in a 3D Graphical Interface to a Binary Relational Database*, Interacting with Computers, Vol 4, No. 2, 1992, pp147-162.
- (Michelitsch 91) Georg Michelitsch, *Providing a Shared Environment to Distributed Work Groups*, In Proc. GLOBECOM'91, Phoenix, Arizona, December 2-5th, 1991.
- (Morningstar 91) Chip Morningstar and F. Randall Farmer, *The Lessons of Lucasfilm's Habitat*, in *Cyberspace: First Steps*, Michael Benedikt (ed), MIT Press, 1991, pp 273-302.
- (Navarro 92) L. Navarro, M. Medina and T. Rodden, *Environment Support for Cooperative Working*, In Proc. IFIP 6.5 ULPA'92 Conference, Vancouver, Canada, May 1992, North-Holland.
- (OMG 90) The Object Management Group (OMG), *Object Management Architecture Guide*, OMG Document Number 90.5.4, Available through the Object Management Group, 1990.
- (OSF 92) OSF, *Distributed Computing Environment: An Overview*, Available from the OSF January 1992.
- (Pankoke 89) Uta Pankoke-Babatz (ed), *Computer Based Group Communication - The Amigo Activity Model*, Ellis-Horwood, 1989.
- (Pehrson 92) Pehrson, B., Gunningberg, P. and Pink, S. *MultiG-A research Programme on Distributed MultiMedia Applications and Gigabits Networks*, IEEE Network Magazine vol 6, 1 (January 1992), pp. 26-35.
- (Root 88) R.W. Root, *Design of a Multi-Media Vehicle for Social Browsing*, In Proc. CSCW'88, Portland, Oregon, Spetember 26-28 1988, pp25-38.
- (Sarin 91) Sunil Sarin and Irene Greif, *Computer-Based Real-Time Conferencing Systems*, in *Computer Supported Cooperative Work: A Book of Readings*, Irene Greiff (ed), Morgan Kaufmann, 1988, pp397-421.
- (Sarkar 92) Manojit Sarkar and Marc H. Brown, *Graphical Fisheye Views of Graphs*, In Proc. ACM SIGCHI'92 Conference on Human Factors in Computing Systems, May 3-7 1992, pp 83-91, ACM Press.
- (Smith 92) Jennifer Smith (ed), *Frequently Asked Questions: Basic Information about MUDs and MUDing*, Posting to the alt.mud USENET News Group, July 1992.
- (Ståhl 92a) Olov Ståhl, *Mdraw - A Tool for Cooperative Work in the MultiG TelePresence Environment*, Technical Report T92:05, SICS, 1992.
- (Ståhl 92b) Olov Ståhl, *Implementation Issues of Aura Based Tools*, In Proc. 5th MultiG Workshop, Stockholm, December 18, 1992.
- (Takemura 92) Haruo Takemura and Fumio Kishino, *Cooperative Work Environment Using Virtual Workspace*, In Proc. CSCW'92, Toronto, Nov 1992, ACM Press.
- (Victor 91) F. Victor and E. Sommer, *Supporting the Design of Office Procedures in the Domino System*, in *Studies in Computer Supported Cooperative Work: Theory, Practice and Design*, J. Bowers and S. Benford (eds), Elsevier Science Publishers, 1991.
- (Winograd 86) T. Winograd and F. Flores, *Understanding Computers and Cognition: A New Foundation for Design*, Norwood, New Jersey, Ablex 1986 and Addison-Wesley 1987.