WP4.3: Collaborative Interactive Techniques

Project Number:	ACTS Project N. AC040
Project Title:	COVEN - Collaborative virtual environments
Deliverable Type:	P*

CEC Deliverable Number:	A040-UCL-CS-DS-P-043.b1
Contractual Date of Delivery to the CEC:	1 October 1997
Actual Date of Delivery to the CEC:	6 October 1997
Title of Deliverable:	Collaborative Interaction Techniques
Workpackage contributing to the Deliverable:	WP4
Nature of the Deliverable:	P**
Author(s):	Martin Usoh (UCL) Emmanuel Frecon (SICS) Adrian Bullock (Nottingham)

Abstract:

Interactive collaborative virtual environments has become an area of research which has received recent attention. Developing appropriate methods for communication and collaboration between participants is therefore an important factor of this line of investigation. This document presents techniques for collaboration and interaction in these environments. We present and investigate the following: 1) collaboration and a better understanding of a large scale virtual environment through the use of virtual maps 2) the sharing of resources and access control. The work is constructed under the scenario of a business traveller who, using virtual reality acquaints themselves with a virtual city before initially visiting it in reality.

Keyword list: Collaboration, Interaction, Presence, Map Drawing, Access Control

***Type: P-public, R-restricted, L-limited, I-internal**

****Nature: P-Prototype, R-Report, S-Specification, T-Tool, O-Other**

1 Introduction

The surrounding three dimensional space of a virtual environment (VE) can permit interactions which are similar in nature to the real world in comparison to 2D metaphors where this is not the case. These interactions may be with virtual objects, or with the visual embodiment of other participants. Whether with virtual objects or visual embodiments, the interactions are important for greater realism, heightened belief and presence in the virtual environment [BARF95a, SLAT93a].

Given the collaborative nature of the COVEN proposal it is important to maximise across all modalities (ideally) the set of interactions which are present between participants. This may be achieved by a system which makes available the tools or techniques to develop these interactions. For example, a system may provide for full body tracking and the ability for visual embodiments to be solid therefore permitting hand shaking interactions. Such enhancements of the VE can encourage participants to behave in a manner similar to their behaviour in the real world under similar conditions. This can lead to an increased sense of presence which in turn may produce a tendency for greater interaction.

In this work we are concerned with presence and interaction within multiparticipant environments. We distinguish interaction paradigms which are "appropriate" to the environment and investigate their influence on presence. By "appropriate" we mean those interactions which have a real world bearing. For instance, we believe it is more appropriate to navigate through a virtual building by bodily movements (virtual walking) rather than "magical" flying. Conversely it is more appropriate to navigate over virtual buildings by flying rather then virtual walking.

As well as investigating the usability of different interaction methods we will investigate the relationship between interaction and presence in collaborative situations. This will be conducted under the scenario of a traveller lost in a virtual town who, through interaction with other participants, is able to find the way (University College London) . The scenario provides for a set of tools with behaviours which permit the dynamic creation of three dimensional maps. The maps may be interactively amended and explored by inhabiting its space.

We have also developed (at SICS) a Path Planning Map (PPM) which is based on the world in miniature metaphor. This provides participants with a hand held miniature representation of the world, somewhat like a map in the real world. The virtual map metaphor benefits from the notion of everyday maps, something most people are familiar with, and thus, a concept that can be grasped quickly. The PPM can be used for locating and orienting the user, searching for other users, and for moving directly to a point on the map.

Finally we present our work which develops a spatial approach to access control in collaborative virtual environments (University of Nottingham). Here we are interested in general techniques for managing individual and group access to resources located within a virtual environment, assuming that related security issues such as authenticating user's identities and auditing their actions have already been dealt with.

2 Aims

This work package is concerned with collaborative interactive techniques in the context of the business scenario - making a virtual visit to a place in preparation for the real visit. A primary aim of UCL’s work in this context is therefore to provide a set of facilities and tools to enable the virtual traveller to explore a large scale city environment with the purpose of retaining as much information from the experience. This can then be applied when visiting the real city environment.

Interactive techniques have been based around the concept of the 'lost traveller' in VE who encounters another participant (the helper) and proceeds to obtain directions. During such human encounters a number of modes of communication are employed, e.g. verbal, gestures, maps. Techniques adopted should either be reproducible in the real world or permit information transfer to the real world.

We propose to develop techniques which will enhance the information retained from the experience for later application to the real world, i.e. when visiting the same virtual place for real. In this effort we propose techniques which can increase the behavioural presence. These will be developed on top of an interactive map drawing system which will allow ideas to be communicated between participants. These ideas may be abstract, for example exploring the relationship between two concepts.

We have integrated the idea of body centred interaction [SLAT93b] for application to a worlds-within-worlds or depth of presence concept [SLAT94]. That is, after interactively creating a map outlining the route to the point of interest, the lost traveller may explore this map 'internally' by shrinking down to the scale of the map and navigating it from within. This will allow us to investigate a number of things including:-

- the interaction between the lost traveller and the helper in way finding;
- the perception of the lost traveller experiencing a world-within-worlds;
- the influence on the sense of presence on the lost traveller;
- the influence on retention of information gained in VE.

In the following sections (3 to 6) we will discuss the virtual body and aspects of interaction and presence pertinent to the goals of WP4.3. These are necessary in the development of an integrated system for collaborative interaction. Section 7 describes the interactive map drawing technique and methods of collaboration.

Further aims of the work package are to enhance communication and awareness in the VE. Section 8 describes the work conducted at the Swedish Institute of Computer Science on a Path Planning Map which may be used by the lost traveller to realise their direction. In addition, the goal of managing access control of shared resources is important in interactive collaborative environments and is pursued in Section 9 by the group at Nottingham University.

3 Presence: The Significance of Virtual Worlds

Virtual environments have advantages over other traditional forms of human-computer interaction in that they allow a person to perform tasks in a manner similar to everyday reality. Therefore no special learning of skills is required in order to carry out common place activities (such as picking something up) in such environments. To go through a door in the virtual world would be similar to the real world. This intuitive approach is a key factor to the use of immersive virtual environments (IVE).

While in the virtual environment the participant can experience psychological changes of state which are consistent with the environment delivered by the computer displays. For example, if the visual display of the computer presents a consistent view seen from a cliff edge to the visual system of the participant, it is possible that the participant will undergo similar changes of state as if actually on the real cliff edge - e.g. effects of vertigo [SLAT95]. Note, however, that the person must suspend belief in the physical world and accept the images conveyed by the displays.

A virtual environment can therefore induce the psychological state of "being there" in the environment displayed by the computer. We take the issue that this sense of *presence* is the central feature of virtual reality which makes it a unique form of man-machine interface. The sense of presence is a psychological state of consciousness. It is an emergent property of an IVE system which is affected by a number of factors [HEET92, LOOM92] such as the behavioural realism of the environment, the quality and update rate of the graphics [BARF95b], and the possession of a virtual body by the participant - i.e. a computer generated representation of the participant's physical body or self [SLAT95, SLAT93a]. The behaviour of this virtual body is important when considering collaboration where participants must be immediately identifiable through their individual body representation and movement.

4 The (Virtual) Body

An important component of everyday reality which gives us physical presence within the world is the body. Since sensory information from the outside world is received and processed through the body it plays the role of an anchor to the world generating the sensory information. In addition the body acts as three important mediums.

- It is a medium of interaction and its possession permits the ability to change the environment. Given current day technology it is also a tool for manipulation in near and remote sites;
- It is the physical embodiment of self. Through it we gain a personal and social identity of being;
- It is a medium of communication which allows interaction with others on a higher level. It is used to exchange feelings, ideas, and opinions through sounds and gestures.

In dealing with collaborative environments equal importance of the body must be realised [SLAT95, SLAT93a]. This is done through the virtual body placed at the egocentric location in the virtual world. Just as the physical body grounds the participant to the real world, so does the virtual body help to ground them to the virtual world. An ideal virtual body should therefore not only map accurately to the limb positions of the physical body but should also reflect facial expressions and gestures for communication by the participant [PAND96].

WP4.3 was developed under DIVE 3.1.0. This provides the facility for visual embodiment of participants. Verbal communication is available through the application but our development was based on the RAT audio communication system [HARD96].

Although verbal communication is the primary mode of information exchange between participants involved in way finding, it is often reinforced by hand gestures. These may include:-

- pointing - to denote the direction of motion;
- a forward motion of the hand with a deviation to the left or right - to denote a left or right turn;

- a horizontal circular motion of the hand - to denote a roundabout, circular structure or an open space;
- an upward motion of the hand - to denote a raised structure.

These gestures can often surpass the barriers of language. Of course not all gestures are based on single hand movements. An immersive system should however provide for at least a single tracked limb. The development environment was based on the tracking of the head and the hand.

As part of the overall usage of DIVE we are planning for the incorporation of a virtual body with limbs responsive to body tracking information provided by an immersive system. Figure 1 shows two types of virtual bodies which are currently in use in the development.

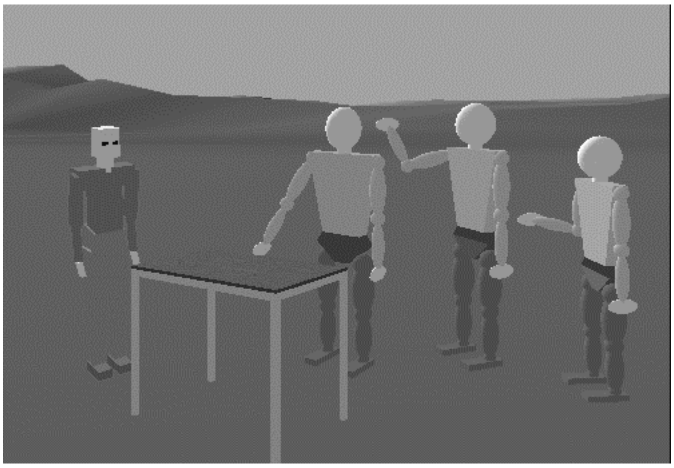

Figure 1: Virtual bodies currently used in desktop (static) development (far left body) and immersion tracked (dynamic) development (from far right - tracking allows different configurations of the body)

5 Body Centred Interaction

The principals of collaborative interaction techniques incorporates the idea of body centred interaction (BCI) [SLAT93b]. This method encourages the exploitation of interactions which make use of the body since this is the locus of the participant both in the real and virtual world. Examples of these try to reduce the proprioceptive/sensory data mismatch between the physical body and its action in VE. For instance, a ubiquitous way of navigation in VE is to point and make a hand gesture or press a button on a hand held device. The system may respond by moving the participant along the point vector. A mismatch is now immediate between what the kinesthetic system experiences when compared to the visual. Optic flow generated from the motion informs the brain through the visual system that motion is occurring. However, the kinesthetic and vestibular system tell it that the body is stationary (but pressing a button or making a hand gesture). This mismatch is detrimental to the sense of presence. The BCI paradigm has been used to reduce this mismatch in ground based navigation applications where a participant feigns walking by walking in place [SLAT95]. A neural network based system recognises the distinctive set of head movements associated with this (from the HMD sensor) and translation occurs in the direction of gaze.

In the context of the map exploration of WP4.3 the BCI paradigm is used in dynamically sweeping out a road in the 3D space and in interactive shrinking and growing of the lost traveller. For example, the traveller may be placed into the newly created map by being shrunk down to the scale of the map. This is done by the helper virtually pushing down on the head of the lost traveller. At the same time the traveller may squat down to facilitate this action and their virtual body is shrunk. This process of squatting down is such as to provide congruency with the shrinking of the virtual body (as a result of which the lost traveller experiences the world as getting bigger). After shrinking the helper may pick up the traveller and place him/her in the map. The traveller is then free to explore it from within. The inverse process will hold for bringing the lost traveller out of the map and scaling them to their original height.

6 Depth of Presence

The BCI paradigm for shrinking and growing of participants has been demonstrated as "stacking environments" for enhancing depth of presence [SLAT94]. That is, a participant can simulate the process of entering the virtual environment while already in such an environment, which can be repeated to several levels of depth. This exploits the notion of *presence transformations*. In order to enter from everyday reality into VE the subject must go through a particular transformation procedure, e.g. donning a HMD and gloves. Given E to be the environment immediate to the subject, and T to be the transformation when carried out in environment E, we obtain a new environment $T(E)$. It follows that the environment of the VE when entered from everyday reality is $T(R)$ where R denotes the environment of everyday reality.

Since a VE provides the potential for simulation it follows that whilst immersed we can repeat the transformation process of donning a (virtual) HMD and gloves and entering another simulated environment. Thus while in $T(R)$ the transformation T may be applied resulting in $T(T(R)) \equiv T^2(R)$. This procedure may be repeated to degree i to give $T^i(R) \equiv E_i$, where $E_0=R$ and E_i is at depth i in the stack. We have shown that the dimension through which the subject moves may be regarded as corresponding to presence [SLAT94].

In the context of WP4.3 we hypothesise that the process of shrinking and inhabiting the virtual map E_2, will place the participant deeper in the simulation, thereby increasing their sense of presence when later brought back to E_1.

7 Interactive Map Drawing

We have investigated several methods for the process of interactive map drawing. These range from a 2D based approach to a 3D metaphor. Whilst creating the map collaboration may occur in several ways:-

- the helper may construct a map and the lost traveller annotates it.
- the lost traveller constructs a map of the destination (e.g. tall building with gas tower nearby) and the helper constructs the path to the location.
- the environment contains *n* number of helpers and *m* number collaborate in the map construction.

The map drawing approaches have included a number of metaphors: Pen and Paper, PDA (2D), PDA (3D), and the 3D Interactive Space. A description of each is given below.

- *The Pen and Paper Metaphor (2D)*

A simple method of construction is by pen and paper. Here the helper produces a virtual sheet of paper on which a 2D map is drawn on with a virtual pen. The lost traveller may still inhabit and navigate this map.

- *The PDA Metaphor (2D)*

Based on a Personal Digital Assistant (PDA) the helper draws on a virtual PDA which is able to recognise gestures. For example, it will replace a 2D horizontal circular gesture with a 2D interpretation of a roundabout and a zigzag gesture with a zebra crossing.

- *The PDA Metaphor (3D)*

This is similar to the 2D version but with the 2D gestures being interpreted as 3D models. Figures 2 shows a map constructed using this metaphor.

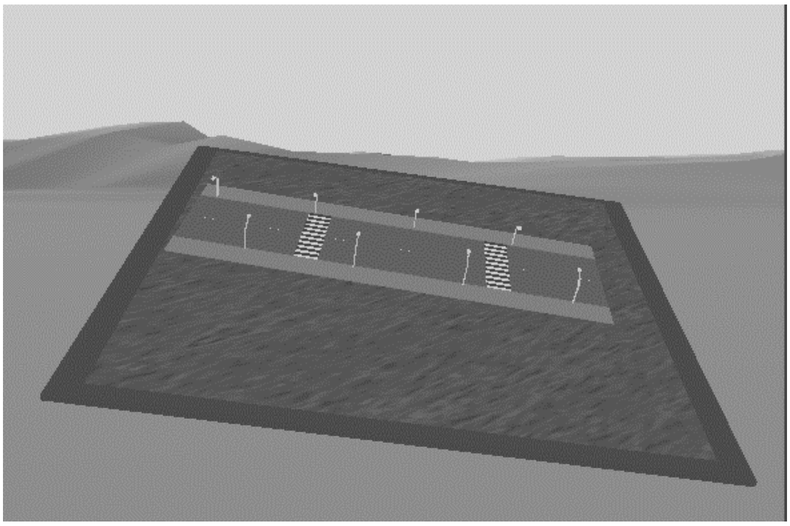

Figure 2: A road segment created with the PDA metaphor showing zebra crossings

- *The 3D Interactive Space*

Here the interactive space surrounds the participants and is part of the immediate environment. 3D gestures are interpreted as the helper sweeps them out in the surrounding space. The map is constructed in a more dynamic space - whole body and arm movements are realised as virtual objects which may be manipulated. We have implemented a range of gestures which include:-

- a horizontal sweeping action to represent a road which is dynamically created as the helper moves their hand. This road is created with sidewalks and street lighting;
- a "cross" gesture to denote a pedestrian crossing with a set of traffic lights;
- a circular gesture to denote a roundabout;
- a zigzag gesture to denote a zebra crossing.

These artefacts are placed within the 3D map at the position the gesture took place. Certain gestures may only be active based on the state of the scene. A pedestrian crossing gesture, for instance, will be recognised if made on the virtual road (Figure 3).

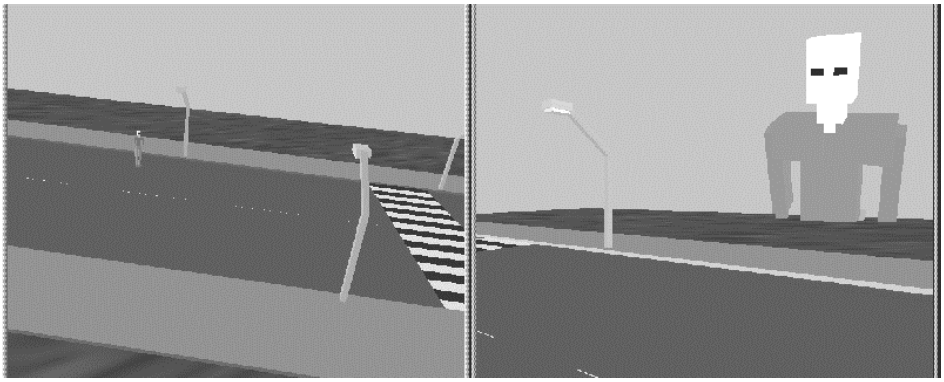

Figure 3: Two views of the map showing (a) the traveller exploring the map and (b) the viewpoint of the traveller looking up at the helper

We note that not all descriptors of the map are constructed by gestures. Generic buildings and trees for instance may be inserted from a library of objects which is available from a toolbelt worn by the helper or held in the non-tracked hand of the virtual body.

In this section of the work our attention has been focused on those interactions which are conducted immersively and are therefore based on a 3D space. In the following sections we will describe interactions which are still based on a 3D space but need not be conducted immersively.

8 Path Planning Map (PPM)

8.1 Motivations

As in regular environments, participants operating in a virtual environment are highly aware of their immediate surroundings. They will pay attention to the different participants and reactive objects located within a small radius and will use distant objects as reference points for navigation. The radius of this awareness sphere will change with the size of the environment and its type, e.g. participants located in a room will mostly interact within the room, while participants sharing an open space, such as a big room or a landscape will have a bigger awareness volume.

One way to let the user gain a larger context of the virtual environment is by using the World In Miniature metaphor (WIM, [STOA95]), with some added features, which we in our implementation call a Path Planning Map (PPM). This provides the user with a hand held miniature representation of the world, somewhat like a map in real life. The virtual map metaphor benefits from being something that everyone is familiar with and has shown to be a concept that most people catch on to quickly, hence it is natural in the sense that it builds upon knowledge the user already possesses.

As opposed to Section 7, the PPM does not require users to interactively present their mental representation of the environment by drawing a map. Instead, the PPM presents a hand-held miniature version of the world, gaining in accuracy compared to interactive map drawing, but loosing in simplification.

8.2 Purposes

The PPM works well for enhancing communication and awareness, and can be used for different purposes, such as:

- Spatially locating and orienting the user: in the PPM the user can see where s/he is and which way s/he is facing.
- A visual searching tool to find other users.
- To see what/who is at other places not in the users vicinity.
- Moving directly to a point on the map.
- Path planning: A person (within the environment) familiar with the world can together with the lost virtual traveller draw paths from one point to another on the map while discussing the path.

8.3 Functionalities

The PPM uses a concept called "subjective views" [SMIT97] which allows the world to be presented differently to different users. "Subjective views" are a result of the COVEN work. In this case, we use the technique to show the paths created using the map in the world only to map users. Indeed, users currently not using the map should not see environmental additions especially made for map users. The term "subjective views" means that a users can be arbitrarily grouped and have the objects in the world represented differently, e.g. an electrician walking through a building may want to have the walls transparent and the wires in the wall visible. Another example would be to have all the text in a virtual environment shown in the user's language.

When initiated, the map makes a simplified[1] three-dimensional snapshot of the environment in which it is started, at a smaller scale. This makes it possible to use the map in any environment, immersively or not, and to use it as a helper tool in any application.

The map has the following functions:

- It can be grasped and rotated like any other object.
- The map can be iconified and/or scaled up and down, in order to make it fit in all kinds of environments.
- It monitors, in real-time, the different users present in the environment and represents them by simple icons within the map-environment.

---

1 Simplification is achieved by using what is usually a rendering optimisation. Only the worst visible level of detail of objects to be added to the PPM is kept when copying them. The level of detail technique consists of having different representation for the same object. The rendering routines choose an appropriate (and usually simpler) representation for an object when it is visualised from a distance. The worst visible level of detail kept for the PPM generally gives an accurate idea of what the object is, while not increasing the rendering load and achieving part of the simplification which occurs on traditional maps.

- Any user, one at a time, may take the map with him. The user holding it will be represented in a differently colour on the map.

- It can be used for navigation by smoothly transporting a user to a given location on interaction.

- A specific mode allows anyone to draw a path on the map (several people can interactively together make the path). This path is shown on the map and also subjectively within the environment.

The path within the virtual environment can be used in two different ways:

- All members of the "map" subjective group are able to follow the path freely, since it is drawn within the environment.

- All users being members of the subjective group are able, on interaction, to follow the path automatically.

The map in path drawing mode is shown in Figure 4. The map is at the front, at the bottom of the picture, while its control buttons are floating in the air, in the centre of the picture. The regular virtual environment occupies the rest of the picture. It is possible to see a path drawn on the map from the water tower to the church. This path is also drawn in the virtual environment, and can be seen in the background.

Figure 4: The map in path drawing mode

In the Figure 5, the PPM is iconified and only represented in space by a floating button. The user is currently following the path from the water tower towards the church. The path is composed of joint three-dimensional arrows.

Figure 5: Path following

8.4 Setups

We have also implemented a way of having a tracker attached to a physical plate (representing the map) that the user holds in his/her hand, which is directly linked to the virtual map. This allows for a more natural interaction with the map in immersive and semi-immersive environments, such as HMDs or CAVE-like environments.

We have also experimented with having this map displayed on a touch screen desktop in a CAVE-like environment (see Figure 6). With this setup several users can collaboratively discuss and interact with the map while having the view of the virtual world on the big vertical projection screens. This setup addresses the problems of many-to-many interactions, i.e. when several users, located at the same physical site, meet several other users located at a remote physical site. It is interesting to notice that the map environment, in this case, serves interaction within a single group rather than in-between remote groups.

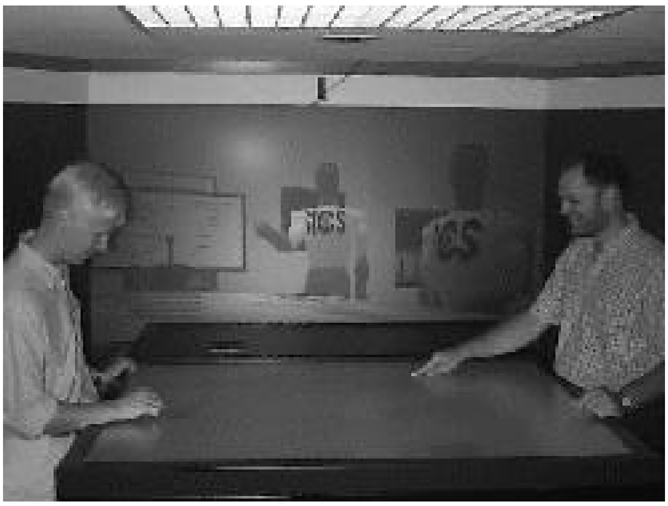

Figure 6: Map on a touch screen in a CAVE-like environment

8.5 Future work

The PPM could, in the future, be extended to allow different filters to be applied to the representation. A user acting in one world but interested in what is happening in other worlds could use multiple PPMs as "radar coverage" on those worlds. Another feature that could be added, is that a user could leave a trail behind him, which could be used as a "history". One could also place a grid over the map for measuring distances, areas, and volumes.

9 Access Control in Collaborative Environments

Access control is an integral but often neglected aspect of interaction with and navigation within collaborative virtual environments. This section develops a spatial approach to access control in such environments. This means that we are interested in general techniques for managing individual and group access to resources located within a virtual environment, assuming that related security issues such as authenticating user's identities and auditing their actions have already been dealt with. Traditional access control mechanisms restrict access to resources such as data and applications by associating access rights with the resources or with the users who are trying to access them. These two alternatives are referred to as Access Control Lists and Capabilities respectively [LAMP74]. Our spatial approach differs from these in that we associate access rights with spatial boundaries, such a portals between different virtual worlds or the boundaries of regions within worlds, so that gaining access to a resource is equivalent to gaining access to the space within which it is currently located.

This ties in closely with ideas on navigation and wayfinding. A natural consequence of creating a full, clear picture of the access related information pertaining to a virtual space is that you also create a full, clear picture of the possibilities for traversal in the space. The access rights associated with various routes are stored, and we can use these routes in our everyday navigation of the space. So, rather than restrict and constrain the possibilities for movement in a space we argue that our approach offers the possibility to enhance movement and wayfinding in a virtual space through the ability to provide end users with not only access information, but also the navigational information to go with that access information. The end user can not only find out what credentials are need to go from one place to another, but can also find out which routes are available to them as well.

We introduce the model, first looking at how boundaries are use to segment space. We then introduce the idea of an access graph, an abstract mathematical representation of the structure of a virtual universe and of the possibilities for access between its different spaces. Access graphs allow us to reason about the overall access properties of a complex structured environment and about the effects that individual changes might have on these properties (e.g., understanding the global consequences of changing the access rights associated with a local boundary). They therefore form the basis for developing powerful access management tools which might aid users in applying and administering access rights. The access graph is a formal map of the environment.

Section 9.2 presents the motivations for adopting this spatial approach. Section 9.3 introduces the key concepts of boundaries and access graphs. Section 9.4 then discusses the problems of understanding and managing access in complex environments. Following this, section 9.5 explores a concrete example of our approach which we have implemented in the DIVE system. Section 9.6 offers some ideas for future work.

9.2 Why a spatial approach?

Why is it important to consider access control from a spatial perspective? In this section we identify some of the potential benefits and problems with the spatial approach.

First we need to emphasise how the spatial approach differs from others. The essence of the spatial approach is that access to an object depends upon where it is located and upon whether a given user can gain access to this space by crossing a sequence of boundaries. Although existing access control mechanisms support this idea to some extent (e.g., the role of directory permissions in the UNIX filestore), our approach takes it to the extreme by making spatial boundaries the primary mechanism for governing access.

We propose that this offers a number of potential advantages in the context of collaborative virtual environments. First and foremost is its natural integration with the whole philosophy of virtual reality. VR is about the exploitation of people's spatial reasoning to improve their interaction with digital information. VR applications rely heavily on the use of spatial metaphor in both realistic and more abstract applications (e.g., simulations and database visualisations respectively). Through the *appropriate* exploitation of the spatial metaphor, access control may be easier for people to understand than with other approaches. Indeed, it can be argued that the abstract and relatively complex nature of some current approaches actually represents a threat to security as it may be difficult for many people to apply them in an appropriate way or to understand the consequences of specific actions (e.g., of changing a directory execute permission high up in the UNIX filestore). Second, our approach can be realised as a direct extension to the existing VR mechanisms of portals and spatially partitioned virtual worlds. The linking of discrete virtual worlds through portals, a simple kind of boundary, is a common feature among VR systems. More recently, several systems have introduced spatial partitioning and therefore some form of spatial boundary in order to divide large virtual worlds into more manageable areas. For example, the Spline system [BARR96] consists of a collection of regions (locales) which can be combined together using arbitrary 3D transformations. In NPSNET [MACE95] the worlds are constructed of hexagonal cells each with an associated multicast group. RING [FUNK96] scopes interaction based on occlusion and utilises a number of servers, each of which are responsible for their own local region of the world. Finally, MASSIVE-2 uses the concept of third party objects to create nested and dynamically evolving structures of regions which apply different awareness adaptation effects across their boundaries [BENF97].

However, there are some potential disadvantages to the spatial approach. It is more difficult to specify fine-grained restrictions on the ability to perform specific actions on individual or small groups of objects. This would require wrapping objects in ever smaller spatial boundaries and finding appropriate metaphors for associating specific actions with these boundaries. It is extremely difficult to group such objects together according to non spatial relationships (e.g., where objects that are far apart need to be treated in a common way). It makes more sense to employ access control list and capability based techniques at this level; using a broad spatial approach for general access control and a specific approach for the detailed fine-grained access. Furthermore, the possibilities for accessing an object can change as it migrates around a virtual environment making it more difficult to specify globally consistent and stable access rights (this would at least require the ability to move bounded spaces through other bounded spaces as realised in the MASSIVE-2 system). Another problem concerns action at a distance. As described above, the spatial approach assumes that one moves to the space in which an object is located in order to access it. Support for action at a distance requires some relaxation of this principle. Indeed, this is a general problem with spatial metaphor and we will offer a possible solution in the following section.

We therefore argue that the spatial approach potentially offers a natural and understandable access mechanism for collaborative virtual environments, although at the cost of some flexibility. Of course, it would be possible to combine it with other approaches (e.g. to retain individual access rights for key objects within an environment).

9.3 SPACE: SPatial Access control for Collaborative Environments - the model

The basis of our approach was originally presented in [BULL94]. In this section, we offer a brief overview of its main concepts, before going on to present an implementation within the DIVE system. We begin with two key components of the model: boundaries and access graphs.

9.3.1 Boundaries

The most fundamental component of the model is the *boundary* as defined by Bowers [BOWE93]. Boundaries provide a way of segmenting virtual spaces into distinct regions and may control both the traversal of and awareness within these regions. As noted above, boundaries might take the form of portals between worlds as in systems such as DIVE and MASSIVE and also the VRML standard, or might be used to partition single worlds as in Spline, NPSNET, RING and MASSIVE-2. Our access model introduces the fundamental idea that in order to traverse such a boundary, one must first possess necessary and adequate *credentials*. In the most general case, these credentials involve matching an arbitrarily complex combination of attributes of the user (or indeed group of users). Examples might include the use of simple numeric values as in the worked example below which is based on an extended clearance-classification scheme [BELL73]; matching attributes of a user's description such as their name and status; or possession of other objects such as keys and passwords. Whatever the case, these credentials are presented to the boundary which determines whether or not the user can cross.

9.3.2 Access graphs

Multiple boundaries might be used to create complex spatial structures such as many worlds inter-linked by different portals or nested and tiled regions within a single world. In such cases, it is possible to construct a useful framework called an "access graph" which provides a concise summary of the possibilities for and constraints on movement around the spatial setting. In particular, access graphs allow us to reason about the possibilities for accessing the objects contained within a given space.

In general, each discrete region of space is represented by a node in the access graph and each boundary by an arc. The arcs are labelled with their associated credentials and may be unidirectional if one-way boundaries are supported. Where there are multiple boundaries between two spaces (e.g. if two virtual rooms are separated by a door, window and a wall), there will be multiple arcs between their two nodes in the access graph. However, in those instantiations of the model where the possible credentials can be ordered in some way according to their degree of power to grant access (i.e., where there is a well defined hierarchy of credentials), it will be possible to replace multiple arcs with a single arc which represents the easiest boundary to cross (i.e., the one requiring the minimum credentials).

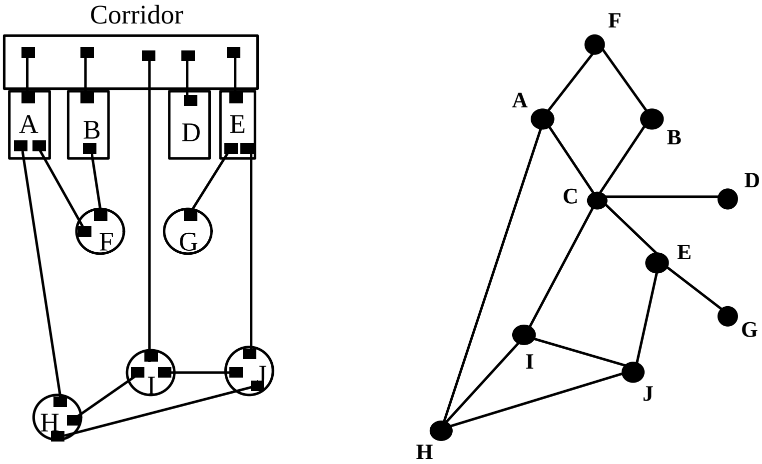

Figure 7: An example space and its associated access graph

For example, in a conventional physical space where a door is combined with a wall, it is the door that is the easiest to cross and therefore becomes the critical boundary. Furthermore, the arcs of the access graph can then be labelled with a number which represents the position of its credentials within the overall ordering (i.e., a label on an arc indicates the relative difficulty of crossing the associated boundary). Then, the application of standard mathematical techniques to this access graph will allow us to easily reason about and manage access rights across the entire space. As we shall see in section five, it is this ability to apply standard mathematical techniques to access graphs in order to predict the consequences of alterations to specific boundaries that makes this approach to access control particularly interesting.

Figure 7 shows a hypothetical virtual universe, that of a virtual office and several associated information visualisations. Some of these visualisations are local to the office, others stored in a separate virtual world. Boundaries exist between different rooms (A, B, C, D and E) and also surround the visualisations (F, G, H, I and J). A portal exists between the Office world and the Archive world. The access graph for this space is seen to the right of Figure 7. In this case, all boundaries, and hence arcs, are bi-directional and we have assumed no ordering of credentials so that the arcs are unlabelled.

Now that the two most fundamental components of the model have been introduced, we define two further concepts related to access graphs, relative and absolute classifications, which are required for the subsequent discussion.

9.3.3 Relative and absolute classifications

The *relative classifications* of two regions describe how difficult it is to move between them. More specifically, the relative classification of region A relative to region B is defined as the minimum credentials required to move to A from B via *any* valid route. Where directional boundaries are supported (e.g., one way portals), this need not be the same as the relative classification of B relative to A. We derive this relative classification through two steps.

1) The relative classification of a particular route between two regions is the maximum of the classifications of the arcs traversed along the route. In other words, whether one can traverse a specific route is determined by the most difficult boundary to cross that is encountered on the route.

2) The relative classification of region A relative to region B is then the minimum of the relative classifications across *all* possible routes from B to A. In other words, whether one can move from one region to another is determined by the easiest of all possible routes between them.

We define the notion of *absolute classification* to be the minimum credentials required to enter a given region. In the most general case, this is at least the minimum of the classifications of all directly incoming arcs to node representing this region. In order to enter this region a person must have credentials at least equal to this absolute classification. To increase the absolute classification of a region, the boundary with the minimum credentials must be modified (see Figure 8).

An additional possibility is to constrain the possible entry points into a virtual universe so that, in effect, a user can only join its access graph at one of a designated set of start nodes. In this case, the definition of the absolute classification of a node can be extended to be the minimum of all of the relative classifications across all possible paths between a valid start node and the target node. This will be greater than or equal to the minimum of all of the *direct* incoming links as defined above.

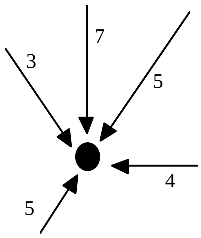

Figure 8: Absolute classification

9.3.4 Modes of movement

We raise one final issue in this section - that of the mode of transit adopted and, in particular, the mechanisms of teleportation and action at a distance. It would seem that our model requires that a user must actually move to an object's current location in order to access it, perhaps crossing many intervening boundaries on the way. Although this might often be the case, two other possibilities need to be considered:

- *Teleportation* - users directly specify the address of a destination region in order to move there without ever being present in any intervening region (e.g., directly supplying a URL in a VRML browser).
- *Action at a distance* - users act on a remotely located object without ever leaving their current location.

Both of these possibilities are allowed by our access control mechanism, so long as the user's credentials would have allowed them to traverse at least one valid route through the access graph. In other words, the access control mechanism is concerned with whether the journey could logically have been made; how the system actually enables navigation and interaction, given suitable access rights, is another matter.

9.4 Understanding and Managing Access Control

In order to provide effective access control, users must be supported in understanding the effects of different boundary configurations and of any proposed changes to these. For example, changing the credentials associated with a boundary at a critical junction might have a major impact on the access characteristics of an entire environment. A key goal of our approach is to enable the system to actively encourage this kind of understanding by automatically answering a range of access related questions. In this section, we identify a number of such questions. In later sections, we show how these can be answered by inspection of an access graph. Indeed, this list of questions defines the functionality of an access control management tool called the security broker which we have implemented to work with DIVE. Referring to Figure 7 as an example, we identify the following key questions that would be relevant to the users and managers of a virtual environment access control mechanism.

- Is it possible for a given user to move to visualisation J from office A?
- What credentials would a user need to move from visualisation J to office A?
- How secure is office A relative to B?
- Where are the most and least secure regions in which to place an object?
- Where can the object currently located in A be moved to so that its level of security is not lessened?
- What paths can be taken when moving this object, so that its level of security is not compromised during the move?
- If I have security credentials x where can I go to within an environment and where am I unable to go to?
- What properties do I require to be able to traverse the whole environment? (i.e., to become a Super User)
- What is the effect of increasing or decreasing an object's level of security?
- What effect does changing the credentials associated with the boundary between the Corridor and office E have on the rest of the space?
- If I add a region and associated boundaries what properties do the new boundaries need to have to maintain confidence in the rest of the graph?
- What is the effect of removing a region (node) from the graph? Do certain conditions have to be met for this to occur?

This concludes the introduction to our spatial approach to access control in collaborative environments. The following section presents an implementation of this approach.

9.5 SPACE - an implementation of our approach

We have implemented our model as an extension to the DIVE system [CARL93]. Our overall design is shown in Figure 9. The key features of our implementation are:

- The adaptation of the Bell and LaPadula clearance-classification model to associate simple numerical credentials with gateways (i.e. portals) between DIVE worlds.
- The extension of DIVE's world description files to include the specification of these numerical credentials as part of gateway definitions.
- The extension of the DIVE run-time system to implement the checking of credentials before allowing users to move through gateways.

- The implementation of a separate security broker application which inspects a set of extended DIVE data files, having first converted them to an appropriate internal format, in order to an construct access graph for a given DIVE installation. The security broker then manipulates this graph in order to answer the questions raised in section four and allows users to make changes to the access graph and to update the underlying data files if happy with their changes.

9.5.1 Credentials, gateways and the Clearance Classification Model

We concern ourselves with traversal of different DIVE worlds through gateways (i.e., DIVE gateways are our boundaries). To this end, credentials have been associated with gateways by extending the world definition format for DIVE. Our initial implementation of credentials is derived from the Bell and LaPadula Clearance Classification model in which each resource in a system is given a classification and each user a clearance. Classifications and clearances are ordered and can therefore be represented by numerical values. Access is then controlled according to two axioms [BELL73, LAND81]:

1) No user may read information classified above his clearance level ("No read up");

2) No user may lower the classification of information ("No write down").

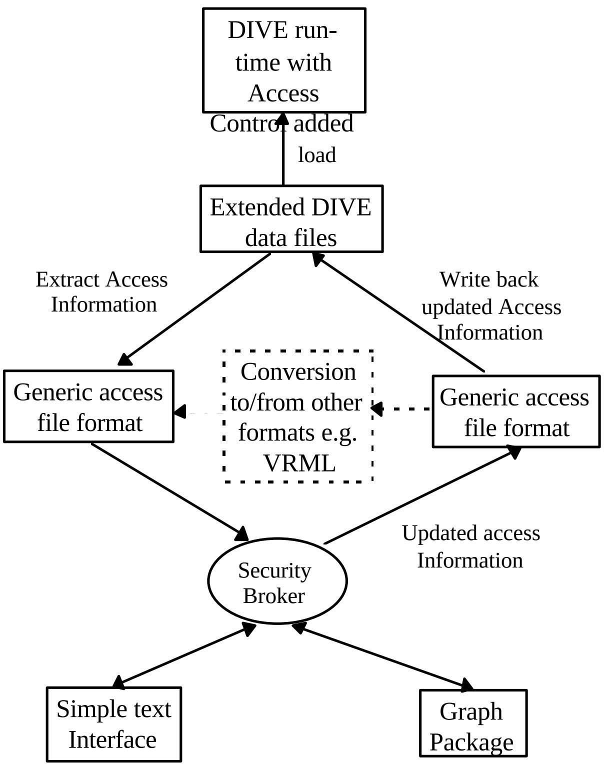

Figure 9: The structure of the implementation

In our interpretation, users are given clearance values and boundaries are given classifications. Following axiom 1, no user is permitted to traverse a boundary unless their clearance (integer value) is at least equal to that of the boundary's classification. The DIVE run-time system has been modified to support access control at gateways between worlds. When a user tries to pass through a gateway the system checks their clearance, if it is equal to or greater than the classification at the gateway they are able to change worlds. If their clearance is insufficient they remain in the same world. Clearance has been implemented as an additional field in the DIVE user embodiment description file and classifications as an additional field in gateway definitions in DIVE world definition files.

9.5.2 The Security Broker

The security broker is an access control management application. It allows users to manipulate access rights (e.g., to update credentials), to experiment with changes before they are applied and to answer the different questions raised in section four. The security broker operates by first inspecting the available extended DIVE world description files and converting them into a general intermediate format which retains only topology and access information. The advantage of working with such an intermediate format is that converters to and from other world description formats such as VRML can be easily implemented in the future. It inspects the resulting files in order to construct an access graph and then applies well known mathematical techniques to this graph in order to implement the user's requests.

The first step in constructing the access graph is to identify all possible paths in the environment, and their associated access rights. The security broker is able to identify these paths using an adjacency matrix for the environment, A.

Definition: The adjacency matrix of a graph G which has vertices (v_1, v_2, ..., v_n) is the n x n matrix A(G) = (a_{ij}) where a_{ij} is the number of edges joining v_i and v_j.

This is a matrix which lists all of the one-step paths in the graph of the space (i.e., which DIVE worlds are connected to which others), and from it we can generate all of the longer paths in the space. A^2, simple matrix multiplication of A by itself, lists all the two-step paths in the graph; that is if $(a_{ij})^2$!= 0 (the matrix entry in the i-th row and j-th column), there are that many two-step paths between i & j. For a space with n worlds (hence n nodes in the access graph) we therefore only need calculate (n-1) matrices from the second power of A up to A^{n-1} (Figure 11) in order to generate all possible paths between worlds. Given a few simple constraints that we shall now describe, the complexity of this task can be reduced considerably.

9.5.3 The Constraints Applied to the Universe

The first constraint we apply is that cyclic arcs are not permitted. A consequence of this is that for each power of the Adjacency matrix we can always set the main diagonal to zero.

The second constraint concerns cyclic paths and paths containing cycles. Such paths are simplified to paths which only contain one iteration of the cycle. If we remove the extraneous cycles early enough it drastically cuts down the complexity in the higher powers of the adjacency matrices.

Our final constraint concerns re-entrant paths, which pass through a particular world more than once on their way to their destination. Such paths are simplified to paths which do not contain the re-entrant section of the path. Again, this constraint helps reduce complexity in the calculation of the powers of the adjacency matrix.

To summarise:

1) cyclic arcs are not permitted e.g. AA is not defined.

2) cyclic paths are simplified e.g. HIJHIJ becomes HIJ.

3) re-entrant paths are simplified e.g. ACIJECB becomes ACB.

We now consider the process of identifying possible paths and their classifications in more detail. The starting point is the example adjacency matrix in Figure 10, which is labelled version of that from Figure 7. The first step is to set the main diagonal of this matrix to zero (applying constraint 1). We then store all the one-step paths and their classifications. Multiplying the modified A by itself gives us all the two-step paths in the universe. Again we set the main diagonal of the matrix to zero. To identify the paths we make use of the information we already have about one-step paths. For a two-step path, the end node of the first step must be equal to the start node of the second step, and we know what the start and end nodes of the path are so it is a case of simple inspection to identify all of the two-step paths (e.g. a path from A to B consists of start step AX and end step YB, so for a valid path we must find one-step paths AX and YB where X=Y). There can be no cyclic or re-entrant paths.

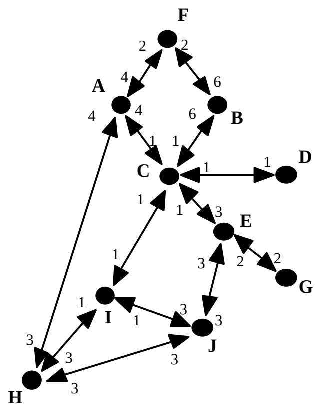

	A	B	C	D	E	F	G	H	I	J
A	0	0	1	0	0	1	0	1	0	0
B	0	0	1	0	0	1	0	0	0	0
C	1	1	0	1	1	0	0	0	1	0
D	0	0	1	0	0	0	0	0	0	0
E	0	0	1	0	0	0	1	0	0	1
F	1	1	0	0	0	0	0	0	0	0
G	0	0	0	0	1	0	0	0	0	0
H	1	0	0	0	0	0	0	0	1	1
I	0	0	1	0	0	0	0	1	0	1
J	0	0	0	0	1	0	0	1	1	0

Adjacency Matrix

	A	B	C	D	E	F	G	H	I	J
A	0	0	1	0	0	2	0	3	0	0
B	0	0	1	0	0	2	0	0	0	0
C	4	6	0	1	3	0	0	0	1	0
D	0	0	1	0	0	0	0	0	0	0
E	0	0	1	0	0	0	2	0	0	3
F	4	6	0	0	0	0	0	0	0	0
G	0	0	0	0	2	0	0	0	0	0
H	4	0	0	0	0	0	0	0	1	3
I	0	0	1	0	0	0	0	3	0	3
J	0	0	0	0	3	0	0	3	1	0

Classification Values

Figure 10: Access Graph to Adjacency Matrix

Three-step paths are similarly dealt with, only we identify a one-step path between the end node of the start step and the start node of the end step. We must be careful to remove all re-entrant and cyclic paths before we store the results (e.g. ACAF is a valid three-step path, as is ACBC and AFBF, but all fall foul of our constraints).

By the time we get to four-step paths, a pattern is emerging. We just consider the start and end nodes of each path, identify a valid path between these nodes and store the result. For four-step paths we look at all three-step paths between start and end steps, but we have already identified all of these, and as we were careful to remove all cyclic and re-entrant paths before we stored the results, we just use the stored results to build up the broker. This way we are guaranteed to reduce the complexity of the powers of A. We only need consider powers of A up to n-1 for a universe with n worlds, because the longest possible path which satisfies our constraints can visit each world only once, so for n worlds this path has length (n-1).

To store the classification values associated with each path, we determine the relative classification of the start node to the end node along that particular path (i.e., we find the single step in the path which has the maximum classification.

As a result of these calculations, the broker now holds the collection of the powers of the adjacency matrix (after the simplifications have been applied), all the paths which exist in the virtual universe, and the classifications associated with these paths. This information is all that is needed in order to check for the existence of any paths between two nodes, to identify the paths should they exist and to inform us of the classifications associated with these paths.

As an example of this process, consider the example space from section three. Figure 10 shows the access graph for this space, and the adjacency matrix for the space which has been derived from this. Classifications have been added to the boundaries, and a matrix holding these classifications has also been constructed. Figure 11 shows the powers of the adjacency matrix which have been calculated and simplified.

A	A^2	A^3	A^4	A^5
0010010100	0201100021	0020111113	0122401221	0212214223
0010010000	2001100010	0010011302	1001200243	2010312234
1101100010	0000021202	2100100222	2000211221	1100122011
0010000000	1100100010	0000021202	2100100222	2000211221
0010001001	1101000120	1010020311	4221010112	2312040221
1100000000	0020000100	1102200031	0010102315	1121401433
0000100000	0010000001	1101000120	1010020311	4221010112
1000000011	0020110011	1322301010	2222133013	2202241020
0010**0**000101	2101**2**00101	1020**1**32101	2422**1**11101	2312**2**31201
0000100110	1020001110	3222110010	1312251310	3411132010

A^6	A^7	A^8	A^9
0101122211	0000111110	0000001000	0000000000
1021423211	0012214121	0001002000	0000000000
0200111011	0100101000	0000001000	0000000000
1100122011	0200111011	0100101000	0000001000
1411040011	1211010121	0001000000	0000000000
2212404242	1101104000	0000001000	0000000000
2312040221	1411040011	1211010121	0001000000
2200022002	1100100010	0000001000	0000000000
1111**1**42001	1201**2**01101	0000**0**02000	0000**0**00000
1111121210	0101101010	0000001000	0000000000

Figure 11: The nine powers of the adjacency matrix

If we wish to know whether we can move between worlds I & E we inspect the a_{95} entries in the matrices in Figure 11 (the entries highlighted in bold). There are two each of two-step five-step and seven-step paths, and one each of three-step four-step and six-step paths, so in total there are nine possible paths between I & J. By inspecting the paths database we can identify what these paths are and what their classifications are, and this will tell us if we are able to move between them given our current credentials.

9.5.4 The Questions Revisited

We now return to the questions identified in section four and show how, given this local database of information about paths and classifications, the security broker is able to answer them.

Is it possible to move to visualisation J from Office A? and *What properties do I need to do this?*

To find out if we can move between two worlds we examine the powers of the adjacency matrix. If there exist any non-zero entries in any of the matrices, it is possible to move between the worlds. To find out the clearance required to move between the worlds we pick out the path with minimal classification from the list of possible paths between the worlds.

If I have security credentials x where can I go in a space and where am I unable to go? and *What properties do I require to be able to traverse the whole space?*

By returning all nodes whose absolute classification is less than or equal to our clearance we identify where we can go. All other nodes (those with classifications greater than our clearance) are no-go areas. To traverse the whole space we need a clearance equal to the maximum absolute classification for the space.

How secure is Office A relative to B?

After confirming that at least one path exists between the worlds we examine all the paths between A and B to find the path with the lowest relative classification.

Which is the most secure region in which to place an object? and *Which is the least secure region in which to place an object?*

We identify the node(s) with the highest absolute classification for the most secure world(s) and the lowest absolute classification for the least secure world(s).

Where can an object be moved from its current location so that its security rating is not lessened?

We identify the absolute classification of the object's current location, and list all locations whose absolute classification is equal to or greater than this value.

What path can be taken in moving this object, making sure it is not compromised during the move?

We identify all paths whose relative classifications are greater than or equal to the absolute classification of the object's current location.

The final four questions posed have a common theme. They are concerned with changes to the space, and their effects. To answer these questions, the security broker applies the requested changes to its local database, re-evaluates the access information about the space, regenerating the paths and classification database, and compares this to the original information. As a final step, if the user is happy with the effects of their changes they can be written back to the appropriate DIVE files.

9.6. Future work

This section has developed a spatial approach to access model for collaborative virtual environments based on the attachment of credentials to the boundaries which separate or partition virtual worlds. In order to access an object, a user must first be able to make a journey (at least logically) from their current location to that of the object, crossing any intervening boundaries. A key aspect of the model is the way in which the access properties of a given virtual environment can be represented as an access graph. The application of standard techniques to this graph allows a range of access related questions to be answered and supports the understanding and management of access rights. Our approach has been implemented as an extension to the gateway mechanism within the DIVE system and as a separate security broker application to support access management within DIVE. We conclude this chapter by identifying some areas for future work.

9.6.1 Interface issues

Greater consideration is needed of interface design, both in terms of how access rights and violations are presented within a virtual environment and also in terms of the security broker. There is much scope for different feedback models when a user's clearance is insufficient. Does an alarm sound? How is the user told they have insufficient clearance? Should they be informed of alternative routes to their eventual destination?

In improving the interface to the security broker, we might exploit graph drawing techniques to graphically show and query an access graph. Examples include the daVinci graph package for 2-D graph drawing [FROL94] or the use of the Force Directed Placement algorithm for constructing 3-D graphs inside the virtual environment itself [FRUC91].

9.6.2 Group access rights

We might extend our approach to explicitly support notions of group access. What should happen when a group of users tries to cross a boundary together? Should their access be determined by the credentials of the weakest member; by the strongest (i.e., we can all cross if one of us has a key); or by the average or sum of their credentials (e.g., we can “storm the boundary” if enough of us get together)? A general solution to this problem which introduces different aggregation techniques for determining group clearance from those of individual members would open up many interesting possibilities. For example, by taking the sum of individual credentials, we could create spaces which no one could enter or leave on their own (i.e., whose contents were so sensitive that no one could be left alone with them). One approach to this aggregation problem might be to extend the crowd modeling framework described in [BENFb] which defines techniques for generating aggregate representations of dynamic crowds of participants in collaborative virtual environments.

9.6.3 Scaleability and distribution

The application of our approach to large-scale distributed collaborative virtual environments raises the issues of scale and distributed access to world data. For example, the construction of a single local access graph representing a globally distributed environment of millions of inter-linked worlds (such as a future VRML environment) will be impossible due to the scale of the computation involved and problems with gaining access to the necessary world data. Future work should therefore address the problem of developing a distributed access model and security broker so that responsibility for managing and presenting access information is spread across many local but communicating security brokers.

9.6.4 Potential applications

We conclude by reflecting on the potential scope of application of our approach. We anticipate possible application within a wide range of VR systems including those that exploit the notion of portals and other links between worlds (e.g., VRML) as well as those emerging systems identified above which support some form of spatial partitioning within individual environments. There is also scope for application outside of VR. For example, Rodden recently discussed techniques for modeling the collaborative properties of a range of systems, including 2-D windowing and hypermedia systems, in terms of abstract graphs [RODD96]. Our approach to access control could be introduced as a direct extension of this work. Finally, like VR in general, our approach could be used to simulate the access properties of real world environments and to provide an analysis and management tool for security within everyday physical spaces.

10 Achievements

In this work we have constructed a set of tools which may aid the lost virtual city traveller in way finding. We have run a number of collaborative trials with two participants using the PDA metaphor (3D) and 3D Interactive Space in the UCL COVEN model [COVE]. The collaboration was made using an SGI High Impact running a desktop VE, and an SGI Onyx running an immersive scenario. The participants were able to interactively construct a set of roads and populate them with traffic based artefacts.

We have also developed a Path Planning Map for enhancing communication and awareness in a collaborative environment. This has used the notion of "subjective views" which has been a result of the COVEN work. In addition we have developed a method of access control in shared environments based on a spatial approach which ties in closely to the work package goals on collaboration and interactive techniques based around the subject of wayfinding.

As a consequence of the workpackage a number of publications have been produced; [HANS97], [STEE97], [BULL97]. For example, in order to provide interaction with large scale scenes, it is necessary to build efficient data structures to support the fundamental actions of detecting collisions or proximity of objects. One significant problem with interaction with large scenes such as building interiors, is providing viewpoints which are at legitimate positions in the model. We wish to constrain the viewpoint not to lie in invalid areas such as the interiors of walls, and to maintain the correct eyelevel above the floor. [STEE97] describes an efficient technique (developed by UCL) to allow surface following and viewpoint collision detection within very complex scenes, which has approximately constant run-time cost. The underlying structure exploits the 2D coherency of natural navigation, and is constructed in a manner that makes use by collaborative systems simple. These properties contrast with standard techniques which have worse run-time cost, and do not scale have efficient scalability properties.

The work has also contributed to a project proposal (UCL) on dynamic collaborative surface creation for the design of virtual clothes.

References

[BARF95a] Barfield W, Sheridan T, Zeltzer D, Slater M (1995) Presence and Performance Within Virtual Environments, in W. Barfield and T. Furness (eds) Virtual Environments and Advanced Interface Design, Oxford University Press.

[BARF95b] Barfield W, Hendrix C (1995) The Effect of Update Rate on the Sense of Presence within Virtual Environments, Virtual Reality: The Journal of the Virtual Reality Society, 1(1) 3-16.

[BARR96] Barrus, J. and Anderson, D. (1996), "*Locales and Beacons*" in Proc 1996 IEEE Virtual Reality Annual International Symposium (VRAIS'96), San Jose, 1996.

[BELL73] Bell, D.E. and LaPadula, L.J. (1973), "*Secure computer systems: Mathematical foundations*", ESD-TR-73-278, vol.1, ESD/AFSC, Hanscom AFB, Bedford, Mass., Nov. 1973.

[BENF97b] Benford, S. and Greenhalgh, C. and Lloyd, D. (1997b), "Crowded Collaborative Virtual Environments", Proc CHI'97, Atlanta, Georgia, US, March 22-27, 1997 (in press).

[BENF97] Benford, S. and Greenhalgh, C. (1997), "*Introducing Third Party Objects into the Spatial Model of Interaction*", to appear in Proc. ECSCW 97.

[BOWE93] Bowers, J.M. (1993), "*Modelling Awareness and Interaction in Virtual Spaces*", in Proceedings of the 6th MultiG Workshop, Stockholm-Kista, May 1993.

[BULL94] Bullock, A and Benford, S. (1994), "*An Approach to Access Control for Collaborative Virtual Reality*", 6th ERCIM Workshop, Stockholm-Kista, June 1994.

[BULL97] Bullock A, Benford S (1997) Access Control in Virtual Environments, Proc. VRST 97, Switzerland, pp29-35.

[CARL93] Carlsson, C. and Hagsand, O. (1993), "*DIVE -- A Platform for Multi-User Virtual Environments*", Computers and Graphics, 17 (6), 1993

[FROL94] Frolich, M. and Werner, M. (1994), "*The Graph Visualization System daVinci - A User Interface for Applications*"; Technical Report No. 5/94, Department of Computer Science; Universitat Bremen; September 1994

[FRUC91] Fruchterman, T. M. J. and Reingold, E. M., Graph Drawing by Force Directed Placement, Software Practice and Experience, 21 (11), 1129-1164, 1991.

[FUNK96] Funkhouser, T. (1996), "*Network Topologies for Scalable Multi-User Virtual Environments*" in Proc 1996 IEEE Virtual Reality Annual International Symposium (VRAIS'96), San Jose, 1996.

[HANS97] Hansson P, Wallberg A, Simsarian K (1997) Techniques for "Natural" Interaction in Multi-User CAVE-like Environments, ESCW97.

[HARD96] Hardman V, Iken M (1996) Enhanced Reality Audio in Interactive Networked Environments, Proc. FIVE 96, pp55-66.

[HEE92] Heeter C (1992) Being There: The Subjective Experience of Presence, Presence: Teleoperators and Virtual Environments, 1(2), MIT Press, 262-271.

[LAMP74] Lampson, B.W. (1974), "*Protection,*" in Proc Fifth Princetown Symposium on Information Sciences and Systems, Princetown University, March 1971, pp. 437-443, reprinted in Operating Systems Review, 8, 2, January 1974, pp. 18-24.

[LAND81] Landwehr, C. (1981), "*Formal Models of Computer Security*", Computing Surveys, Volume 13, Number 3, September 1981.

[LOO92] Loomis J.M (1992) Distal Attribution and Presence, Presence: Teleoperators and Virtual Environments, 1, winter, MIT Press, 113-119.

[MACE95] Macedonia, M. Zyda, M., Pratt, D., Brutzman, D. and Barham, P. (1995), "*Exploiting Reality with Multicast Groups: A Network Architecture for Large-scale Virtual Environments*", in Proc 1995 IEEE Virtual Reality Annual International Symposium (VRAIS'95), 11-15 March, 1995, RTP, North Carolina.

[PAND96] Pandzic I.G, Capin T.K, Thalmann N.M, Thalmann D (1996) Towards Natural Communication in Networked Collaborative Virtual Environments, Proc. FIVE 96, Pisa, pp37-47.

[RODD96] Rodden, T (1996), "*Populating the Application: A Model of Awareness for Cooperative Applications*", Proc. CSCW'96, pp. 67-76, November 1996, Boston, US, ACM Press.

[SLAT93a] Slater M, Usoh M (1993) Presence in Immersive Virtual Environments, Proceedings of the IEEE Conference - Virtual Reality Annual International Symposium, IEEE Neural Networks Council, Seattle, WA, September, pp90-96.

[SLAT95] Slater M, Usoh M, Steed A (1995) Taking Steps: The Influence of a Walking Technique on Presence in Virtual Reality, In ACM Transactions on Computer-Human Interface, Special Issue on Virtual Reality Software and Technology, 2(3):201-219.

[SLAT93b] Slater M, Usoh M (1993) Body Centred Interaction in Immersive Virtual Environments, Virtual Reality and Artificial Life, University of Geneva, 25-26 November. Also in "Artificial Life and Virtual Reality", Eds. N. Thalmann and D. Thalmann, John Wiley & Sons, England, 1994, pp.125-147.

[SLAT94] Slater M, Usoh M, Steed A (1994) Depth of Presence in Virtual Environments, Presence: Teleoperators and Virtual Environments, MIT Press, 3(2), pp130-144.

[SMIT97] Smith G. "Cooperative Virtual Environments: lessons from 2D multi user interfaces". Computer Supported Collaborative Work 1996 (CSCW '96). Boston, Massachusetts, USA. November 16-20, 1996. Pages 390-398.

[STEE97] Steed A (1997) Efficient Navigation Around Complex Virtual Environments, Proc. VRST 97, Switzerland, pp173-180.

[STOA95] "Virtual Reality on a WIM: Interactive Worlds in Miniature", Richard Stoakley, Matthew Conway, Randy Pausch, ACM CHI'95 Conference on Human Factors in Computing Systems, 1995.