LONG TERM RESEARCH PROJECT 25377

Presence and Representation in Multimedia Art and Electronic Landscapes

.

Document ID	eSCAPE -D1.1
Status	Final
Type	Deliverable
Version	1.0
Date	June, 1998
Editors	Monika Buscher, John Hughes, Jon O'Brien, Tom Rodden
Task	1.1

Project coordinator:

Tom Rodden
Computing Department
University of Lancaster
Lancaster LA1 4YR
United Kingdom
Phone: (+44) 524 593 823
Fax: (+44) 524 593 608
Email: tom@comp.lancs.ac.uk

The eSCAPE project comprises the following institutions:

Swedish Institute for Computer Science (SICS), Stockholm, Sweden
University of Lancaster, Lancaster, United Kingdom (Coordinating Partner)
University of Manchester, Manchester, United Kingdom
ZKM, Germany

Editors of this report:

Monika Buscher
John Hughes
Jon O'Brien
Tom Rodden

ISBN 1-86220-052-1
Lancaster University, 1998
This report is available from http://escape.lancs.ac.uk/.

Table of contents

Chapter Four:
Interaction and Presence in Shared Electronic Environments: fieldwork at ZKM

Chapter Five:
Developing a framework for e-scapes

Chapter Six:
Session support for e-scapes

Chapter Seven:
Open Support for shared spaces based on e-scapes

Chapter Eight:
Using presence to manage
shared e-scape displays

Appendix One:
Conceptualising and Explicating ‘Presence’

Appendix Two:
eSCAPE Art Works : the Commissioning Context

Preface

Jon O’Brien, John Hughes

Lancaster University

The purpose of this Deliverable is methodological not in the sense of presenting a method for constructing electronic landscapes so much as illustrating a way in which the project has explored some of the issues involved in constructing such technologies. Its aim is to present a framework, along with appropriate discussion and illustration, as an approach toward informing the design of what are collectively, if misleadingly, referred to as ‘virtual worlds’. It is methodological in the widest sense of this term in focussing on *how* the study of ‘real world, real time’ social practices can be brought to bear on the design of various types of ‘virtual worlds’. Although this will involve some theoretical discussion, the motivating spirit is very much a practical one, namely, that of meeting the needs of system designers. Like all frameworks, understanding how it may be used, understanding how to apply it, is a matter of learning, with some sensitivity, just what it intends to accomplish, and what it can and cannot do. A framework of the kind that is presented here is intended to focus attention on key aspects of informing the design of ‘virtual spaces’, to retain this rather vague characterisation for the moment.

In the above respects it seeks to contribute to a better understanding of the principles and constraints involved in the design of ‘electronic landscapes’. Although ‘electronic landscapes’, as a species of ‘virtual worlds’, offer the potential of overcoming the constraints to ‘real world’ spaces - such as the ability to transcend distance cheaply, be in more than one place at once, represent information in innovative ways, and so on - it is clear that the design of such ‘landscapes’ needs to be informed by an understanding of the ways in which ‘real world’ social space is constructed and used. This is not an argument for making ‘electronic landscapes’ a replica of ‘real world’ social spaces, nonetheless, in our view there are cogent reasons for gaining a better grasp of the force of the metaphor which, after all, does draw on a familiar way of talking about aspects of social space, namely, the ‘landscape’. Whether electronic or ‘real’ there are some common problems which will need to be addressed, though not necessarily in the same fashion. For one, *navigation* is a common problem. Finding one’s way about a ‘world’, irrespective of whether this is ‘real’ or electronic, has to be provided for and understanding more about the ways in which this is afforded in social space will be informative for the design of ‘electronic landscapes’. Similarly, registering and representing *presence* is a problem common to both types of ‘worlds’, though again the solution to these problems can vary between the two modes. However, and once again, the burden of this report is to argue that knowledge of the ways in which presence is registered and represented in 'real' social spaces can be informative in the design of 'virtual spaces'.

Another dimension to the study is the use of art exhibits. The purpose of this is twofold. The first is to explore at least some of the ways in which the insights of the visual arts can be brought to bear on the design of electronic landscapes. Not merely in the sense of enhancing the aesthetic quality of design – important as this is – but also to see what can be learned from artifacts constructed from an artistic motivation. The second and related purpose is to study the responses to artistically informed electronic artifacts to aid an understanding of facets of the social organisation of space. In a sense, such artifacts can be seen as 'breaching experiments' (Garfinkel, 1967) in that they are often designed to 'violate' taken-for-granted conceptions of space in order to make an artistic statement, and investigating how people respond to such creations can be a valuable clue to understanding the more familiar features of spatiality.

It is the relationship between these two dimensions – the sociological and the artistic – which is the most exploratory and tentative and it is important to be as clear as we can be about how, at the present time, we understand this relationship as far as this project is concerned. Artistic creation and sociological investigation are motivated in very different ways. This is not to say that they cannot be connected but it is to emphasize their different purposes without, we hope, making too much of the distinction. An artistic point of view – and what we have in mind here are the visual arts – has its own standards of adequacy and achievement derived from the traditions of artistic creation. The standards are not necessarily the same within all traditions and, indeed, often the point of some creation may well be to challenge those traditions, but the point being that these standards belong to, and derive from, the traditions and it is crucial to understand these in order to understand what it is a work of art seeks to do and how well it does so. Similarly, a sociological point of view has its own standards of adequacy which are not aesthetic but tied more to what we might construe as those of science, such as rigour, objectivity, empirical adequacy, and so on.[1]

The point we are trying to make here is that we need to be careful not to confuse, or to conflate, the different motivations, though this is not to say that we cannot learn from a juxtaposition of the two points of view. Quite the contrary. Nor is it to suggest, as indicated earlier, that there are always clear lines of separation between the two perspectives. After all, there are disciplined activities – architecture and industrial design being but two prominent examples – where the aesthetic and the practical intersect. However, and hopefully without overdoing the point, an important part of the eSCAPE Project is working through what the connections can be between the two points of view.

[1] Bearing in mind that what constitutes the standards of science, especially as regards sociology, are very much a matter of dispute. Nonetheless, the point being made here stands.

The structure of the deliverable

The eSCAPE Deliverables should be read in conjunction with the material provided within the associated year one eSCAPE project video. Where supporting information is available on the video the symbol shown in figure 1 appears in the margin. Material on the video is ordered in terms of each deliverable and in terms of their presentation within each deliverable.

Figure 1: The Video Material Symbol

As indicated, this Deliverable is an exploration of an approach to the building of electronic landscapes. There are some features which it is important to stress at the outset. The first, and as indicated above, is that the construction of such ‘landscapes’ need to be informed by studies of the ‘real time, real world’ use of spaces. This is the topic of the first chapter, *The Social Organisation of Space.* What this does is outline and illustrate the lineaments of an approach to the study of the social organisation of space and spatiality which informs the ethnographic studies. Its focus is spatiality and everyday life and seeks to identify some key properties of the social organisation of space which would have to be considered in the design of ‘virtual worlds’ while acknowledging that how such properties might be instantiated in ‘virtual worlds’ is an open and researchable question.

The two following chapters *Presence and Representation in Interactive Artworks* (Chapter Two) and *The Legible City* (Chapter Three) can be read as a related pair, with the former providing a rationale for the interactive artworks exhibited and studied at ZKM, whilst the latter offers a detailed account of the historical development of one of the pieces covered in the preceding chapter, Jeffrey Shaw’s *Legible City*. This account explicates the ways in which the artistic work developed in relation to technological capabilities of the multimedia equipment integral to the piece and moves on to outline the projected development of the Legible City as a multi-user in future eSCAPE project work.

The following chapter *Interaction and Presence in Shared Electronic Environments: fieldwork at ZKM* (Chapter Four) should be treated in parallel with the two preceeding chapters, rather than sequentially. It reports on field work done at ZKM observing users interact with a range of art exhibits exploiting shared electronic environments. The aim was to gain a sense of how users of the exhibits ‘resolved’ the practical issues ‘where am I?’ and ‘who is with me?’ The presence of a consideration of the Legible City from a range of perspectives in Chapters Two, Three and Four, reflects a consistency in eSCAPE’s orientation to ‘the artistic’ as an arena for consideration and exploration.

The next four chapters are explorations of the evolving framework which has emerged out of the project thus far. These chapters very much reflect work in progress in the sense that they reconnoitre, as it were, some of the technical implications of the previous discussions. They do, however, move tentatively towards offering a framework for the construction of e-scapes.

Also of key importance to the Deliverable are two appended chapters. The first of which reflects upon phenomenological accounts of presence, augmenting the issues discussed in Chapter One and informing the fieldwork reported in Chapter Four. The second appended chapter reflects upon the eSCAPE commissioning process, emerging from the themes and issues covered in Chapters Two and Three and pointing towards the future work of the project.

Chapter One:
The Social Organisation of Space

John Hughes, Jon O'Brien

Lancaster University

The purpose of this section is to outline some considerations about spatiality as a feature of the social organisation of daily life. As was indicated in the introductory section, the aim is to better understand the prosaic and highly distributed social skills of spatiality in order to inform the design of 'electronic landscapes'. In particular, we focus on the notions of 'presence' and 'representation' as organising principles for directing attention to problems common to both 'real' and 'virtual' worlds. So, the task in this section is to gain a better sense of what these concepts might indicate as features of spatiality in the social organisation of the routine interactions of everyday life.

An important preliminary set of considerations is to understand what is meant by 'everyday life' in this context. While aspects of this notion will be dealt with as the discussion proceeds, the sense in which we use the term refers to the world as experienced by a social actor in the course of going about their daily practical affairs in 'real time'. The contrast here is with a world constructed under the auspices of some theory.[1] That is, it is a focus which begins from describing the world as understood and as used by ordinary persons as a resource for 'going about their daily business'. As indicated, further aspects of this idea will be elaborated during the discussion.

In what follows we have concentrated upon presenting as clear a discussion of the main ideas as we can make it. In so doing we have avoided overburdening the exposition – excepting one or two footnotes - with the qualifications and elaborations that would be appropriate in other academic contexts. There are, it needs to be said, other approaches to space and spatiality that are not considered here, but are covered in detail in Appendix One of this deliverable.

Spatiality and everyday life

A first point to make, and perhaps an obvious one though no less consequential for that, is that when we talk of space in the context of daily social life we are not speaking of an 'emptiness', as it were, but rather of the spatial and extended

[1] This distinction is owed to Phenomenology, particularly the work of Alfred Schutz. See, for example, work discussed in Appendix One of this Deliverable.

character of things, be these people, buildings, desks, topographical features, events, or whatever. That is, we are referring to what might be termed the ‘coordinate distribution of things’. As ordinary social actors we know and presume as a condition of the social organisation of our daily life, that people, objects, places, etc., are distributed spatially. It is a feature that we design for architecturally, aesthetically, economically, spiritually, etc.. As a feature of daily life, spatial arrangements in the sense indicated are intimately connected to temporal organisation though in complex ways. What we have in mind here are the manifold ways in which our sense of spatial arrangements is coupled with our sense of such practical matters as, ‘how long it will take to get to London?’, ‘how long have I got before the last train?’, or ‘what time are we eating?’. The general point is that as far as the social organisation of activities is concerned, space and time are not so much abstractions as wordly and embodied in the activities we do.[1] It is in this sense that we want to explore the phenomena of spatiality.

Hitherto, sociological studies of space have largely been subsumed under the rubric of urban sociology and social geography. Here the concern is mainly with the charting of the geographical distribution of various social characteristics, such as income, industries, social classes, religion, ethnicity, population types, mental disorders, and so on. The pioneering research of the Chicago School, for example, in the 1920s developed this emphasis in ways that have not been much exceeded since.[2] Their ‘ecological approach’ introduced the concept of territoriality into sociological analysis and showed, mainly using the city of Chicago as the source of their data, that urban growth was characterised by nonrandom patterns of settlement, land use, and life styles. For them, the ecological patterns of the city were very much patterns denoting significant variations in the character of social life and the differential possibilities for the development of, for example, natural communities, life chances and moral sensibilities.

Our own approach to spatiality is more interactionist in focus in trying to explicate how sociality is displayed and interwoven as a feature of the ‘real time’ construction of courses of action. One might say that we are conceiving of social organisation and spatial arrangements as coordinate resources in the practical construction of the scenes of daily life.

One way of thinking about spatiality as a feature of social organisation is as an arena, a stage upon which social actors construct their courses of action. This is often the sense conveyed by expressions such as ‘the environment’, ‘surroundings’, ‘territory’, and so on. Spatial arrangements become, as it were,

---

[1] This is not to say, of course, that we cannot form abstractions out of the ordinary experience of space and time. Indeed, it is manifestly the case that we do as is testified by the development of spatial measurement systems as well as systems to measure time. Although such systems have assumed some autonomy in that they possess properties of calculability which do not depend upon social experience, the stimulus for their development was originally located in social practices such as the construction of buildings, the demarcation of land rights, navigation, and more. Our focus is on space and time as resources used in the construction of everyday courses of action.

[2] See, for example, Park (1926), Park and Burgess (1925), Faris and Dunham (1939).

the settings in which social activities of various kinds occur. This view, of course, is consistent with the mundane observation that, within social life, certain spaces, certain places, are strongly tied to particular activities. Thus, and for example, classrooms are spaces organised for teaching; restaurants places organised for eating, as are dining rooms; libraries are spaces where books are stored; roads are spaces for the movement of vehicles, and so on. There is, in other words, a strong sense to the notion that particular spaces and their arrangements are tied to particular social activities. Goffman (1961: 3), for example, comments:

> "Social establishments - institutions in the everyday sense of that term - are places such as rooms, suites of rooms, buildings or plants in which activity of a particular kind regularly goes on....Some establishments, like Grand Central Station, are open to anyone who is decently behaved; others, like the Union League Club of New York or the laboratories at Los Alamos, are felt to be somewhat snippy about who is let in. Some, like shops and post offices, have a few fixed members who provide a service and a continuous flow of members who receive it. Others, like homes and factories, involve a less changing set of participants. Some institutions provide the place for activities from which the individual is felt to draw his social status, however enjoyable or lax these pursuits may be; other institutions, in contrast, provide a place for associations felt to be elective and unserious, calling for a contribution of time left over from more serious demands."

What Goffman is drawing attention to here is, first, the ways in which places, spaces, can become *institutionalised*, that is, spaces in which particular activities are widely understood to take place in; second, that these spaces are variable in their social character. In other words, the social character of spatial arrangements are not happenstance arrangements but ones created, constructed and used in all the variety of the social organisation of activities.

One possible implication of the stage, territorial, arena metaphors - though there are others - is to see spatial arrangements as a 'container', as a set of constraints, on social activities; a conception reinforced, no doubt, by the physicality of many of our spatial arrangements. We cannot walk through walls, travel from A to B takes time, we cannot be in more than one place at a time, etc.; a physicality which is exploited in a myriad of ways to facilitate, impede, constrain, hide, show social activities. The physicality of buildings, for example and as indicated earlier, exploits physical properties as a means of denoting or marking off territories of activities. However, while accepting such features of spatiality our approach is to understand these as they are elements within the shared social understandings pertaining to spatial arrangements. In other words, and to state the point in general terms, our interest in the physicality of spatial arrangements is as a socially relevant feature of everyday social organisation. To stress the point: this is not to discount physicality but to understand it from the point of view of the social organisation of spatiality.

The legibility of spatial arrangements

Attending to the social organisation of spatial arrangements brings out what can be termed the ‘legibility’ of spatial arrangements. This notion points to two related generic features of the social organisation of spatiality:

* spatial arrangements are manifestly visible and constructed for their visibility;
* this visibility is public and widely known

The visibility of social arrangements is a precondition of their sociality. For the ordinary member of society matters to do with spatiality – for example, walking, shopping, displaying intimacy, driving, finding a bathroom, etc. – is not the exercise of some deep mystery open only to adepts, but the practical use of what ‘anyone’ knows about the organisation of the world in which they live. ‘Legibility’ points to the mundane fact about daily life that we are involved in a world which is encountered as a recognisable, reportable, observable, preservable, available state of affairs and that spatial arrangements exhibit a mutual intelligibility for members in social life. In the normal course of daily events we can recognise places where we can catch buses or trains, places where we can eat, places where we can report a crime, places where we can buy groceries, places where we cannot go without invitation, places which are private and which are public, places where we can drive cars, places where we are likely to be able to obtain illicit substances with some safety, and so on throughout the immensely rich ways in which spatial arrangements enter into, and are intimately connected with, a huge variety of social activities.

The notion also emphasises the public character of spatial arrangements in much the same sense - though an analogy which ought not to be taken too far - in which writing is a public medium. It is public in that it makes use of conventional signs and operations of combination which are widely accepted and learned as a condition of literacy. In much the same way, the legibility of spatial arrangements has a conventional character which means that spatial arrangements are overwhelmingly ‘readable’ for what they are and for the activities which take place within them.[1]

Subjectivity and intersubjectivity

At this juncture it is important to clarify a key element of the point of view being elaborated here, namely, the notion of ‘intersubjectivity’ and the contrast we want to make with another ambiguous but widely used concept, ‘subjectivity’.

Within much of social and cognitive science, the term ‘subjective’ typically, and in various ways, refers to the point of view of the individual. The general theoretical interest here is engendered by the contrast, sometimes implicit but

[1] This is not to say that we are always and invariably correct in our ‘reading’. Anyone can make mistakes. Moreover, the conventions of visibility can be used deliberately to conceal such as illegal activities in the way that, for example, during Prohibition the illegal sale of alcohol was fronted in various ways in order to escape the gaze of the law – or at least those who had not been bribed.

more often than not explicit, with the 'objective' world or reality. From here on in matters become complicated, not only by the various ways in which this distinction becomes a matter for theoretical and empirical inquiry in the social and the cognitive sciences, but also because these are very often intertwined with philosophical, particularly epistemological, debates and presuppositions. However, in an effort to keep things as simple and as brief as we can, the general issue is about how unique individuals, using and living within their own subjectivities, develop the capacity to see and know about 'the world out there', the world independent of their perceptual and cognitive capabilities. An important ingredient of this is that widely known fact that individuals can and do differ on what they see and know about this objective world. For example, even our basic perceptual apparatus can mislead us, people can see and hear different things, etc.. In sociology, the fact that individuals can hold different beliefs, have different aspirations, disagree about 'what happened', have different attitudes, different interests, and so on, means that subjectivity is taken as a key element, along with an individual's location within the 'objective' social structure, in the explanation of social action.

An essential presupposition of this kind of viewpoint is that which concerns the objectivity of the 'external world', the world beyond the viewpoint of the individual, the world, as it were, beyond our senses. This is an issue which has been the focus of much of western philosophy for centuries and it is probably fair to say that most of the human science approaches to the issue ultimately depend upon some commitment to an epistemology, such as materialism, to mention but the most prominent. While such a philosophical doctrine might have considerable purchase in the natural sciences – though it is not a necessary commitment for their practise – in the human sciences matters are less straightforward. Materialism, for example, has a number of variants, that of the marxist being the most prominent, but its inheritance of connotations of an external *physical* world means that it sits uncomfortably with the candidate phenomena of the human sciences. Nonetheless, it typifies the urge to want to claim – and so direct inquiries into – that there are factors external to the individual which causally determine the patterns of behaviour. It is this kind of claim which gives emphasis to 'subjectivity' as, variously and among others, the repository of flawed versions of reality, the need for science as a corrective to subjectivity, and itself the outcome of external causal forces.[1]

The approach exemplified here departs in significant ways from the above conception in being constructionist in orientation. That is, it takes the methodological viewpoint that the world experienced by the social actor is an *intersubjectively* constituted world. The term 'intersubjective' is no mere change of vocabulary but indicates a marked difference in the kind of attention given to sociality. It starts from the position that social actors are active constructors of courses of action along with the sense and intelligibility of the world using the

[1] It should go without saying that this is a very truncated account of what are complex and detailed arguments which, especially in sociology, have been the focus of intensive debate since the discipline was founded. See Hughes and Sharrock (1998) for a discussion of many of the arguments.

socially acquired competences necessary to the performances of social activities as we ordinarily know them. Such competences, since social performances involve cooperation and collaboration with others, have to be shared and known in common in much the same sense that the use of language involves knowing, and being able to use, a shared medium. This means that while we acknowledge that as individuals we have our own view and perspective on the world, much of this is under the auspices of a practical presumption of a *reciprocity of perspectives.* In other words, social actors, as a condition of being able to operate in the world 'as we ordinarily know it', act on the principle that the 'world as I see it' will be the 'world as others see it', for all practical purposes. Our experience of the world has an *egological* character in the sense that I experience it as 'my experience' but this is embedded in a presumption about the reciprocity of perspectives. Indeed, the condition of it being 'my experience' is furnished through and by the intelligibility of the shared world known in common.

There is, of course, a great deal more that could be said about these matters but now we want to attend more directly to spatiality, a discussion in which further elements of the approach will, it is hoped, become more apparent.

A world known in common

The legibility of the spatial arrangements reflects, and depends upon, a 'world known in common', a resource which is an aspect of the commonsense knowledge of social actors. The concept of a 'world known in common' is not merely a concept for analysis but a necessary presumption about the way in which social actors orient themselves to the world as an ordinary and mundane feature of their social lives. That is, social actors use this presumption in the practical organisation of their daily lives. It is, to put it another way, a precondition of the very possibility of sociality and social organisation and is exhibited in the most elemental characterisations of social activities. Take, for example, 'driving in traffic'. This is massively and intricately dependent on the presumption of a world known in common; a presumption displayed in knowing which side of the road to drive on when going in a particular direction, that traffic lights are signals and convey instructions which have to be obeyed, that flashing lights on cars indicate the direction in which a driver intends to turn, that other lights on cars are break lights, and so on. These, and more, are recognised and trusted by drivers who are expected to share in the common understandings of the conventions and practices of driving, and through which the orderliness of 'driving in traffic' is produced.

Of course, social actors are not naïve about the fact that the world is not invariably known in common. The notion does not imply that everyone knows the same things nor that the presumption can always be relied upon. After all, and to use the earlier example, those who are not car drivers will only know some of the things involved in 'driving in traffic'. Drivers also know that sometimes some drivers will not obey traffic conventions. However, despite the fact that social actors know that the presumption of a world known in common cannot always be relied upon, it is not possible to dispense with it in the practical conduct of daily life since these depend upon the *reciprocal understandings* of parties to the occasion, whatever it might be.

The presumption of a world known in common also furnishes members with a condition for the intelligibility of the world. Earlier we stressed that the focus of this approach is coming to understand the ways in which the world is given sense and orderliness particularly in respect of its spatial arrangements. The contrast here with orderliness is not so much chaos (or as it is in some other sociological approaches, conflict) as bewilderment. The world makes sense to us. Even ‘disagreement’ presupposes a ‘world known in common’ to be disagreed about. Even ‘puzzlement’ is typically confined to aspects of the world rather than holistically. The presumption of a world known in common is, for members, a very practical matter. They live within a world where the presumption operates as an integral part of the intelligibility of the world and as a resource for the accomplishment of courses of action with others.

Furthermore, we do not learn about a world known in common as a once and for all matter. Indeed, what the presumption of a world known in common points to is not so much what we might call propositional knowledge - though some of what we know could certainly be listed in this way - but to an essential condition for enabling us both to know what we know and to learn what we need to know. For example, members exploit the presupposition of a world known in common about driving not by learning a list of propositions about driving, but under highway conditions.[1] By engaging in the demanding task of driving at speed, maneuvering in heavy traffic, perhaps contending with mechanical difficulty, having to assess the risk that the driver about to be overtaken might do something unexpected, and so on. Exploring the driving world known in common is not something done instead of or as well as driving at speed, coping with mechanical difficulty, dealing with heavy traffic, etc., but is part and parcel of doing these things.

As we have indicated earlier, it is important to note that the ‘world known in common’ does not imply that everyone knows ‘the same things’ or that the knowledge each member possess is identical. Some of this knowledge - Schutz refers to ‘stocks of knowledge’ - are widespread, though constantly open to revision. Some of this knowledge is knowledge of typicalities rather than of specifics and often this is all we need to know. Also, our knowledge is distributed biographically in that there are some things that only we know or have experienced or give our own particular accent to.

An essential ingredient of ‘the world known in common’ are the practical methods we have for ‘finding out’ or adding to our ‘stocks of knowledge’. The world, we might say, is so organised that we can find out about it, investigate it and know about it as practical matters in the conduct of our daily affairs. We do

[1] They do consult and learn the Highway Code, or equivalent national documents, but this is only an aspect of ‘learning to drive’.

not need, as it were, to store everything we know or have experienced. What we do need to retain are the competences which enable us to ‘find out’, ‘recall’, ‘remember’, ‘look up’, ‘make further inquiries about’, ‘test out’, etc. the things that are relevant ‘at the time’.

The social organisation of spatial arrangements and social activities

Inquiry into the social organisation of space is not an inquiry into a set of phenomena which is to be construed as independent of inquiries into the social organisation of activities. As was suggested earlier, from the point of view of coming to understand the social organisation of spatial arrangements this is inseparable from an understanding of the social organisation of activities; the activities which accompany, go along with, take place in, are part of, spatial arrangements. To illustrate, Garfinkel (Garfinkel and Wiley, 1980) notes how people do ‘being together’ in queues as a collaborative affair between queue members and those serving them. For example, a queue in a fast food outlet typically operates on the principle embodied in the question, ‘who’s next?’ This question on the part of the server will normally produce a response by the person whose turn it is. In Garfinkel’s example, there are four juveniles in the queue, so ‘who’s next?’ provides for the sensibility of the response ‘we’re together’. But it is the presence of the social organisation of the queue that makes relevant and sensible the response ‘we’re together’. This description elaborates, that is, orders and organises as a sensible matter, the circumstances it describes and is itself elaborated by them. It is in such ways that participants and others constitute the spatial arrangement of the queue as a social phenomena and do so for practical purposes.[1]

Persons are oriented to spatiality as a visible arena, incessantly taking account of others, of their categorical memberships, of their spatial positioning, their projected trajectories, ownership and the like. Visibility is a crucial feature of the legibility of spatial arrangements. Sudnow (1972), for example, describes the fine subtleties of the ways in which persons, in public spaces, orient to each other and to the fact that they themselves are subject to the possible attentiveness of others. Our own and that of others’ demeanour in public places can be seen as oriented to the fact that we are open to observation - ‘being glanced at’, ‘looked upon’, ‘avoided’ - if only as a condition for successful navigation. In other words, we know, and assume that others know, that our own visual activities are taken account of by those who may be the object of those activities and, as such, is an instantiation of the reciprocity of understanding that is a feature of the world known in common. ‘Looking’ is not just looking but is itself a public display and itself an action for which we may be held accountable by others. There is, we might say, a ‘minimally acceptable looking’ in public places: too extended a look may well be seen as threatening, whereas too little

[1] See Schenkein and Ryave, 1974 for a discussion of related matters involved in ‘walking’.

may not provide us with the information we need.[1] 'Looking' is a normatively regulated activity and which is also used to furnish us with the means of making the world intelligible. The common knowledge of what we have termed the 'minimally acceptable look', to use the example again, serves us as a rule-like formulation for determining 'what is going on'. Non-conformity to the rule may become noticeable and, in being so, precipitate matters for further inquiry: 'why is that person staring at me?', 'does that woman fancy me?', 'does that policeman suspect me of something?', 'have I done something wrong?'.

Co-prescence and co-orientation in public places are in many ways visually constituted. Parties to the space are visually 'doing something' - 'walking', 'hurrying to work', 'window shopping', 'gazing around', 'on their way to the pub', 'visiting the town', etc. – and as part of this visual constitution are spatial arrangements. In the following sections – immediately after the 'cautionary note', we discuss further aspects of the social organisation of spatiality in more detail.

A cautionary note about 'world' and 'reality'

However, it is important to add a qualification as to how we should understand the rather general picture just outlined. Although we have used terms such as 'reality' and 'world' these should not be treated as technical terms or as terms denoting any real entities. They have been used as shorthand expressions for the purposes of giving a general exposition of the main principles of the approach. Members, of course, do not typically orient to the things around them as features of a 'world' or a 'reality' but as what they are: 'cars', 'trees', 'friends', 'Kenny Rodgers', 'Paddy's', etc., etc.. This is important if only to avoid being led astray by what might seem to be terms conveying a misleading sense of something very fundamental going on - as in much of the hype surrounding VR and 'cyberspace'. Although it would be fruitless to try to change what is a widespread practise, we are uneasy about the term 'virtual reality' for example, not least because it can convey the impression that it is about very elemental things which constitute a deep and thoroughgoing shift in our basic conceptions. While sometimes we may forgive hyperbole - 'multiple realities' or 'different worlds' are other candidates - as rhetorical expressions intended to dramatize a point of view, we should be very careful not to be misled by them. Rather, and on the approach being developed here, matters involved in and surrounding expressions such as these should be matters of inquiry not rhetoric. Accordingly, we need to understand the presumption of a world known in common as a feature of the orientation of members as both a very practical matter and as populated with the particularities of what they understand 'the world' to consist in. After all, we learn about 'the world' *from within* our society and our culture and learn it in the course of, and under the conditions of, the conduct of its practical affairs. It is known for the purpose of carrying out those practical

---

[1] The possibilities here need also to take into account contextual factors. Failure to give an extended look between lovers may indicate that something is amiss between them, for example.

affairs. We do not learn about ‘the world’ as a general ‘thing’ but through learning about a great many particular things.

Categorical membership and spatial arrangements[1]

An important aspect of the legibility of spatial arrangements, and their visibility, is what it makes available for membership categorization activities. Following Sacks, these are commonsense categories for referring to persons (Sacks REFS). They are usable as part of members’ commonsense, known in common, knowledge of the social world; that is, used for describing socially organized states of affairs.

Sacks starts from the observation that in everyday life any person can be categorized in a very large variety of ways – certainly in more ways than are typically privileged in sociological accounts which tend to feature ‘class’, ‘gender’, ‘ethnicity’, ‘age’, and one or two others, as generic categories which are ubiquitously causal and without reference to the ways in which categorizations are done as interactional features. By contrast, in everyday life persons not only make use of an extensive variety of categorizations but do so as integral elements in the construction of courses of action. The question is: why is a particular categorization chosen rather than another?

In principle, any person can be described using a very large variety of membership categories - ‘woman’, ‘tennis player’, ‘civil servant’, ‘daughter’, ‘sister’, ‘driver’, ‘Londoner’ etc. – all of which may be correct in the sense of being true of that person(s). What, then, does determine the choice? According to Sacks, members typically select according to a rule of referential adequacy, that is, according to an economy rule which if a single categorization is both minimally adequate and also sufficient identification for a given person set of persons, then use it. Which single categorization is selected is an occasioned matter, that is, contingent upon members orientation to the context in which the categorization activity is taking place.[2] Thus, and for example, in a queue the salient contextual features, ‘head of queue’, ‘second in line’, ‘tail end’ will be foreground categories, while other categories, such as ‘woman’, ‘smart dresser’, etc., will be, at best, background items and frequently disattended to for all practical queuing purposes. Thus, queue categorization, following the economy rule, will have practical adequacy for reference to persons in this context as well as constituting the queue as a queue. In other words, categorisations will be dependent upon, tied to, the activity being done.

However, the economy rule does not exhaust the properties of membership categorization as a procedure for the social description of, in this case but in all other cases of social description, spatial arrangements and activities. Categorizations are, in any culture, conventionally grouped together in what

[1] The discussion which follows owes much to the unpublished paper by Lee and Watson, 1991.

[2] This is a very truncated outline of the extremely subtle and sensitive ways in which categorisations get used in everyday interaction and are related to conversational topic and a host of other activities including insults and ‘put downs’, making jokes, identifications, and more..

Sacks terms, ‘membership categorization devices’, or MCDs. These are collections which members of a culture treat as ‘going together’ in some way under a category title. ‘Mother’, ‘father’, ‘son’, ‘daughter’, for example, ‘go together’ under the MCD ‘family’. These MCDs provide for members’ ordinary use of categories in orderly, that is, in publicly intelligible, ways. Thus, ‘head of queue’, ‘second in line’, ‘tail end’ fall under the collection ‘parties to a queue’.[1]

This conventional grouping of categories into MCDs, provides the consistency rule which operates through, but in addition to, the consistency rule and is a built-in feature. The consistency rule comprises cultural procedures for the co-selection of membership categories and is, accordingly, a relevance rule. In procedural terms, the consistency rule states:

– if a membership category is selected, either visually or conversationally, and then,
– a second or more categories are made available in proximity to the first; and if,
– these categories can be heard or seen as deriving from the same MCD, then;
– see them that way; that is, see them as relevantly co-selected from the same device. If they can be seen as relevantly co-selected, do not see them as randomly selected or unrelated.

While for Sacks this rule operates in conversation, it can also be treated as a viewers’ maxim, that is, as a maxim for making visual sense, for assembling a visual order. Thus, and for example, members’ co-selection of queue categorizations, according to the consistency rule, serves to constitute the queue as a visible order, as a coherent set of visibility arrangements.

We are focussing on queuing here to illustrate the intimate connection between spatial arrangements and social activities and how knowledge of spatial location, particularly in terms of categorization procedures, gets to be legible as, in this case, a queue. There are other features which are worth consideration. The MCD ‘queue’ can be seen as ‘duplicatively organized’ into team-like units. To be a member of a queue is, typically, to belong to that ‘team’ and not to another ‘team’, that is, to any other queue in the vicinity. This duplicative organisation of MCDs often means that there is a proper number of incumbents, though in the case of queues it is more relevant to talk of a minimally adequate set of categories with different rights and responsibilities, such as ‘the head of the queue’ who has different rights and responsibilities from second, third, or n^{th} in line; rights and responsibilities which change as the person moves up the line. The rights and responsibilities, the normatively proper activities and orientations, or states of attention, to the queue, and the category-bound states of knowledge, all change in real time for any given member of the queue until he or she has left the queue.

---

1 Categories can belong to more than one MCD as, for example, ‘baby’ can belong to the device ‘family’ and ‘stage of life’.

The queue can also be seen to have itself device bound predicates. For example, the moral requirement to pay attention to the movement of the queue, to keeping one's place, and to the timely taking of one's turn, is a predicate applicable 'across the board' to all incumbents of the queue regardless of their categories. The queue is a device which advances its members along a categorically-defined turn organisation in order that they may eventually leave it, and quit the 'team'. Thus, referring to the queue as a set of categories is a dynamic not a static description.

What we have hoped to illustrate so far is some aspects of the finely tuned cultural machinery used to constitute the legibility of spatial/social arrangements. This machinery, we want to reiterate, is part of that provided by the practical orientation to a world known in common as a set of competencies for persons finding their 'way around' the spatial arrangements they inhabit or visit. After all, one of the remarkable things about our social competencies is our ability to 'find our way around' in places we have never visited before. And it is to a consideration of this aspect that we now turn.

Legibility, standardisation, and the world known in common

Throughout we have been insisting that the social world is an immensely public world which gains its sense from the shared and known in common knowledge. A further aspect of this is the way the world is organized so it is available to be practically investigated by its members, for example, to 'find their way around'.

One of the features of the social world which is insufficiently regarded by cognitive and subjectivist approaches is that various aspects of the world are intentionally organized in ways that make its features publicly available and usable as a world known in common. Road signs are a classic example of this, as are room numbers, floor indicators in lifts, warning signs, addresses, maps, house numbers, etc.. In other words, using more or less standardised formats to guide members around the spatially distributed social world. These are, if you like, practical representations which 'stand for', indicate, places and sometimes their distribution in commonly understood ways.

Of course, such standardised formats require knowing how to use them and often to use them along with what 'anyone' knows about the activities that take place at these places. What we have in mind in this latter case includes things such as knowing that people live in houses and, hence, will have addresses; knowing that room numbers in a public building will typically follow a 'logical' ordering; knowing the difference between an office building and a dwelling place; knowing driving through a town is likely to be difficult late in the afternoon, and so on.

It involves knowing not so much standardised formats, as in the sense used a moment ago, but also knowing conventional arrangements of places. Knowing, for example, in a market that one will see 'passers-by', 'potential buyers', 'browsers', 'traders', and so on; categories which are commonsensically tied to the kind of place it is. Of course, incumbents of such places, as in any place, can give off other signs as to other categories they might occupy, and this, too, is commonly known. It is knowing how to conduct oneself in a crowded area, what ‘body glosses’ signify, how to ‘keep with the flow’, etc..

‘Representation’ seen as a commonplace feature and resource of everyday life is firmly rooted in the practicalities of the known in common world and the standardised conventions used to indicate, plan, signify, etc. features of that world for members.

Presence

In this section we offer some general remarks on the notion of presence in light of the preceding discussion. These remarks are guided by the presumption that ‘electronic worlds’ will be populated, in which case identifying ‘what others are doing there?’ will be a major consideration for ‘inhabitants’. Such a question can be motivated by a number of concerns: looking for someone, either someone known or of potential interest, navigation, or assistance or just curiosity. What the following discussion is intended to provide is an understanding of the interactional considerations involved in the management of copresence in ‘real world, real time’ public spaces.

One of the ways in which we might consider presence is from the point of view of navigation in public places where the issue is not simply ‘getting from A to B’ but doing so among other persons who are likely to have their own schedules and destinations. A condition of successful navigation is knowing the placement and movement of other occupants of the space. As even casual observations of the most crowded thoroughfares illustrate, persons display a remarkable ability to avoid collisions with others using the thoroughfare. As we reviewed earlier, Sudnow (1972) notes that in the coordination of pedestrian traffic a single-glanced monitoring of others’ actions is not only sufficient but a ‘required procedure’ - or, to use an earlier phrase, ‘ a minimally acceptable looking’. He goes on to argue that persons exhibit, through their actions, an orientation to their appearance, activities, relationships, etc., so that there will be a correspondence between what the portrayal is intended to convey and what is observable in that portrayal to a glancing observer.

Sudnow goes on to identify two constraints on the production of appearances and which operate simultaneously:

- the temporal situation of observation
- the order of information about an action, status, relationship, person, etc., that is intendedly conveyed to an observer

Temporal situation of observation

A feature of many interpersonal settings in public places is that an orientation to an observer or a performer in such settings is an orientation to an ‘unknown’ observer or performer. Accordingly, if ‘a glance’ is to furnish the information that is required, then it is an orientation to ‘anyone’s’ glance. Further, it is a presumption that ‘a glance’ will not only furnish an hypothesis regarding others’ appearances but will ‘yield expectably correct and sufficient, interactionally relevant information’ (Sudnow, 1972: 261). As Sudnow goes on to remark:

> The glance is not regarded by members as an imperfect, poor substitute for preferred extended looking. Rather, it is for members of this society at least, a natural, non-incomplete, normatively governed unit of observation, often a maximally appropriate monitoring procedure, known and known in advance to be sufficient and sufficient for ‘anyman’ for furnishing interactionally relevant information.

However, if glances are routinely and correctly sufficient for a vast range of concerted actions, their timing is important and, moreover, as far as the production of appearances is concerned then their orientation to timing is likely to be a relevant issue. This would be especially the case when ‘glancing’ at ‘courses of action’ which are spatially and temporally extended as it might be in the case of waiting in a queue to get served. Accordingly, to the extent that a glance is capable of furnishing some level of information about courses of action whose temporal extension exceeds that of their observation, then the timing of the glance in the course of such action might well be critical. In which case it is to be expected that the production of appearances will be oriented to the availability of that appearance to a glancing observer, that is, sensitive to its ‘to be glanced at observation’. Further, such sensitivity must be organized in a standardised way so that it operates for ‘strangers’ or ‘anyone’.

It is important to note, however, that there are a number of extended activities where there may be continuous, relatively continuous and expectedly repeated observation; situations in which the extent and manner of observation is uncertain. Sudnow cites the example of cases where a person seeks to provide that an acquaintance whom he/she does not want to greet will be observed by the acquaintance as absorbed in ‘window shopping’ and who will take that activity as a warrant for the person not having noticed his or her presence. Another instance is that of students sitting an examination where the observer, in the person of an invigilator, may merely glance or may look more extendedly but not in any predictable fashion. The ‘structured anxiety’ of such situations at least partly resides in the fact that the scope and duration of possible observations is unpredictable making it difficult to avoid appearing posed or staged, on the one hand, and too mobile on the other. Sudnow suggests that enough mobility is required to ensure that ‘posing’ is not seen as occuring - as, for example, a ‘cover’ for some illicit activity - if more than one glance should occur, but not such a degree of movement that the activities being done become questionable should only a glance occur, and occur at points over which the performer has no control. Excessive movement may convey the impression of an illicit motive.

However, take some commonplace activities where there is no suggestion of any undercover work going on but situations of coordination where single-glance monitoring occurs - ‘walking across the street’, ‘having a conversation’, ‘window shopping’, ‘rushing to a lecture’, ‘looking at a watch’, ‘waiting for someone', etc. - and consider having to produce these in natural settings in which we anticipate only the briefest look by another, and provide that this 'briefest look' will be sufficient for that other to see these activities as 'walking across the street', 'having a conversation', etc.. In the case of 'walking across the street' an orientation to be clearly seen is held by noting the presence of an approaching vehicle. Here the concern for a correspondence between the details of what the performer is doing and what we are seeing at a single glance to be doing, is of paramount concern. The performer's orientation to the possible mere and single glance requires that an appearance be produced that at any point at which the single glance occurs within the observation period, we are immediately clearly seen as 'walking across the street'. In such cases close attention must be given to assuring that all features of the conveyed impression are consistently in accord with the sought after identification. If a walker, in the course of a sequence of actions involved in 'crossing the street', stops to scratch his head before entering the stream of traffic and should a glance occur at that instant, the equivocality and the potential danger is apparent.

By contrast, being seen as 'engaged in a conversation' does not depend, within some limitations, on any particular item of behaviour at any instant in time. During the conversation, persons may scratch their heads, adjust clothing, shift body positions, etc., so long as the primacy of the conversational involvement is conveyed. Being seen to be 'with' another means that appearances must be organized such that whatever else the person might be doing, the 'withness' is observable by 'anyone'. Further, 'being engaged in a conversation' may be available at any of a series of glances directed toward the conversationalists at various points on the course of the conversation. If a glance happens to occur even when neither party is currently speaking, their bodily alignments, facial expressions, distance, etc. will likely be organized, and unselfconsciously, as to allow for the determination of the status of their encounter.

Accordingly, given some situational need for making activity observable, the purpose - to be 'engaged in a conversation', to be 'about to enter the traffic flow', 'rushing to a lecture' - will structure and be structured by the temporal relationship of the producer's activities and the observer's witnessing. Thus, the sufficiency of a glance will partly depend upon the ordering of the activity and its temporal structure. In some natural interactional settings, the momentary 'freezing' of a position, awaiting some indication - usually eye contact, that what has been portrayed has been seen, is regularly found in coordination situations, such as 'crossing a traffic stream'. The frozen pose compensates for the glance-timing's criticalness providing that, within some time span, the action will, in being held constant, be unambiguously available. For interactions where assuring a correct interpretation is critical, special signalling efforts will often be made to insure that a look will occur in the first instance.

Order of Information

Sudnow suggests that there is a need to discriminate between kinds of ‘glances’. The first is a ‘focussed look’ which allows for ‘seeing’ such actions as ‘person crossing the street’, ‘man window shopping’, ‘man and woman having a conversation’, and a ‘look’ which is unmotivated, undirected and relatively unfocussed. He goes on to suggest that some special or situated interest is required before a ‘focussed look’ provides persons with an interest in making out what is happening.

The distinction is relevant to the extent to which ‘glances’ furnish sufficient material for specifying such features of a scene as categories of persons involved, their activities, their appearances, and the like. It is ‘focussed looks’ which, Sudnow argues, furnish the materials for such ‘detailed formulations’. ‘Unfocussed looks’ he regards as peripheral monitoring behaviour which, at best, yield scenic properties that ‘things are normal’, ‘strange’, ‘nothing is happening’, ‘striking’, etc.. Of course, it may be that such ‘unfocussed looks’ precede a more focussed ‘looking at’ which furnishes more information.

Of course, persons may orient to producing an appearance so as to insure that ‘nothing’ gets seen of their activities. ‘Glances’ are interactional phenomena and their occurrence and the degree of ‘interestedness’ they exhibit is a matter that persons seek to control as a normative requirement. To do a ‘glance’ is to show some level of interest in the glanced-at-object and there are various techniques that persons can use to influence the degree of informativeness of that level of interest. For example, by trying to appear ‘not to glance’ but showing some ‘disinterested’ way of visually taking in the environment.

In the traffic example discussed above where ‘posing’ is done to minimise the import of the glance’s timing, an additional interest may be to get a look generated in the first place - that is, to be seen. The interest is in getting a ‘focussed look’ – in this case, by drivers - and, in some settings, this interest may be sufficiently strong and crucial for engaging in special efforts to generate it. The task of ‘getting noticed’ often requires producing an appearance - such as ‘eye catching’ movements - that an unmotivated look will notice and result in a ‘focussed look’.

Appearance and categorisation

Earlier we discussed the legibility of spatial arrangements and the categorization of persons. We can also think of the production of appearances as solving various categorization problems. However, it is also clear that no matter how ‘obvious’ some activity may appear to be by virtue of the appearance a person presents, there are also available numerous alternative formulations for describing what someone is doing. For example, no matter how strenuously someone may seek to display the fact of ‘window shopping’ that person can also be said to be doing an almost infinite range of things. No matter what effort is put in to presenting the appearance of ‘window shopping’ there is no guarantee that this interpretation will be uniquely fitted to that activity. The person may be seen as ‘covering up’, ‘feigning’, ‘spending the day’, ‘getting fresh air’, ‘on a lunch break’, ‘a typical tourist’ and so on, and so on. To the bulk of passers-by the person may simply be a blurred figure.

However, when persons go about the business of constructing an appearance, and do so with a sensitivity to the temporal features of witnessability, they do feel that they are controlling the categorization of their activities for others. But this cannot be a matter of physical appearance alone. Physiological appearance cannot determine among such alternative formulations as ‘working on a paper for a conference’, ‘a hard working academic’, etc. What seems to be at work as the source of relevance is the interactional situation itself. For the driver of a vehicle, ‘pedestrians’ have such courses of action as ‘crossing the street’, ‘not crossing the street’, ‘waiting for the lights to change’, and so on. Persons, that is, formulate the categorical status of others’ actions from the relevancies established by the context in which such formulations matter. In other words, it is inappropriate to speak of temporal constraints on the production of appearances apart from concrete interactional contexts.

Some implications for building Virtual worlds

Although the above is a very brief illustration at best of what our approach to the social organisation of spatiality might yield, it does, nonetheless, provide some suggestive ideas and issues. We stress that at this stage of the project the ideas and issues are very much suggestive even though we mean them to be taken seriously. We hold to the idea that there are things to be learned from studies of the everyday social organisation of space even for the most radical of virtual worlds.

At this point it is worth considering such issues alongside ethnographic work previously undertaken (Bowers, O’Brien, Pycock x2) that focused on both interaction within Collaborative Virtual Environments, and upon the kinds of real world collaborations necessary for the Virtual Environment to come into being.[1] These studies describe the holding of a meeting in VR, and outline the activities of meeting participants within the virtual and real world at multiple sites as the meeting progressed.

The resonance of this previous work with the kinds of issues outlined in this document centre upon the concerted ordering of the virtual world by those who were party to the meeting taking place by way of their mundane, commonsensical interaction skills. The virtual world was ‘worked up’ as an orderly arena for action through the interactional competencies of meeting participants, rather than by purely technological arrangements; the utility of the virtual space and the embodiments within it is additionally *and indispensibly* brought about by the interaction of meeting participants coping with the contingencies associated with a variably ‘trustable’ set of resources for

[1] These documents have been be made available on the BSCW server as background material.

interpreting the actions and utterances of other meeting participants and their 'embodiments' in the virtual world:

> ... the utility of the embodiments is *not* given by technical arrangements alone. Their usefulness as bearers of interactional significance has to be worked up, worked on, maintained and repaired in the light of ongoing activity in both real and virtual worlds.

The kind of skills involved in working up such utility, and in providing these environments with their collaborative nature, were not 'new' ones in the sense that they were, in essence, no different from the kinds of mundane interactional competencies displayed in the social organisation of everyday *real world* life.

> A great deal of excitement surrounds VR and, in some quarters, this has turned into hype and exaggeration. For each published account of a working VR system, we read many more announcements that VR comprises a technological 'brave new world' or that the possibility of 'electronic embodiment' ushers in new forms of human subjectivity. Our research makes a different impression. We see researchers and meeting participants engaged in their day-to-day business, exploiting local knowledge, mobilising everyday skills and competencies, artfully managing contingencies and problems as and when they come up as best they can. We see ordinary interactional competencies (methods for managing turn taking, displaying attentiveness and orienting bodies, using another means if one fails) deployed so as to make less familiar, less trustable arenas for action more recognisable and reliable than they otherwise might be. In all this, the virtual world is but one domain and the management of multiple arenas appears in many ways a normal and unexceptional task and in that sense mundane—which is not to say that it requires no skill, indeed quite the reverse. But the skills it requires are those of appropriately competent researchers and meeting participants, rather than 'cyborgs' or 'postmodern subjectivities'. Activities—be they in the real or virtual worlds remain *worldly*. How could it be otherwise?

Precisely what is being outlined in these studies then, is an account of the virtual world as a 'worldly' place, a place that relies on the kinds of fundamental understandings of 'a world known in common' that we discussed earlier. This virtual world was, then, 'worldly' in the sense that it existed as an *intersubjective* world, the intelligibility of which depended upon its social organisation by members, predicated upon their practical orientation towards it as a 'world known in common'.

These observations, alongside the account of the social organisation of space herein, make plain the central importance of one of the main issues within the VR community - if it is sensible to speak of it as a 'community': the extent to which the parameters of the 'real world' can be relaxed to take advantage of the functionalities provided by an electronic medium. Since communication at a distance is virtually instantaneous, for example, do we need to represent distance in a VR world? What minimal embodiments can we use effectively? How far do we need to reproduce the proxemics of bodily spacings? and so on. At the other extreme would be the ideas of violating most if not all of the parameters of the 'real world'.

However, a crucial issue is how users are to learn a world which violates, or suspends, their taken-for-granted, known in common intersubjectivities about space and social arrangements. From what has been said earlier, it should be reasonably clear how deeply ingrained spatiality and sociality are, which means that suspending any sense of our common understandings is likely to produce bewilderment unless the world is so constructed that it can be investigated and learned. In other words, it must provide for itself as a learnable world in some way.

It should be stressed again that this observation is an important one within the context of the eSCAPE project - to speak of an inhabited virtual ‘world’ is to speak of a public place in at least the minimum sense that the ‘world’ must be learnable by ‘anyone’. To reiterate, in making this contention we are drawing attention to the *intersubjective* nature of the social organisation of space, and are arguing that the most fantastical of environments will require such an underpinning for it to be engaged with in any kind of meaningful sense.

These arguments take on particular significance when we consider the proposed role of the e-scape as ‘the place where places meet’, a ‘learnable’, ‘knowable’ environment that provides links between a range of heterogeneous other virtual worlds: clearly we already see crucial elements of these requirements reflected in the characterisitics of ‘the world known in common’ and we see furthermore that we ignore such principles at our peril. It seems reasonable to contend that an e-scape affording such action must preserve important elements of the principle of intersubjectivity that we outline and explore in this document.

Accordingly, if the e-scape is to fulfil these aims it might be worth considering some of these principles as ‘sensitising concerns’ for those constructing such environments as we consider the technical work undertaken in producing, for example, subjectivity (in a technical sense) in differently presented virtual worlds. Such technical undertakings might be informed, for example, by an outlining of a range of interactional principles that must be afforded by a given environment, in whatever form it might be presented if it is to remain an intersubjectively intelligible arena for interactions. One might also consider the preservation of such a fundamental intersubjective basis when switching between very different worlds and undertaking other such manouevres within the e-scape. The most fantastical of worlds, if they are to be worlds that can afford interaction and collaboration must be underwritten by the kinds of principles we have been outlining.

In essence we see the ethnographic work undertaken by Lancaster within the eSCAPE project as a means of further exploring such notions through the fieldwork at ZKM, contributing to the technical work of the project, whilst also providing the kinds of insights to the artistic work seen as characteristic of the project[1]. These issues, and the technical arrangements arising from them, constitute the main focus for exploration in the final four chapters of this Deliverable.

---

[1] Indeed an interesting set of intellectual relationships are clearly traceable between the orientation of this kind of analytic work to notions of subjectivity and spatiality, and that found in the artistic work. To point this out is in no way to contend that the two share an intellectual perspective, indeed quite the opposite.

Bibliography

Bowers, J., O'Brien, J. and Pycock, J. (1996), Practically Accomplishing Immersion, in *Proceedings of CSCW96*, ACM Press

Faris, R.E.L. and Dunham, H.W. (1939), *Mental Disorders in Urban Areas*, Chicago, University of Chicago Press.

Garfinkel, H. and Wiley, N. (1980), 'Discussion', *Unpublished manu*.

Goffman, E. (1961), *Asylums*, New York, Doubleday and Co.

Hughes, J.A. and Sharrock, W.W. (1998), *The Philosophy of Social Research*, 3rd ed. London, Longmans.

Lee, J.R.E. and Watson, D.R. (1991), 'Interim Report on the Use of Public Space in Urban Areas', *Unpublished. Manu*. Department of Sociology, Victoria University of Manchester.

Park, R.E. and Burgess, E.W. (1925), *The City*, Chicago, University of Chicago Press.

Park, R.E. (1926), 'The Urban Community as a Spatial Pattern and a Moral Order', in E.W. Burgess (ed), *The Urban Community*, Chicago, Chicago University Press.

Ryave, A.L. and Schenkein, J.N. (1974), 'Notes on the Art of Walking', in Turner, R. (ed.), *Ethnomethodology*, Harmondsworth, Penguin, pp. 265-274.

Sudnow, D. (1972), 'Temporal parameters of interpersonal observation', in Sudnow, D. *Studies in Social Interaction*, New York, The Free Press.

Chapter Two: Presence and Representation in Interactive Art Works

Annika Blunk

ZKM

In this chapter we focus on artistic approaches as embodied in interactive artworks exhibited – some of which produced by artists in residence – at ZKM. In particular these are examined for their conceptions of ‘presence‘ and ‘representation’. We follow this chapter with an extended discussion of Jeffrey Shaw’s Legible City, considering its development as a piece of interactive art. This discussion focuses upon the relationship between the development of the artistic work itself and the technological developments that underpin it, and relates to issues highlighted in Deliverable 3.1, contributing to the development of a framework for the construction of heterogeneous e-scapes.

The visual arts have always been technologically rich even if their contribution to disciplined scientific knowledge has largely gone unheralded. Artistic work has always had more of the character of a craft in which the mastery of skill is regarded as more important than the acquisition and application of formal knowledge. Nonetheless, much of this skill – depending, of course, on the visual art concerned – is intimately involved with what we would normally regard as scientific or engineering principles such as, and to mention but a few, the chemistry of colour and texture, visualisation, the plasticity of materials, and constructional principles. It should be no surprise, then, that visual artists have taken to electronic technologies as a medium for the expression of artistic visions and aesthetic experiences.

Most of the works at ZKM use the new media technologies as channels for the communication of artistic ideas. Interactive artworks are designed to involve an audience in the work of art as part of the experience that the work seeks to explore. The entire range of multimedia expressions are used: still and moving images, written and spoken language, sound, as well as a variety of visual effects, exploiting expertise drawn from technology, computer science and robotics. Interactive artworks are aimed at participation and always describe their own presentation space. In order to unfold the exhibit needs at least the participation of one viewer via the interface. The exhibits have a close analogy to the theatre in that each creates its own space for the performance. The role of the actors and the spectators has to be decided and the exhibit itself is like a stage upon which participants perform and ‘play out’ their respective roles.

These characteristics make already the closeness to the theatre clear. Also the theatre needs a space and asks there for a relation of interaction of the participants to decide for the role of the actors and the role of the spectators, since only in the dialectic of playing and watching the theatre reveals itself. However, many more analogies between interactive art and theatre can be determined. The vicinity of the theatre to interactive artworks will serve in the following to determine and to analyse presence and representation in interactive artworks. The theatre serves here as a model for a better understanding of the individual characteristics, concepts and structures of the interactive installations. Also the analogies can serve as measurements in how far the understanding of presence and representation has changed through the usage of innovative media art.

The interest of the eSCAPE Project in interactive artworks is their exploratory and visionary character. Motivated by artistic considerations such works can explore ideas which are unconstrained by the concerns of, say, industry, commerce or other more ‘practical’ interests. In this respect, interactive artworks can extend our understandings of ‘presence’ and ‘representation’.

In what follows ‘presence’ and ‘representation’ are explored from a different but not wholly unconnected approach to that presented in the first section, *The Social Organisation of Space.* The main difference is the reliance on Phenomenology and, thus, is more philosophical in tone.[1] However, since at this stage of the Project the aim is very much exploratory, we see this difference as potentially fruitful rather than an ‘intellectual trouble’.

The Requirements.

Seen from a historical point of view, a good deal of attention has been paid to the concepts of ‘presence’ and ‘representation’. There is an extensive literature in which the numerous attempts of their definition can be found. There is a substantial literature on both concepts, covered in detail in Appendix One of this Deliverable, aspects of which will be expanded in the following discussion. However, the main focus will be on Heidegger’s conception to be found in his *Grundproblematik der Phaenomenologie*

Presence.

What we meet in our daily contact has the character of the unobtrusiveness. We do not continuously explicitly perceive the objects around us within a familiar

[1] The first section is Ethnomethodological in approach and which owes some of its inspiration from the Phenomenology of Alfred Schutz. However, Ethnomethodology sees itself as an empirical discipline not a philosophical one.

environment. Especially since an explicit establishment and insurance of their presence does not occur, we have them around us in a strange manner, the way they are in general. Availability [Zuhandenheit] and unavailability [Abhandenheit] (of the being) define adaptations of a basic phenomena that we formally characterise with ‘being there’ and ‘being absent’ and more generally as presence(Heidegger 1975) Heidegger defines presence within time, absence, according to Heidegger, is a modification of presence. So only a time reference allows the existence, the being-in-the-world to be transcended and makes its intentionality possible. Presence is finally something what we perceive in the current moment, what is there - or what we do not perceive but what we miss.

Representation.

To represent [vor-stellen] means to bring what is present at hand [das Vorhandene] before oneself as something standing against, to relate it to oneself, to the one representing it, and to force it back into this relationship to oneself as the normative realm. Wherever this happens, man gets into the picture in precedence over what is. But in that man puts himself into the picture in this way, he puts himself into the scene, that is, into the open sphere of that which whatever is must henceforth set itself forth, must present itself [sich praesentieren], be a picture. Man becomes the representative [der Repraesentant] of that which is, in the sense of that which has the character of objects (Heidegger, 1975). Consequently representation defines itself only by putting oneself in relation to the represented and by this is necessarily subjugated to a temporary subjectivity. If the term re-presentation is used for an optical image this would mean for Heidegger that something having been there is recalled via a picture. Establishing a reference between oneself and the object of the representation is updated.

Presence and Representation in Interactive Installations. A Theatrical Approach.

Interactive installations reveal themselves in the dialectics of action and reaction; the interaction requires the willingness of the spectator to get involved in the play and, in the beginning, to obey the dictats of the instructions. Action and reaction are the two offers for participation that are given with interactive installations. The individual artworks will be examined for their architecture, the conditions of perception and their presentation forms, the hybrid forms of multimedia. The analysis will show that all components of the artworks are present at the moment of action, completely independent from the reference established by the design of the space.

The space itself can be either composed as a black box, a room delimiting the surrounding area, or as an environment, which the artist has designed to match the virtual space. During the plot, real space and virtual space, objects, actors and spectators are equally present. Emphases are made by the artists' individual conception which can evaluate the presence of individual elements differently. For example, selectively positioned light can focus perception and concentration on specific components. Also the presence of the spectator changes when he changes from a passive role to an active role; he is not more or less present but rather on a different level, he is not anonymous anymore but an explicit participant of the action.

The representation repertory determined in the work is of an entirely synthetic nature. These forms of representation describe less the phenomenal side of the real. Not an object precedes a picture but a model, a formalised description, an interrelation of the real that is literally or visually elaborated and filtered. And through this, the presence of the artist is included in his presentations. In *The Legible City* the architectures of the cities are represented by the text, Perry Hoberman chooses everyday objects as agents for the participating viewers, in *ConFIGURING the CAVE* the various conceptions of space are represented in the various simulations. Bill Seaman's objects most obviously show what share of the representation is due to the artist and what is due to the active viewer. The repertory, which Seaman offers, consists of objects, which can be chosen and modified. The acting spectator is able to charge them in light of his own interpretations of the real within the virtual world. However, these representations must necessarily be temporary. Like this not only the things latently present within the plot are being held present, but merges with the elements of representation incorporated by the artist in the temporary representation of the acting viewer, his interest, his culture and his individual world of imagination.

Jeffrey Shaw: The Legible City (1989-91)

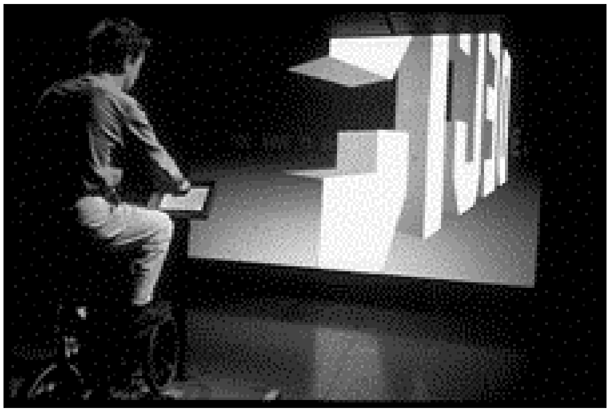

An Objective Observation.

A black box from the outside that does not reveal anything about its interior. If one enters the dark room one is confronted with a slightly curved nearly 16m² projection screen. In front of this a firmly installed bicycle is mounted. A flat

LCD monitor is attached to the handlebars. The projection screen displays large three-dimensional letters strung together to words and positioned on a dark ground. The bicycle invites the viewer to mount it, and once he/she has taken a seat the eyes fall on the LCD monitor and an apparatus rather like a bicycle gear. The monitor displays a ground plan, while the ‘gears’ with the push-buttons function as subtitles: Manhattan, Amsterdam, Karlsruhe. And if the viewer is familiar with the maps of these cities, he is easily able to assign the three maps to the metropolises, which alternatively pop up when the button is pushed. Pushing the button also causes a change of the projection: the size and colour and position of the letters are altered. Once these functions are clear, the viewer begins to pedal and the projection to transmute. Soon the viewer realises, that the map corresponds to the projected three dimensional space, though the arrangement of buildings according to which the map is laid out has been replaced by chains of words. The actual position in relation to the projected database is indicated by a pointer on the display. However, the letters conform to the actual scales and proportions, at least in Amsterdam and Karlsruhe. With the handlebars, the viewer is able to steer through the projected letter-landscape. And in each of the three cities the reading-journey makes sense, once the viewer has let himself into the work and follows the displayed streets. Each journey gives - maybe a new identity, though definitely for the bicyclist a very individual - view of the cities.

And each of the virtual cityscapes has its own characteristics. The pedalling and attentive viewer realises in Manhattan, for example, that there are eight different ‘tours’, which can be followed from 66th to 34th Street between Park and 11th Avenue. Like on hiking, skiing or cycle tours, different colours indicate the various paths. In Manhattan the paths are monologues, views of eight characters whose paths of life are only randomly crossing: the ex-major Koch, the billionaire Donald Trump, the architect Frank Lloyd Wright, the editor of the Webster-Dictionary, Noah Webster, as well as an unnamed tour guide, a confidence trickster, an ambassador and a taxi-driver. Though the texts are fictitious, all eight characters turn to the virtual visitor on the bike, offering him their individual view on the places the cycling person is moving through. Despite their distinction, the unifying characteristic due to their professions is that they mediate the city.

The virtual Manhattan

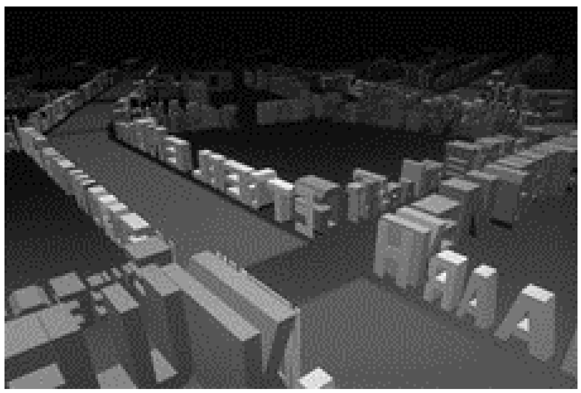

The virtual Amsterdam

If the viewer switches to Amsterdam, the words transform into the Dutch language and its colouring corresponds to the characteristic canals and brick buildings. Here the bicycling reader is touring through the history of Amsterdam, touring from the 15th to the 19th century. The locations of the various historical actions coincide with the textual visualisation map, as well as the chosen language corresponds to the historical moments - contents and expressions were taken from archive documents. His starting point (in front of the central station) assures him that he actually is ‘in’ Amsterdam: he is confronted at first sight with the city’s name in its old spelling.

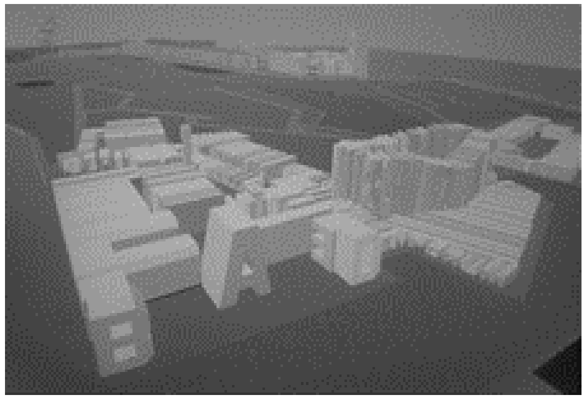

The Virtual Karlsruhe

The Karlsruhe version is comparable to ‘Amsterdam’. Also here historical archive material served as the basis, and historical scenes and events can be ‘accessed’ with the bicycle, though here the texts are displayed in German. And in this city the starting-point is the actual starting-point of the town itself, in front of the castle facing all the streets that lead towards it.

The work is an interactive and hypermedia work and soon the viewer realises that although the text is displayed in a linear way, he is able to move freely within the three dimensional database, in other words, not according the actual layout of the streets. The individual trip always also means composing an arbitrary text.

The interface of the work is the bicycle itself. The handlebars and pedals are linked to a computer via electronic sensors. The projected image reacts in real-time to the ‘commands’ regarding direction and speed as a result of the viewer’s actions. The bicycle demands from the viewer physical effort in order to move through a universe without any resistance. To make a city legible, he must adjust his speed. The map in front of the bicycle serves for better orientation, rather like in real space. Indicated by a small moving dot, it reminds the viewer where he is moving virtually in the real space. Since the map refers to the real location as well as to the virtual, it keeps the real city present through associations - visually as an abstraction. Like this the LCD-displayed map becomes the true interface between real and virtual space, and the viewer moves in-between and in both.

Once the viewer is involved in this journey of reading, into this play, he will very soon realise that each row of houses builds a line in a narration which is meant to be read while strolling along it. Inevitably one feels reminded of classic modernism, reminded of exactly that figure that became a topos some time ago (See for example Mitchell (1995), Taylor & Saarinen (1994), Glaser (1995), Roetzer (1995)). The cult figure of the flaneur, as Benjamin has called him, has once again become a cult figure.

In order to explore the virtual space in *The Legible City*, the viewer in the real space has to use his physical strength even though he actually is not moving at all. But it is in that moment - the moment of immersion - in which the borders between real and virtual space are blurred and are threatened to dissolve, the real locomotion merges with the electronic flanerie. The model for subjective experience of an unstructured navigation within the structured hypermedia or in virtual metropolis seems to be found in the aesthetic model of Baudelaire's flaneur from the 19th century. However, as banal and ordinary the actions of the flaneur may appear - it is the flaneur after all who transmutes the ordinary of a city into something extraordinary, who transmutes the banal into something aural. The message of the flaneur is that he sees what everybody else sees as well, but much clearer, truer and better. And is this not what actually happens in *The Legible City*? The flaneur cycles along those familiar streets and discovers 'another' face of the space, a face that lies beyond the facades of the architectures: the fourth dimension, the voices, the movements, the time passed.

The Legible City: Seen from a theatrical Point of View.

The Participants (I): Viewer/Actor.

To leave the 'real' space behind, to enter a 'Twilight Zone' and begin from here a journey through virtual space. In this piece the artist does not impose an 'Aha-effect', rather the layout strives for recognisable familiarity. The viewer need not experiment before finding out how this virtual reality can be entered, an environment that is displayed promisingly on the frontal projection screen, which stretches from the ceiling to the floor. Intuitively he knows how to operate the interface. Since the viewer does not need to concentrate on the interface, he rapidly leaves the environment that surrounds him behind. His actions are entirely motivated by the virtual space, and he acts as on a bicycle in the real world, because the parameters are nearly the same: you are cycling on a road, gravity cannot be surpassed and perspective corresponds to the position of the bicycle in the virtual space at all times. And even though the viewer is able to drive through the rows of houses, any other co-ordinate within the space can only be reached by cycling there, through physical strains. The fascination of the newly discovered space derives from the texts that substitute the architecture whilst being part of the city at the same time (as its history or as its personality). While the interpretation-work asks the reader to 'read in-between the lines', the reader of *The Legible City* is actually invited in-between the lines. And the viewers climb down in-between historical and authentic lines and by moving within the text, the narrative components become part of their experiences. Since each viewer in a specific sense actively participates in the writing process of the 'history' it becomes even more part of his experience. New meaning is brought about with each individual path.

Active viewer riding through Manhattan.

The immersive impression breaks up in those moments, when a new cityscape is chosen. For one moment the viewer clearly realises that he is moving in front of a projection screen and that he is the one who chooses the projected contents via a push button. The artist has created a frame of action for the navigating viewer, within which he can freely move according to the mechanisms of navigation. Once the artistic concept is understood the active viewer is able to develop his personal choreography, is actor and director at the same time. He creates sense by riding the bike and atmosphere is established through his driving style.

The Participants (II): Spectator.

However, the active viewer is able to discard the interacting role or exchange it with the observer's perspective whenever he wants to. However his advantage over all other spectators are all those experiences that he made while he was interacting. Because of this he does not disappear in the audience but rather takes on a special role amongst the spectators. For all observing people in the environment, the interactive installation turns into a cinematographic event. They perceive a film that one viewer brought into motion, and by this makes the absent space present. The isolation from the outside world makes it easier for the viewers to participate in the play and to follow the dictation of the projected images: the typical cinema conditions. The rhythm of the interacting viewer reduces the audience's action to watching. He transforms them into perception apparatuses: if they want to follow the 'story', they have to acquire the speed of the projection and to keep it up.

For the initially observing viewer the actor appears nearly sculptural in this otherwise empty black box and by this is the main attraction for any newly entering spectator. The spectators on the other hand are able to collect passive experiences by observing the actions: they see how the movements of the images are triggered, how the plot of the projected narration is decided upon and they consciously perceive the actor as protagonist and presenter.

The actions between 'audience space' and 'stage' continuously creating new constellations, letting a new scene begin. Dramaturgically this is expressed by the changing directions the plot receives with every new actor. The situation can change at the moment the actor becomes a member of the audience and starts commenting: he takes over a function that comes close to the chorus in classic theatre. This in return can influence the role of the actor insofar as he pursues the direction suggested by the audience; the actor becomes its agent. Increasing knowledge divides the spectators into various groups of perception. Alliances are concluded, they change from the passive observer, to an adviser, to an active director.

The Legible City: Interaction in a Black Box

The frame of action defined by Shaw consists of the bicycle and the virtual cityscapes, which are projected on the screen within a black room excluding the outside world. Even though the room itself does not resemble a stage area, on entering the space one immediately knows what defines the space for action and the space for the audience. The illusionistic image space, the three cityscapes, are also easily identified by the active viewers and the spectators. The theatrical and dramaturgical moment is developed when the image space is becoming pure actuality, when the actor defines the space as his acting space, when he climbs on the bicycle and the ride, and by this the action, starts. Doing this the actor uses the entire repertory offered by the artist. Within the image space he discovers the unknown less in discovering the simulated space but rather in exploring the simulated content. The artist has to rely on the spectator's experience to fix representational content, the integration of a linguistic level in a figurative way protects the work against idiosyncratic experiences that actual spectators may have within the environment through deficiencies or peculiarities of knowledge, sensibility or biography. The total immersion of the viewer and the spectators in the virtual world is the aim of the artist, and once actor and spectator get involved in the play and have 'lost' themselves in the virtual space, the definition of the theatre space is lost. Real world and virtual world are not standing separately opposite each other anymore but the virtual is becoming part of the real.

Real-Time.

The interplay of the parameters space, time and action - essential for classic theatre - is also relevant for Shaw's *Legible City*. While the virtual databases of Manhattan, Amsterdam and Karlsruhe describe the space for action, the real time structure is assumed in the virtual space. This component contributes to maintaining the immersive impression. In real-time the virtual space is kept present. Physical strains, speed and time passed correspond to their own everyday experiences. The unknown and new has been consciously reduced in favour of the immersive impression. Operating within the virtual space is simplified by adapting everyday experiences and soon one moves there as here, the own acting structures are corresponding. The time frame of the plot is again defined by the actor. In real-time he joins the action sequences together to a linear whole, determines like this the dramaturgy of the installation.

Action: the Narrative Element.

The image space, as could be established, is always defined by one of the three literal cityscapes. For this Shaw uses letters and words figuratively – they are never present only for formal reasons. Indeed a text piece seems to serve as basis for the action. But now the text as such is held present while the interpretation work has to be done by the actor in his individual handling of the text as action space, the historical and narrative elements of the architectures. The narrative paths read as unified plots represent information packages on the history of and lifestyle in the individual cities. However, they do not have to be read as such. Even though linearity seems to be established, the work is not laid out to the maintenance of this linearity. If this had been intended by the artist, then the modelled letters could have been defined as solid, so the viewer could not drive 'through' them. Instead the viewer is requested to compose the text himself. What Walter Benjamin under the pressure and impression of new media in the 1920s assumed, namely that the reader himself is on the verge of turning into a writer, seems to be achieved in this digital city. It is said that in digital media writing and reading functions are integrated to the extent to which one is able to switch without transition from one to the other. According to Wooley (1994)

> "In cyberspace everyone is an author, which means no one is an author; the distinction from the reader disappears. Exit author…"

According to this author and reader cannot be perceived as separated anymore but only as one. Also for Norbert Bolz the question, 'What is an author?' resolves in 'Docuverse' where everything written is entering the databases and can be re-used by the writers. This is where unauthorised texts will emerge from, namely texts without authors just as if they are written while reading (Bolz, 1993). Within complex structures such as the Internet, for example, no single author can any longer determine which path the reader has to choose. From each position there are always several possibilities to continue the narration. The reader decides which possibility will be chosen.

Though here are three cities offered, at least for some viewers the text-based representation of history stays in obscurity. The letters representing an unfamiliar language, can be perceived as objects and as such they are purely visual arguments. And at this very moment the representational function of language becomes obvious since this representation system follows specific fixed rules and conventions. For a given image represents what a symbol system decrees it does, no matter how the spectator experiences it and no matter how the artist intended it. And this system will be inaccessible for the viewer until he has broadened his knowledge in this respect. Until then the plot is reduced to the purely visual argumentation of an original production.

Bill Seaman: The World Generator – The Engine of Desire (since 1996)

Describing the Visual Space.

A large projection screen in a darkened room. In front of it at some distance a table with navigation tools, a chair inviting to take a seat. The projection screen is parted into two sections: an upper blue zone and a lower black one. Within the black zone an elongated rectangle is located that is subdivided into individual fields for words and icons.

The viewer is confronted with three different tools: two push-buttons, a trackball and a *spacemaster* (a kind of three-dimensional mouse). By pressing the upper push-button, the rectangle which serves as a virtual menu bar, can be faded in or out as well as it can be enlarged for clearer visibility. The function of the second push-button is to select the options offered by the menu.

The trackball allows the viewer to navigate within the menu in horizontal as well as in vertical direction. A light red frame indicates the actual position. By moving the trackball in the vertical direction, the fields on the selected y-axis move up or down respectively. The amount of the selectable options within one axis varies within the eight selection possibilities and each string of selections forms a loop. In five of these compartments a total of 62 objects are arranged. If the viewer selects one of the objects with the button, it will appear as a three dimensional object in the empty zone above the eight wheels. - The two dimensional abstract surface becomes a three dimensional (virtual) space by integrating the first object.

Set-up of the installation THE WORLD GENERATOR.

The *spacemaster*, the third tool offered on the table, is the true navigation tool: it facilitates movement around the centralised object, as well as facilitating movement far away from it. A black push-button on the *spacemaster* activates a light green sphere that surrounds the central object. The artist calls this sphere the ‘aura’. The viewer is able to modify the enclosed object that he previously placed into the space according to his own ideas using the options offered by the menu.

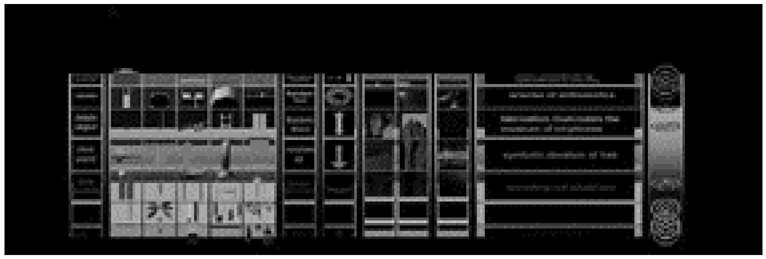

The eight wheels of the virtual interface.

While learning how to handle the complex navigation-system the virtual space permanently changes: as a starting-point serves a surface that presents itself as a combination between abstract surface and a virtual interface (and alternatively can either be interface or abstract surface); the trackball transform the interface into a three-dimensional object, to a spool whose separate wheels have different width; by selecting an option the surface is converted into a three-dimensional space; finally the *spacemaster* facilitates movement within this only recently discovered space, expands it eternally.

Once these spatial basic principles are understood and the various possibilities to navigate are tried and tested, the true interaction - exploring and conquering the artificial spaces - can be started. And consequently the viewer starts with the virtual interface. He is offered a total of thirteen optional reels that are divided into eight different categories. All wheels moves at the same speed when activated, regardless of their circumference. If one classified the wheels by titles, the first wheel would be named ‘Command Options’. Keywords indicate the functions of each field, though most of the commands will only be understood once the possibilities of the all wheels have become clear. The next wheels shows 62 3D objects. These objects chosen and modelled by the artist will be placed into the empty world as white objects when they are selected from the wheels. The following five option-reels can be understood as expansion or manipulation of the objects. ‘Generators’, ‘Behaviours’, ‘Textures’ as still and moving images, animate the static 3D object. And here the possibilities seem to be unlimited: a choice of ten generators and eleven different ‘behaviours’ could theoretically be assigned to each object (and each text fragment). These options, however, can only be activated but not fine tuned: an object can be enlarged (or put back to its original size), an object can spiral around its own axis (or be set back to immobility). The following two compartments, which are covered with textures in form of still images, are the longest reels within the entire interface: 370 textures are represented by the fields and can be mapped onto the objects. The same can be done with the sixteen filmloops, which can be found on the 6th wheel. Text-fragments are displayed on the next wheel. If a text-fragment is activated, it turns into a textual 3D object - or rather 4D object, since it is inseparably connected to its pronunciation which is repeated in slow loops. A total of 317 text-fragments can be positioned into the space and modified with the help of the various manipulation tools.

Sound can also be added to the objects, as well as to the textures and films that can be positioned as a two-dimensional surfaces in the space. 118 sound files are visualised by circular icons, which become audible when they are highlighted with the trackball. These files can be used as features as well as objects. As an object the audio-file can be positioned invisibly into the space and is only audible when it is approached during the navigation through the virtual image world. Volume and direction out of which the sound-object is perceived depends on the distance as well as the direction from which the viewer approaches the file. As qualities the sound-files function like the textures: they wrap themselves around the object; and again distance and direction are substantial.

Apart from the options offered by the first wheel and the ‘Generators’, selecting a field means activating in the sense of switching a function on or off. Not so the ‘Command’-fields. A total of ten options are offered here which are fundamentally different in their function and effect. ‘Clear World’ or ‘Random World’ generate decisions made by the viewer that influence the entire system grown until that specific moment: ‘Clear World’ changes the projection screen back to its two-dimensional initial position. ‘Random World’ randomly develops a 3D world based on all parameters.[1] Additional functions allow the fine-tuning of individual parameters. Also the option ‘Delete Object’ refers to the manipulation of a single object[2]. The options ‘Hiding Sound’, ‘Showing Sound’, ‘Centre’ - are commands that do not affect the 3D world but serves the viewer for better orientation. While ‘Showing Sound’ constitutes a kind of filter that visualises all sound-files as those symbols, which represent them on the 12th wheel, the option ‘Centre World’ sets the perspective back to the initial position.

The World Generator: Analysed from a Theatrical View Point

The apparatus and instruments for designing the space of action, which are put at the directing person’s disposal by the artist, seem to be immense. This offer opens up a seemingly endless variation span of setting up and presenting the virtual space with the various media. While in *The Legible City* the virtual space was defined by the artist himself, it is now a question of developing architectural strategies in order to simulate a non-existing space.

1 This function is altered additionally by the entire *generator*-wheel. Here selected random processes are activated: 3D objects, texts, films, sounds only will be distributed within the space.

2 Delete Object can be applied not only on actual objects but can be applied as well on text-fragments, still and moving images, which have been integrated as objects into the virtual space.

Constructor/Director/Actor.

The step from spectator to actor in this installation is initially not as easy to comprehend as it is in Shaw's installation. If the active viewer sees himself confronted with an abstract image space, his first task is to construct a virtual environment before the action can be enacted. For constructing the space, the viewer has to make use of both interfaces, the real and the virtual one: the chosen steps are translated; direct access to the simulated world is not yet permitted. The constructing viewer perceives the virtual space from a distance, from behind his navigation tools, the real space becomes his reflection area. The *spacemaster* is the 'key' to the virtual world, once the user enters it; his constructing tasks have come to a provisional end. He moves through his constellation of the space, inspecting his own construction. Like an explorer he is now able to discover the new continent, conquer it, modify it. With each modification and implementation the viewer is thrown back into his own, real environment. However, the frame of action has been shifted: while initially the space was constructed through a window, the virtual space has now become the personal one after the first reconnaissance. Now the active viewer does not construct anymore he sets up a scene. He now switches back and forth between the role as a director and as an actor. In those moments when he navigates through the virtual space, and defines the action through his navigation, sets up links between the objects and develops a linear plot, in those moments the actor in his own production.

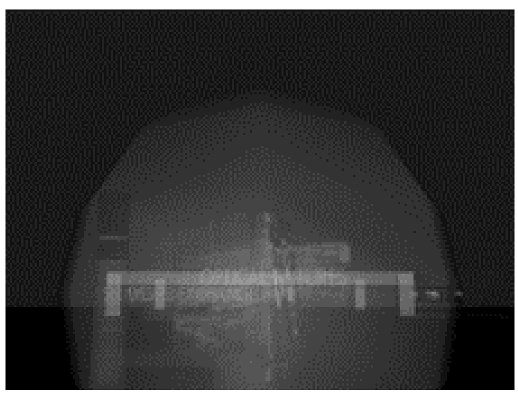

Objects can be modified as long as they are surrounded by the green sphere.

The Audience.

If additional viewers are also in the room beside the active viewer, they are handed over the role of the audience and define their space as the auditorium. However while they were able to perceive *The Legible City* as a filmic narration, they are here confronted with the projection as a variable space for performing. They observe the construction and 'staging' of the virtual space, they perceive the current moment as an improvisorium. And also here the roles can be alternated, one can collaborate with the actor, can direct the scene together with him. Yet, it is impossible to re-establish a moment staged by one actor, since the action space is not static but continuously updated.

The Space for Action.

Just like *The Legible City,* Seaman's *World Generator* is an interactive installation bound to space, is even conceivable within the same dimensions. Abstractly seen both works consist of the same perceivable components: dark room, projection screen, interface - moreover this piece is structured to exclude the surrounding environment. However, unlike Shaw's piece, here a stage for improvisation is included - that empty image space which has to be filled before one can act. There is no uniform stage to be played on, rather the actor arranges his action space, the room is malleable.

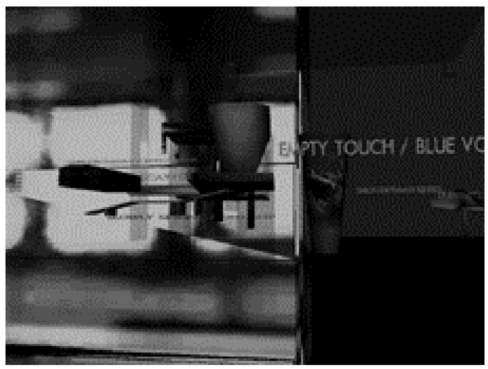

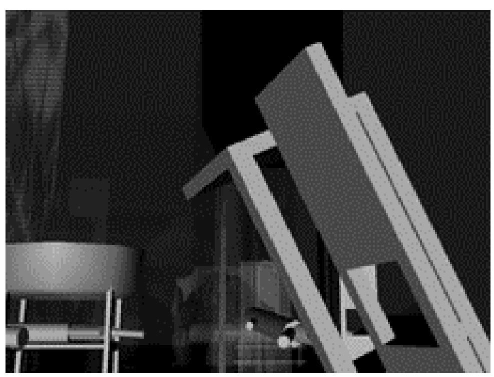

Screenshots from virtual worlds of THE WORLD GENERATOR.

But just how malleable is this room? The wheels serve as Seaman's archive - here the content, his vocabulary, is stored. The viewer is able to choose, to position, to modify the appearance of the props, can even make them move, but as on a stage, as in a museum, each of the objects, each exhibit stays alone and independent. Relationships are explicitly defined by the viewer or curator - or the actor and viewer respectively - either by positioning the objects in relation to each other or through movement within the space. Furthermore the improvising person also decides on the duration of the space's existence and the meaning conveyed through his action, which at a maximum corresponds to the duration of his interaction.

The character of the action space as an improvisation space is additionally emphasised by the fact that the various constructed worlds cannot be stored, not documented. In this piece, the space is not to the fore, but the provisional, the interaction and navigation. The conserving moment is the moment of its emergence, and each moment is destroyed by each additional move, each further interaction. The stage as a playground for the desires of life. Exactly here it is not intended to freeze, not to preserve from death. Instead, this work documents in the temporary production a provisorium of culture, which can only be enjoyed while navigating through it, while one acts within it. The requisitorial character of the objects is additionally emphasised by them staying independent, not being able to react to each other. They are carriers of information in which the viewer can store additional information. Seen from this perspective, *The World Generator* can be defined as a mirror of the culture, since provisionally and temporarily the culture of the viewer is presented.

Once the space is constructed and serves the actor as an activity space he is subjected to the rules of his virtual space: in order to move from one place to another one has to move linearly. At the same time the improvised does not conflict with the real since the rules seem to be aligned. The actions that are performed in the virtual space are perceived in real-time by the spectators. The virtual improvisation becomes pure actuality, exists only in its authentic execution as current situation in the dialectics of playing and watching. And if the action only exists in the respective time of its course, in its actuality, then this time structure is also compelling for the sensory perception of the recipient, is the time structure of his experience. Like this the change between the roles ”actor” and ”spectator” is simple since the limitations of time seem to correspond; the passive viewer is has practised the action, is participating in the plot. Once he got involved in the actor’s offer to play, he includes the own role within the given situation of interaction and understands himself as an element of its dialectics. In this way he is prepared and can act.

The characterisation of the ‘play’ in *The World Generator*, as generally within interactive installations, does not happen as a transformation of a text. The spectator is confronted with a scene within which the events take place on two different levels. On the one hand he observes the interaction between active viewer, interface and projection screen, observes the positioning of the visual arguments within the installation. At the point at which the action does not consist of the permanent alternation of the image planes anymore, but instead defines itself by the linear movement of the actor in the space, concentration shifts into the projected image space. If the actors change so too concentration of the spectator shifts and the action takes on a new direction, the action space gets new spectators. The duration of the plot is defined by the actors as well as by the spectators. Neither a definite beginning nor ending exist, only individual sequences that are developed by the ‘members of the cast’. For the individual the action ends at its latest when he leaves the room, the separate lines of action come to an end when the actors change. The action always ends when there is no actor; when the objects evaporate from the screen, when that which exists is deleted

The action that the actor develops within the imagined space is - similar as in Shaw's *Legible City* - determined by the movement in space, though here the legs do not have to move. The increasing contraction of the dimensions of the straight lines of the real geometry whose width can only be surpassed when motorised, the individual is accelerated, adapted to linearity. The acceleration of the body in motorised capsules leads, as Paul Virilio has stated, finally to absolute standstill. In this standstill the actor of *The World Generator* has arranged himself. He is moving within a world that abandons the building of streets since they have become electronic thresholds. The seated actor imagines total moveability and dynamic and by this absolute power over the space. Virilio (1978) has called this a ”sclerosis of behaviour. Since only mobility and motility of the body", he says,

> ”supply perception with that wealth that is indispensable for the establishment of the ‘I’. To slow down or even completely remove this dynamic of the locomotion, to fix behaviour and motion in the extreme will lead to serious disorders of reality competence."

In *The World Generator* the actor is allowed to repress this successfully, exactly because the space becomes playground and on this raised level it should become a space for reflection.

Perry Hoberman: Bar Code Hotel (1994)

A Descriptive Introduction.

On entering Perry Hoberman's *Bar Code Hotel* one is located in a room in which seemingly every smooth surface is covered with printed black and white bar code symbols: tables, walls, floor. One of the four walls is a projection screen on which in a stereoscopic view an empty room is projected reminiscent of the aesthetics of Rene Magritte. Above the tables five bar code wands are hanging down from the ceiling. On the table there are several polarising glasses as well as numerous white cubes onto which bar codes and an object indication is printed.

Set-up of Perry Hoberman's BAR CODE HOTEL.

If the viewer wants to check into the *Bar Code Hotel*, he has to pick up one of the bar code wands. With each wand the bar code information can be scanned and is immediately transmitted to the computer system. As soon as the information from one of the cubes has been scanned in, it appears as a three-dimensional object on the projection screen. If with the remaining wands other cubes are scanned, the projected 3D world will soon be filled with various computer generated objects, each corresponding to one of the guests. Each agent is provided with a specific sound and can now be animated by its human collaborator using all those bar codes that are distributed all over the room. Each bar code is labelled, either graphically or verbally, in order to help the viewer to orient himself among the black and white stripes, to differentiate them and to get an idea what reaction will result.

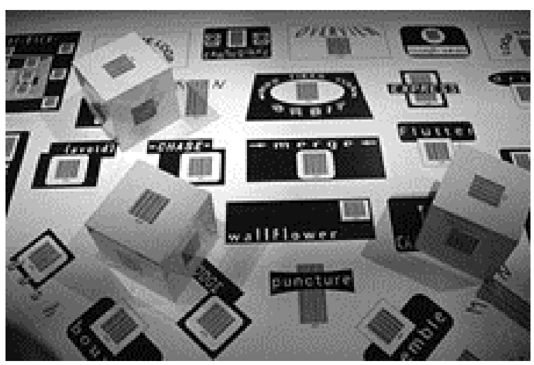

The cubes with barcodes.

The surface of the entire room has become a responsive membrane, has been turned into an immersive interface that can be used by a number of people simultaneously to control and react to a projected real-time computer-generated stereoscopic 3D world.

The bar code information within the installation

The virtual objects are based on static familiar things from everyday-life: a pair of glasses, a hat, a paper-clip, a suitcase, boots, in other words only objects that do not refer to living creatures. The moment that they become characters (and in which they are always also surrogates for the viewer) is only provisional and depends totally on their movement and interaction. Scanning in new information does not only modify the movement of an object but also its behaviour and its appearance (as well as the looped sound effects which are connected to the changes): the objects can expand and contract, can breathe, tremble, jitter or bounce. Certain commands cause the object to move slowly through the space while randomly changing the direction, or that it moves as a ‘wallflower’ into the nearest corner. Other bar code commands describe the relations between two objects: chasing, avoiding, merging or punching. However, once all possibilities of behaviour and movement have been tried and tested it does not mean a thorough understanding of the pattern of behaviour, since each behaviour is also dependent on all other objects’ current behaviour and location.

Because the system can distinguish each wand as a separate input device, each guest can have his own consistent identity and personality in the computer-generated world. Yet, the viewer is only partially able to control his virtual representation. The longer a person is interacting with the agent the more obvious it becomes to him that not only the behaviour of the objects is defined in relation to all other objects but that it starts to develop its own capabilities and characteristics depending on factors like age, size and history. Younger objects respond more quickly and more accurate to bar code scans than older ones. Since the contact between object and viewer is only established by scanning in information, the objects are on their own between these moments of human contact and the duration of these moments is determined by the viewer. By this the object escapes his human collaborator while at the same time his individual behaviour within the virtual space becomes stronger.

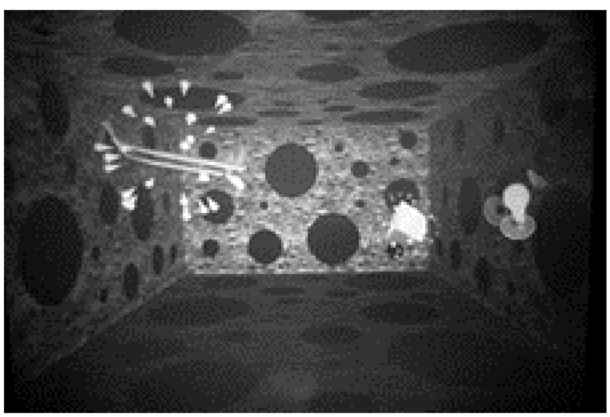

Virtual objects in the virtual space (screenshot).

Besides controlling objects, certain bar codes affect and modify the environment in which the objects exist. The viewer has the possibility to shift the point of view of the computer projection, and also the settings can be switched between various rooms and landscapes. By scanning in an earthquake he can affect all characters, the real and the virtual ones: the result is a temporary state of disorientation of all object.

Consequently, the interactive environment could be described as an audio-visual experience space within which real and virtual ambient sound are merging and real and virtual actions and reactions, alliances and counter-alliances are mutually dependent. Who controls the suitcase? Who is responsible for the earthquake? What is meant by ‘wallflower’? And also the individual way of dealing with and within both spaces continuously changes. The viewer controls, reacts, observes, is curious, experiences, protects his object or permanently manipulates it; he looks for a new wand to try out all those bar codes which were too far away to be reached.

The bar code actions produce a whirlwind of images and sounds, the entire environment becomes a virtual mixer. The resulting audio-visual cocktail apparently exists only for the sake of entertainment, amusement and for stimulating collaborative creativity. Yet it seems that in this work the ‘medium (the bar code system itself) is the message’ since the artists turns to the universal bar code which world-wide has been adopted to classify and serve commerce. Hardly any product exists without it and even though all information is visibly displayed it cannot be deciphered by the human eye. We are so familiar with the bar code that we hardly notice them anymore, and we almost never have access to the deciphering machinery. In Bar Code Hotel Hoberman has reversed the situation. So why did the artist choose this technology? For Hoberman the stripes are

> ”just an unpleasant act of contemporary life. They are ugly, plastered onto countless consumer products, defacing the design of packages, books, magazines. And they don’t seem to have any of the magical properties that often get attributed to advanced technologies. They are, however, one of the earliest infiltrations of the digital infrastructure into the universe of existing objects. They represent a kind of alternate reality superimposed onto the physical world. This reality is not addressed to us, but instead directly to the computer. Potentially everything that has a bar code can be apprehended by the computer. Things can be sorted, classified, priced (...) Eventually I realised that they [the bar codes] were an accessible and cheap technology that could be abused and misused in the service of art.”

By covering an entire room with bar codes and classifying them with terms which one normally would not expect behind the black and white code, he on one hand makes the invisible (e.g. the domination of the bar code within contemporary society) visible. And also the viewer immediately knows how to operate the interface. On the other hand, Hoberman uses this technology to emphasise the limitations of interactivity. When something is said to be interactive, it sounds as if the viewer would be able to change it somehow. But generally the choices are very restricted, and leaving is like checking out of a hotel room: the work returns to its pristine condition, and there is no evidence one was ever there. The only ‘location’ where the interaction of this work leaves its traces behind is in the remembrance of the viewer and the spectator.

The Hotel Seen as a Stage.

The First Visit: the Audience.

The stage area, which is entered, dominates with its black and white stripes the entire situation. If the space is filled with viewers, the spectator easily understands the concept of the action. As an observer he finds himself soon in the role of a film-cutter who looks back and forth from the real to the virtual space, in order to keep up with the actions and the resulting reactions. In the spectacle which is offered to him he perceives viewer and virtual surrogates as equally present, since he is able to link all acting components (humans or virtual objects) with each other: each change allows a reaction of all participants either by scanning in new information codes, by manipulating the parameters of the virtual projection, or by the specific characteristics of the objects themselves. For each spectator therefore a temporary never reproducible play is enacted. Only the stage design, the props and the virtual personnel are determined. And also the spectators themselves, though unconsciously, can become part of the play by taking on the part of the ancient chorus, when they start commenting on the scene or filling in newly arrived spectators.

Active and passive viewers in the installation space.

The spectator intervenes.

Once the spectator himself becomes an active viewer, also the relation between presence and representation changes for him. He becomes part of the play followed by the spectator, becomes present in the (inter-)action between real and virtual space. Using the bar code wand, he obtains a virtual agent within the virtual space of action, a representation that enacts his actions within the unreachable computer-generated world. But here the installation differs from the three others. Neither can the virtual space be completely dominated by one single person - the work is composed for collaboration - nor can the virtual space become entirely one's own space without establishing an alliance with a virtual partner.

Acting in the real space.

A slight distortion of this 1:1 relation arises, because of the objects' individual characteristics, something that has not become clear to the viewer while he was still a spectator. The virtual representation transforms into a semi-agent. Once the viewer has understood his new role as a temporary film director (since each decision he makes manipulates the entire situation) he also understands the various semi-agents rather as dramatis personae whose behaviour can be determined though not their 'personality'. If the active viewer rather wants to stay in the background he is able to change certain parameters to influence the action, without being virtually present: by changing the perspective, the light, triggering an earthquake. Each actor has the choice to design a choreography for the virtual space with the help of the props of the real environment. This always refers him back to the real environment. At the same time this condition stresses the inseparability of real and virtual space; it establishes a closed circuit from both spaces. Walter Benjamin (1977) described the position of the camera, which here corresponds with the interacting viewer's as

> "the only viewpoint where one does not see the 'special' (and utterly artificial) 'procedure'. The audience is only able to look where the camera has already looked for them."

Viewer and spectator perceive different pictorial levels in different ways. For the spectators virtual and real space are spaces within the same unique narration, which is presented, to them before their own eyes. And for him the virtual objects and the viewer form the cast of this action.

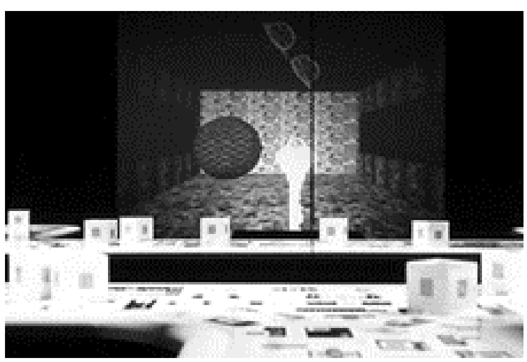

Real and virtual space of the BAR CODE HOTEL.

Space, Time and Action: the Space.

This artwork presents itself in various spatial and pictorial levels, which are mutually dependent and mesh inseparably. If one wants to define them one has to differentiate first of all between real and virtual space, and then one has to divide the real space into viewer and spectator level, and the virtual into a level of narration and one of action. In other words the artist does not deploy a black box which directs the attention immediately towards the projection screen and to the virtual space. Darkened is the space only because of the light sensitive projection technology. Action space is the *Bar Code Hotel* whose real space the actors and spectators have to share while the virtual picture level is exclusively reserved for the virtual actors. The connection between both spaces established by the actors' actions and reactions allows the two separated action spaces to merge. If no action is performed, the spaces are standing unrelated next to each other. In the real space the repertory is kept ready with which the virtual space is to be designed. Simultaneously this repertory defines the basic atmosphere of the real environment. The exclusion of everyday life is pointed out to the entering person. And by this it is also emphasised that here a reflection space in contrast to everyday reality is possible.

Time-Structure.

As much as for the spectator virtual object and viewer coincide in the personnel, time of action and real time correspond to each other. While for the viewer during his actions it becomes clear that the presence of his semi-agent within the virtual space has a limited 'life span', the spectator only realises a permanent appearance and disappearance of virtual objects as well as of real actors. The time of action equals for him the duration of his stay within the stage area. The actor on the other hand, if he wants to react to the actions of the other participating parties learns the rules of the virtual world, and co-ordinates his actions accordingly. This allows two different time structures to be present at the same time: the real time and the duration of the existence of the various objects (around five minutes). However, then the question arises which relevance the visual life span of an object has for the work as a whole. Since the virtual world does not possess a memory and does not remember, the time structure on the virtual level seems entirely insignificant. But in fact the separation of both time structures supports the differentiation between playing and watching, between actors and audience and secures the interactive work its utopian dimension.

What Happens.

The interplay of the components space, time, action and persons form the action, or rather the narrative element of the work. However, also here not the interpretation of an existing text is effective, but rather the artistic conception as the basis for performance. A performance could then be characterised by the combinatorial presentation of the visual arguments. For a characterisation of the *Bar Code Hotel* one cannot apply the classic conception of a tension line, since the work does not have a fixed time frame. The time frame, one could say, has rather converted into an ‘explorative duration’, determined by the duration of the stay. In *Bar Code Hotel* one finds two systems of variables: the projected actions and the actions within the environment, which enclose the projection. Both lines of action become the collaged experiences of the participants and the spectators, the perception replaces the narrative moment. The action results from the relations between the actors and their relation to the environment. Only in the remembrance the perceived turns to the narration. However, if one does not realise any development within the performed action, but only permanent changes of images and sound, why should one remember such a kaleidoscopic event? Yet the *Bar Code Hotel* allows completely different choreographies: one extreme would be the game without limits; and in contrast stands the graceful dance. If the viewer perceives both forms of expression, he recognises the *Bar Code Hotel* as a playground for those desires of life, which are expelled from everyday communication. Or he recognises it as realised dream - definitely it is seen as a platform that is differently used by each actor.

For the narrative moment this would consequently mean that something like the one narration does not exist at all anymore but only the separate experiences of the various interacting viewers, and all those then would add up to the entire (his)story of the work as a whole.

Shaw/Hegedüs/Lintermann/Stuck: conFIGURING the CAVE (1997)

Describing the Scene.

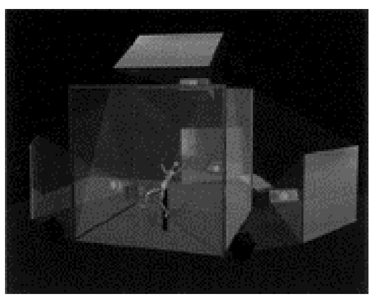

The installation space of CONFIGURING THE CAVE.

A cubed white space without a fourth wall, in the middle a spot-light illuminating a wooden man-sized model puppet, as used by artists, which holds its hands in front of its eyes. On a thin column the model puppet is fixed on an axis at the lower back. With so little words the environment can be described.

Equipped with a pair of polarising glasses (shutter glasses) one is able to move freely in this open space. The emptiness of the room and a consciously positioned spotlight focus the concentration on the puppet. Nearly reflexively the visitor takes away the hands from the puppets eyes; ambient though strange sounds can be perceived and from a distance a spiralling chain of cubes seem to approach, until they finally enter the CAVE space, an effect made possible by the shutter-glasses. After operating the interface the viewer realises that direction and distance of these geometrical forms correspond to specific positions of the puppet. The surfaces of some cubes carry images of Japanese spaces presented in a fisheye-perspective: a tea-room, the entrance of the ICC building, the home of the homeless at the Tokyo railway station Shinjuku. At regular intervals a massive form, a wooden rod, moves in and out the actual CAVE space. In these moments the puppet's position defines whether the viewer, puppet as well as spectators are temporarily located within the rod with its changing bright interior colours. The rotation of the puppet triggers a rotation of all virtual components. If the viewer puts the puppet's hands in font of its eyes, the geometrical world that has been projected onto the white walls, disappears. Once this initial position has been re-established, the play starts anew. As the handling of the interface has been trained in the first 'phase', viewer and spectator now know how to interact with the offered space. The puppet's eyes are uncovered again and text-walls one wall layered behind the other fill the space, which has been empty only seconds ago. Various white sign-systems - alphabets and hieroglyphs - are crossing the space, always depending on the puppet's orientation. On the floor of the room the viewer is able to see a Chinese rubbing stone carved with text. If the puppet is tilted by 180 degrees so the feet are pointing towards the ceiling, the signs seem to be collected on the floor to spiral around the puppet's head, additionally emphasised by buzzing ambient sounds. While the text-pages are moving vertically a calm female voice fills the space as if she would recite an unknown alphabet. Covering and uncovering the puppet's eyes opens the path into the following world: viewer and spectators find themselves among distorted world-maps and the five platonic forms: tetraeder, cube, octaeder, dodecaeder - corresponding to the four elements - and ikosaeder, which Euclid assigned to the entire universe. The position of the puppet's limbs helps to orient and explore, to 'move' through the space. The puppet as a globe. Maybe the visitors are able to recognise the various spherical images of the world which are merging and overlapping: the globe of the Surrealists, a satellite image of the world in which the continents have been exchanged by simulations in the form of clouds, and the astrological map of the heavenly bodies. The platonic forms, which float through these layers, can be virtually entered and deformed, while still being surrounded by the spheres. The spectators are located within the old-worldview, which placed man in the centre of the universe

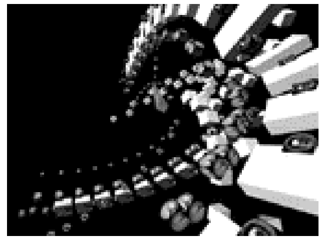

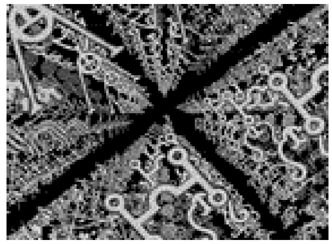

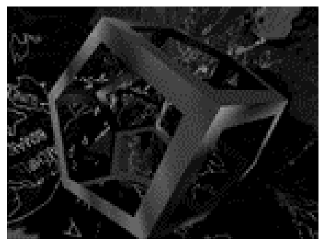

Screenshots form the ‘first’ three worlds: Material, Language, Macrocosm.

Virtually entering the following room the spectators see themselves among entirely concrete images: from a distance four spirals of line-drawn pairs of hands that perform various gestures and photographs - like windows into our world - are approaching. Moving the puppet allows to zoom the images closer and along the walls a panoramic photography of persons sitting in a restaurant is formed. Folding up the puppet’s hands trigger a hand shadow play on the CAVE’s rear wall.

The virtual CAVE worlds: Association, Union, Person, Emergence

Above the floor of the fifth world an image plane is seen of two Daguerreotyped naked women holding each other. This scene can be rotated and lifted via the puppet. Underneath a second image plane becomes visible, a satellite relief photograph of the Hiroshima geographic area. Lifting the first image plane far up by turning the puppet vertically reveals a black vacuum between the two planes, one becoming the heaven of the other. Lowering the plane again, the bodies are sinking into the topographical representation.

And the next world is in flames: the space itself, the spectators, the active viewer, the interface. The shadows of the puppet appear as simulations on the walls and the floor: real and virtual actions are merging. Through the conflagration within everyone seems to be captured, seven objects are permanently moving: a candelabra, a compass, a child's wagon, a gyroscope, a blue sponge, a camera and a cube which shows on each surface representation of the human being in relation to space. Folding the limbs into a specific position makes the invisible visible: some objects are converted into their wire-frame models. On leaving the world of shadows, for some moments the ghostly characters of the objects are floating through the space without any ambient sound.

In the seventh world the spectator is surrounded by spiralling organic shapes moving in-between blue brush strokes. Manipulating the puppet's limps one moves the restless elements in and out the CAVE space. The rather dark world can become a dense agglomeration giving the spectators the feeling of being captured by an ice-sea.

Covering and de-covering the puppet's eyes anew close the cycle: the viewer finds himself at the beginning of the cycle.

conFIGURING the CAVE: an Interactive Performance.

The aesthetic conventions forming the occidental term of theatre has its origin in the fundamental conditions of the occidental culture. According to its spiritual structures this culture is defined by the metaphysics of identity developed in Greek philosophy as well as by its corresponding Aristotelian logic. This kind of metaphysics and its logic assume a dichotomy of spirit and matter, of spirituality and materiality, of subject and object. The differentiation between essence and manifestation is a basic condition of this philosophy, whose approach is presented in the allegory of the cave, the metaphor of the theatre situation in the 7th book of Plato's Publica: human beings are tied in a cave allowing them only to see the shadows of the objects cast on the wall. They perceive these shadows, since their view to the outside is denied, as the true being of the objects. Starting from this basic pattern of defining the world, the occidental philosophy developed its systems, independent from the diverse definitions of the relations of essence and appearance. Hegel and Marx have marked the essential moments of this development in the younger history of philosophy.

It s quite obvious: this reality-model that up to today essentially influences the structure of consciousness and of experience of the occidental human being serves as an essential basis for the aesthetic conventions of those works which have been and will be especially developed for the CAVE-system: the performance within the CAVE as a metaphor of the world. To imagine, to copy, to create 'as-if' realities, all this is happening in this space. Additionally, this view is emphasised by Agnes Hegedues and Jeffrey Shaw when they say about their piece:

> ”This work embodies within its program a meta-language of functional relationships between bodily and spatial co-ordinates. These relationships are both physical and conceptual, creating an anthropomorphic world that reflects that very long history (which is common to most cultures) of conjecturing the body as the measure of all things. At the same time this work is located in a post-modern exigency which throws doubts on these grand symmetries, and releases the body into a vertiginous space of deconstructed co-ordinates and equivocal complexity. The balancing act is the fragile covariance of this representative body within a space (the world) of inchoate forms and reflections - where yet these are the images, which are the interactive extension and coherent confirmation of its (and our) identity. This soft boundary between place and person is aligned in this work with the dramatic immersion of the visitors' own bodies within the CAVE's virtual space of representation.”[1] (personal emphasis)

So what is really happening within the CAVE to space, time, action and persons and which results arise out of this interplay within the interactive environment?

The interface.

The empty environment with its wooden schematic body as an interface in the middle of the room, dramatically illuminated by a spot-light, reminds inevitably

---

[1] Quote from the initial concept paper of the artists.

of a classical stage-environment; one nearly expects the puppet to remove its hands from the eyes and start moving. As nothing is happening in the space, the puppet's immobility nearly forces the viewer to act, the interplay is starting. The resulting constellation is the physical relation of at least two unequal actors: human being and interface. The actors have to deal with their own physicality as well as with the partner's one, and this allows the interacting viewer to learn how to operate the interface while for the spectators the presentation of a choreography is unfolded. The established situation: the actors are not acting towards each other within an (internal) relationship of the performed play, but also (externally) towards the spectators. Once the play has started one has to add the given acting conditions outside the 'relation of the roles' between interacting person and interface, in other words, the performance space itself with all its visible sign-systems: light, music, the unfolded space. However, soon the viewer realises that exactly these conditions of action are actually not to be found outside the conditions of his role, but are on the contrary closely linked, are dependent. And also the spectator soon realises this interdependence: playing within this realm of experience is a pure up-to-date relevance. It exists only in its authentic execution as a current situation within the dialectics of playing and watching. However, these dialectics are conditioned by a permanent shift, since at any moment the spectator can become the actor, or he could be joined by two co-actors meaning that now on their actions have to be agreed on, in order to guarantee the continuation of the performance. The player can become as well a spectator again as a director, giving instructions. Also here it becomes obvious that the forming image-worlds deny any form of representation by their auratic immediacy; each performance is unique, becoming as a current action and experience a moment of life itself. This impression is additionally emphasised by the stereoscopic component: not only the various roles are permanently distributed anew, also the virtual objects which transform the space into a (virtual) world are perceived as being currently located within the CAVE-space. Experiences made in the CAVE-environment are part of life's obligingness and intensity, yet they are with the passing life irretrievably lost forever. The analogy to life results from the conditions that [a] acting viewer and reacting spectator are together performing the action (which is the perception of the various spatial representations) and that [b] because of both their physicality as living beings the projected spaces through which they move are permanently updated anew. But also the virtual objects, images, systems are much more than just virtual spaces. Their three-dimensional presence is effective enough to ask viewer and spectator to act and react: directly by getting out of their way, indirectly by the viewer's manipulation of the interface in a way that the objects are either transported into a distance or zoomed closer. Viewer and spectator as observer and explorer? As grasping travellers through unknown worlds?

The Cave.

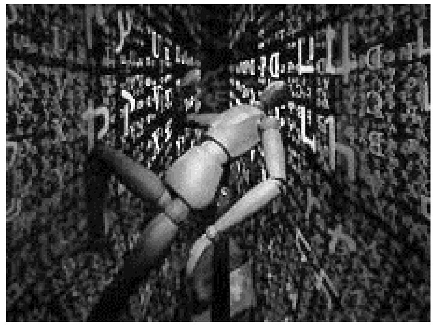

View inside the interactive installation.

In *conFIGURING the CAVE* the space is not existing. There are only spaces. Yet one could contradict, saying that in the initial situation - the moment in which the puppet's hands are folded up in front of its eyes and the white walls of the CAVE become visible and are not dispersed by the projections - one would actually be located within the space where everything is enacted, though, what meaning, what impact would this have? Is this space not in fact the eighth world? This is our daily-life world, which is confronted with the seven other worlds. In each world one has to get used to the newly presented spatial perception, one has to relate oneself to the (virtual) world. In the moments of the projection, the various spaces are present, for the viewer and spectator, independently of them being entirely virtual or not. And they are perceived as real as the real space with its three walls and the floor.

However, it cannot be denied that on another level the CAVE (as a real space) is present while it is impossible for the other seven worlds[1] of *ConFIGURING the CAVE* to be ever present as objects. These spaces as well as the experiences of these spaces are not recorded as a told story is in a book, or a picture or sculpture and accordingly does not confront an 'objective' view distant from its current status. Its character of being established in the art piece does only exist within the dialectical relevance of playing and watching. In relation to this, subjectivity as a viewpoint is unavoidable and cannot be escaped. As little as experience in general, experiencing the various worlds can be mediated they have to be made individually. However, the experience the visitor of the CAVE makes cannot be made in any real-life-context, since the work's topic is the body-space relation throughout history seen from a contemporary viewpoint. This establishes the anthropological basis *of ConFIGURING the CAVE* as a form of aesthetic action in an elementary law of the psychic life: the imagined process of the subject to become real in fantasy and in dream. The CAVE to live out fantasy, yet also to contemplate, to realise the dreams, one nearly gets the impression of an existing hope for a playfully

---

[1]The artists thematically structured the seven worlds, or rather the seven spatial representations of conFIGURING the CAVE which can be experienced during a visit in a cycle. Even though the worlds are numbered starting with Material, followed by Language, Macrocosm, Association, Union, Person and ending with Emergence.

lived draft of freed life. The CAVE, it seems, reveals a program of hope. Seen from the content aspect of the activity structure this moment appears to be the utopian share of VR, based on its nature as play.

On the Fragmentation of Time.

Similar as in all the other interactive installations, it seems that at first glance only one single time-structure exits, the real-time. The period of time in which experiences are made, corresponds to the actual duration of the stay within the interactive environment. However, the duration of the stay within the worlds can be structured individually and this in return influences the structure of the experience. Consequently this time structure is compelling for the sensorial perception of the recipient. Chronology and linearity are established by the artists insofar as individual access to each world is not possible: one follows the other in a cycle. If the viewer wants to visit a specific world, he has to cross - like in real life - the other spaces before reaching the desired one.

Though does the real time actually have any value in the individual worlds, does it play a role, is it perceived at all? How present is time actually? In daily life, time past is measured by the tasks done, covered distances, a ‘sense of time’ is developed from sheer habit. Within the seven virtual spaces one has to orient oneself permanently anew, the space can be manipulated, the relation between time and space seems to be annulled, has to be reformulated. Historical images, captured moments, documentary photographs refer with their presence to specific moments in time, though, as frozen moments they are objects, have been turned to references. As representations of past moments, however, they always also establish a relation to reality. Like this a time-structure is formulated as distance between the current moment and the moment when the photograph was taken, or between current moment and historical time of the belief in the conception of the spatial systems conveyed. Time becomes an additional variable component of the interactive work.

The Hyper-Action.

Already in *Bar Code Hotel* it could be seen that the action in the truly narrative sense does not exist anymore, though of course a play between the performing viewers and between viewers and spectators is unfolded, which is followed as if spell-bound and which is far removed from daily life. And viewers and spectators are removed to a travel through virtual spaces. Consequently one could argue that the action is constituted by the experiences made on this ‘travel’, and the viewer’s performance. These two elements compose together the narrative component of *ConFIGURING the CAVE*, and it is this action which finally is remembered. Yet, the activities the viewer is asked by interface and spectators to perform, always motivate self-portrayal and self-performance. These actions, however, are more willingly enacted because they are not documented but only remembered. Consequently the true action of the art piece has started at the moment the work was switched on and the continually changing spectators and actors of this play each have since been responsible for their current and specific flow of action. The artists have established the experience-space that is inactive without the participation of the viewers and spectators, and once it is activated, the artists seem to have assumed, the action is developed all by itself. As in real life.

Re-Configuration.

So what is the result of the interplay of these various components? The artists seemingly consciously abstain from an action, only the spaces are established to which the visitors have to relate themselves to, and which they perceive differently. And also the museum, in other words the environment in which the interactive environment is located, is a space beyond daily reality. It is a space for conservation, preservation, storing, and a visit of the museum is comparable to the travel through the virtual worlds leading through spaces within which the living movement of past as well as contemporary culture are stored, their inner restlessness, their obsession for new definitions. Since the viewer in *ConFIGURING the CAVE* sees himself in a permanently changing relation to the (virtual) space, the intense relation between utopia and reality seems annulled for the duration of the stay. Distance is created by the various experience-spaces to the real world in order to question and analyse it.

Presence and Representation as Couriers in Interactive Installations?

Throughout the discussion reference has deliberately been made to the tradition of the theatre, that art form which, until recently, was the only one capable of utilising all the dimensions, especially space and time. The installations have been examined in relation to the closeness of the conception of the theatre even though none of the pieces were explicitly inspired by this analogy. Nonetheless, what is decisive is that role of the viewer as an active component of the artwork. The artists withdraw as authors defining only the space for action as a stage with its repertory and technologies for creating the temporal illusion.

Despite the fact that all four artworks are interactive, the similar treatment of space is striking. All works define borders, describe an interior space, within which new and different rules are valid, some kind of stage on which a special event takes place, removed from reality. The event itself is thought as a unity and the individual work describes the form that guarantees this unity (see Ong, 1987). The border which divides the artwork from its environment and which have to stage and to maintain the work itself, stand for the much more general borders, which divide symbolic processes from real ones and which releases the filigree trial actions against the unequally robust real ways of life.

In interactive artworks - as could be seen - the various expressions of presence and representation often overlap, mesh, correspond to each other: They have become fluent. In *Bar Code Hotel* for example, the actor, present in the environment, has a representation - his agent - in the virtual space. However, like in the other three installations he is able to change his position from the pure voyeur, to the actor, to an agent in the sense of a directing person or in the sense of an 'intelligent agent' of the spectators when they are communicating.

Also the space of the individual installations itself oscillates between presence and representation. On one side the objective world is reflected into a virtual world of imagination with its limitations and the restricted virtual vocabulary (the imaginative world of the artist). It is permanently transformed anew into a representation. On the other side as a unity they both form together a training space for communicative acting. Beyond that, the environments create a platform for critically reflecting the controlling technology. This aspect becomes especially obvious in Hoberman's piece: the potentially surveilled (for example, the contemporary consumer) is handed over the tool to control, yet the control over the object controlled is soon lost.

Because the artist uses the universal bar code one can talk about a principle of picture-explosion. The artist sharpens our perception for the omnipresent bar code, nearly asks for scanning in additional bar codes unknown to the system. Yet these actions do not crash the system, but result in an another action within the virtual space.[1]

Such disobedience is definitely not consistent with the general understanding of interface designing. Designers are rather working under the assumption that the users are willing and able to follow specific instructions. The designers try to design user-friendly, sensible, logical, rather obvious. But in this artwork the artist intends to create a convincing artificial world, and for him it is exactly at the moment where the user STOPS following instructions that the effectiveness of the illusion is judged.

Just as Hoberman alienates the familiar, also Shaw alienates the real architectures of the city, just to allow a level of reflection to the available. By adapting the real parameters Shaw makes the transition of the actor from the real to the virtual easier. If he knows 'his way around', he is able to navigate steadily in-between the letters. Conversations within this black box expand the virtual space by the ambient sound of the real. The real is expanded by the virtual since the familiar, invisible to us in daily routine, is made visible in these works.

The World Generator on the other hand offers its actors a completely empty projection screen, onto which all life designs and all visionary worlds can be projected. And as far one moves away from everyday life in ConFIGURING the CAVE, when spectators and actors have to define themselves permanently anew in relation to the space.

Even if all these four works make different demands on their audience - reflection, collaboration, organisation, immersion - still all the artists attach importance on the alternation and communication between spectator and actor:

[1] Of course the possibilities are quite restricted in this case. Unfamiliar bar codes usually affects usually the same result. However what is important is the fact that the artist allows these possibilities and by this not leave the real world outside the interactive environment but rather incorporates it as a part of the work itself

entering the virtual space is generally simply designed, using interfaces that are either easy to learn how to operate them or which are even familiar. Once the spectator has become an active part of the works it is aimed for a situation in which he does not feel lost or even useless. Shaw has chosen familiar spaces which can be explored and experienced in a new way, Hoberman hands over a partner to his actors with whom they have to collaborate, and who manages the virtual space, while in the CAVE the viewer never moves alone in the unknown spheres and has a 'collaborator' with whom he is able to deal with the virtual. Seaman on the other hand creates a 'gap' - the wheels - which helps to create the own familiar environment.

Once the viewer has accepted his participatory role and the virtual space has been developed as experience space, the responsibility the artists took over when they designed the pieces, can be measured: in the amount of sense of the actions, which are carried out by the participants. Because if the formerly passive viewer is asked to become an active component within the work, who else should be responsibility for him but the person or persons who designed the work and have invited the spectators to participate. In interactive installations we are moving already far beyond a situation as in the classic computer-games, when the participants were just pushed into an established circuit, into a network within which already everything is predestined and set into its place. If the viewer is asked to make decisions, these decisions must be meaningful. The common aim must be that the participants are able to bring in the full range of their experience, knowledge and know-how on entering the works. And even if the individual pieces do not fulfil this goal, might not even 'have it in mind', still these installations refer to the necessity. Each work seen by itself only presents one concept with which the artists peruse their completely individual goal. Yet, if the task is to create a space for exactly these kind of spaces, the various concepts and components have to be considered. And one has to consider also the apparatus, which makes the multimedia-characterised performance possible at all, the actor and the spectator while the distribution of the roles must be kept variable, as well as the dimension of the freedom for performance and action, while these two elements must make sense - beside their selective insights which can only be offered to the spectators.

References

Mitchell, W. (1995): *City of Bits - Space, Place and the Infobahn*. (Cambridge, Mass)

Taylor, M & Saarinen, E: (1994), *Imagologies - Media Philosophy* (London)

Glaser, P. (1995) *24 Stunden im 21.Jahrhundert*.. (Frankfurt a.M)

Roetzer, F.(1995) *Telepolis. Urbanit im digitalen Zeitalter*. (Munic)

Benjamin, W. (1977) *Das Kunstwerk im Zeitalter seiner technischen Reproduzierbarkeit*. (Frankfurt a.M.)

Virilo, P. (1978) *Fahren fahren, fahren*. (Berlin)

Ong, W. (1987) *Oralitaet und Linearitaet. Die Technologie des Wortes*. (Opladen 1987)

Bolz, N. (1993) *Am Ende der Gutenber-Galaxis*. (Munic)

Wooley, B. (1994) *Virtual Worlds* (London)
Heidegger, M. (1975) *Die Grundproblematik der Phaenomenologie*, no. 24 (Frankfurt a.M)

Chapter Three:
The Legible City

Jeffrey Shaw

ZKM

The Legible City is an interactive computer based art installation first shown in 1989. It should be seen in the context of an artistic practice beginning in 1967 where research into interactivity and virtual imaging techniques was a central concern. It is probably one of the earliest art works to exploit interactive 3D simulation technologies. In the context of its further development within eSCAPE as a distributed multi-user virtual environment, it is useful to look at the genesis of the original work, particularly with respect to the way in which technological constraints were dealt with in the design process.

We also include this discussion in order to relate the issues covered in this Deliverable to those in Deliverable 3.1. In the latter work package we have been considering ways in which techniques for presenting e-scapes might be developed to support a range of presentational capabilities and constraints amongst users. Here we consider the ways in which techniques of artistic presentation develop within a range of technological constraints. We do not do this in order to suggest that e-scape design might be explicitly informed by such techniques, but rather we report this as an integral element of our explorations in this area, sensitising the development of the e-scape framework within 3.1 where appropriate.

Interactivity.

Body and space

The history of art exemplifies a complex set of negotiations between body and space - negotiations between the actual domain of the real body of the viewer and the real space he inhabits and the virtual domain of the represented body and represented spaces.

In the 60's and 70's the evolution of mixed-media and inter-media researches made it clear that the exigencies of the body were a focal issue. The artist's performative body and the spectators participatory body together constituted the notion of a body that was increasingly being 'immersed' in the dialectics of mediated experience. Today the protagonist body plugged into virtual reality is a perspicuous extension of this conjunction.

Event structures

Early forms of interactivity focussed on the direct physical participation of the viewer in mixed media performative situations. In Shaw’s work these 'event structures' were large lightweight inflatable sculptures incorporating film and slide projection, light and sound effects. The projection surface became an 'expanded cinema' where the viewers could jump into the screen and immerse themselves in the pictorial space.

Video disc technology

The advent of the video disc opened a new dimension of interactivity in the arts. For the first time a medium allowed real time random access to pre-recorded audio-visual data. Most importantly this allowed the creation of non-linear narrative structures *(Lorna, Inventer La Terre)* and also the construction of interactive navigable spaces *(Aspen Movie Map, Inventer La Terre).* For many artists the experiences of working with this medium primed them for the new range of interactivities that soon came to be offered by the computerised digital audio-visual systems.

In *Inventer la Terre (1986)* a chrome-plated column stands on a round black terrazzo base. This column has a viewing aperture, two controlling handles, and a pair of loudspeakers. Looking through the opening in the column, the spectator sees a large virtual image projected into the museum space - this image is overlaid on his view of the real environment. By pressing the handles, the spectator can interactively control the rotation of the column and the movement of this virtual image. Sound and image are interactively accessed from two videodisc players and the virtual image is generated by an optical display system within the column.

Initially there is a 360-degree panorama representing six different symbolic sites in a landscape. As the spectator rotates the column, he moves his point of view in a panorama apparently floating all around him. The viewer can select one of the six sites by pressing a button on the column handle. Linked to each site is an approximately three-minute video sequence which articulates the specific themes embodied by that place. At the end of each video sequence, the viewer returned to the panorama and then, by turning the column, choose another place to go to.

The title of this work *Inventer La Terre* refers to its central theme - the exemplification of archaic, mythological and esoteric world sciences. The six symbolic panoramic sites and their respective video sequences are each concerned with specific aspects of this overall subject - the totemic world, primal creation myths, astrology and astronomy, sacred geometry, hermetic geography, and the creation of fictional worlds in art and literature. 3D image synthesis, digital image processing and digital post production techniques were used extensively in these visual sequences.

A theory of interactivity.

The interactive art work is a virtual space of images, sound, text, etc. which reveals itself (wholly or partly) when acted upon by a user. That interactive space can exist within the frame of a particular medium of presentation (e.g. video monitor), or it can be projected (coherently or paradoxically) out into the actual space of the user. It can also become a wholly synthetic space surrounding the user.

The possible mechanisms of interactivity are varied. On the one hand are all the machineries: sensing devices of temperature, sound, pressure, movement, form, etc. On the other hand all the communication sensibilities of the user: touch, gesture, movement, talking, looking, hearing, etc. The technological tools of interactive visual expression are computer aided systems of digital image processing, synthetic image generation and random access information storage, because these tools can effectively describe the real time coordinates of a virtual space. These tools as language can also make use of the vocabularies of all the earlier art forms (painting, sculpture, architecture, photography, cinema) as well as the new vocabulary of wholly synthesised (programmed) images.

The interactive art work is a latent space of images, sound, text etc., whose aesthetics are embodied both in the coordinates of its concealed potential form, and in the scenarios of its interactively manifest form. The scenarios of such an art work are at the same time the manifest paths of request and desire of the user - the art work is each time re-structured and re-created by the action of each user. It is also possible in the future to consider a complex mechanism that actually modifies its own program in response to its users.

The interactive art work is mechanism, and program, and all the potential forms of its possible scenarios of deconcealment. Each user is raconteur and auto-biographer of one of these possible scenarios. For the non-active spectator who watches the interactive art work being manipulated by a user, there is the unique experience of seeing it elucidated through another’s eyes - so its manifestation is performance.

Virtual imageing techniques.

Film and slide projection

The history of the cinema is also a history of technological efforts to increase the sense of immersion and participation of the viewer - e.g. Vistavision, Cinemascope, Stereovision Surroundvision, Imax, Omnimax. This was supported by specific filming and editing techniques such as the identification of the viewer with the moving camera *(killer film The Man with the X-Ray Eyes).*

Multi-screen slide projection also took this direction. The *Diadrama (Rotterdam 1974)* was an audio-visual performance with slides projected onto a large curved wide screen by three pairs of synchronised projectors. The main themes explored in this work were panoramic methods of representation, the collage of visual and temporal events, and the articulation of various simulated and expanded cinema situations. Most of the photography for this work was made by rotating a camera in three positions to capture a 270-degree field of view in order to create an immersive panoramic effect for the viewers.

Expanded cinema.

Expanded cinema sought to extend the space of the projected image into the space occupied with the viewer. Real space became deconstructed by a great density of mediated forms and images which constituted new, indeterminate, pseudo-virtual environments (the modern ’disco’ is also heir to this strategy).

Movie Movie (Belgium 1967) was an expanded cinema performance specially created for the Experimental Film Festival in Knokke-le-Zoute (Belgium, 1967). It took place in the foyer of the festival building, with the audience sitting on the stairs and balcony. Three performers first brought in the inflatable structure and unrolled it on the floor. Then it was gradually inflated while film, slides and liquid-light show effects were projected onto its surface. The architectural form of this inflatable structure was conical with an outer transparent membrane and an inner white surface. The projected imagery first impinged lightly on the tightly inflated outer envelope and then appeared on the semi-inflated inner surface. In the intermediate space between the transparent and white membranes various material actions were performed to materialise the projected images. This included the inflation of white balloons and tubes, and the injection of smoke.

The intention of this work was to transform the conventional flat cinema projection screen into a three dimensional kinetic and architectural space of visualisation. The multiple projection surfaces allowed the images to materialise in many layers, and the bodies of the performers and then of the audience (many of whom spontaneously threw off their clothes) became part of the cinematic spectacle. In this way the immersive space of cinematic fiction included the literal and interactive immersion of the viewers who modulating the changing shapes of the pneumatic architecture which in turn modulated the shifting deformations of the projected imagery. A sensual conjugation of actuality and fiction was achieved through a mediated dematerialisation of their respective boundaries.

Keying techniques

The cinema (as well as photography and then television) developed various techniques for real time 'matting' of live action into pre-recorded backgrounds. In effect these techniques achieve an effective and immersive conjunction of the 'actor' and 'artificial environment' that prefigures many of the ambitions of virtual reality.

The earliest of these techniques was simply back projection. But this was followed by the development of 3M's special 'retro-reflective' projection screen which used specific optical properties to create a very convincing collage of the real and projected situations.

Viewpoint (9e Biennale de Paris, France, 1975) adapted this technology to constitute what could be considered one of the first 'augmented reality' VR installations. This work created a collage of fictional events within a museum space (the Museé d'Art Moderne de la Ville de Paris) by making projected images of the events appear contiguous with the real space and actual situations. The work was constituted by two structural elements: a large projection screen and an optical viewing console with an automated pair of slide projectors. Due to the screen's retro-reflective properties, the projected image could only be seen through the viewing console - from any other position in the room the screen was a plain grey surface. The inherent brightness of this projection system enabled it to be concurrent with the ambient lighting in the museum, meaning a coherent collage of the projected images with the surroundings was achieved. Furthermore, these images showed exactly that part of the museum room hidden by the screen, creating a seamless continuity between the virtual and actual spaces. Visitors walking by the screen were unaware they had entered and become part of the visual space of the projected image.

Twelve different events had been staged and photographed in the same museum area prior to the exhibition. These projected events were programmed in slides sequences that were triggered whenever a visitor stepped onto the viewing console - one of the twelve events was then shown. For example, one such projected event showed a visitor who made a bed for himself on a museum bench and then went to sleep on it. Another showed the projection screen itself being built up and gradually blocking the view of the museum space behind it. In other sequences a visitor was seen walking up to the museum window and smashing it with a pick-axe.

Optically generated virtual images

The mirrored image is an age old method of projecting a virtual image into real space. Magicians have made good use of it to create all forms of illusions, and the Baroque stage is famous for such devices on a grand scale.

Early experiments with computer-generated stereo imagery led to the development of optical devices that projected these images virtually. A virtual-reality type of apparatus *(Amsterdam,1979)* using a see-through mirror enabled the viewer to rotate and tilt a monitor so that various simple computer-generated objects could be seen floating in different places in the real space. This kind of virtual projection technique was also used during a performance *(Rotterdam,1978)* where the audience looking through a mirrored window saw a computer-animated swing superimposed over a performer seated on a real swing.

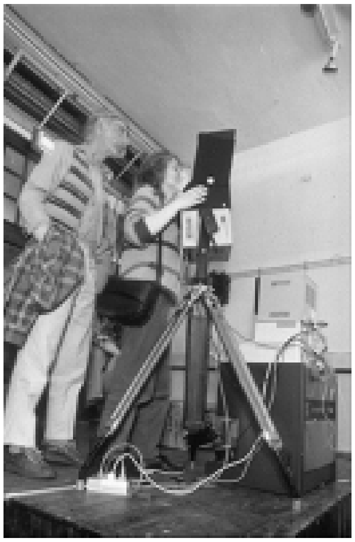
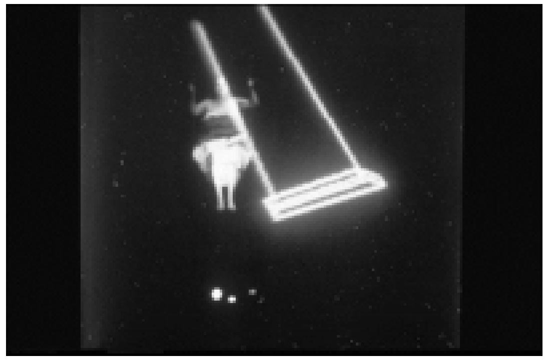

Points of View (1983) - precursor of The Legible City.

Description of the work.

Points of View was a 'theatre of signs' with both stage and protagonists being provided by a three-dimensional computer graphics simulation that was video projected onto a large screen in front of a seated audience. The action of the work was controlled by one member of the audience using two specially made joysticks. Using embodying techniques developed for flight simulation, this work gave the operator the ability to interactively move his virtual point of view 360 degrees around the stage, 90 degrees up and down from ground level to aerial view, and forwards and backwards from a far distance to the centre of the stage.

The representation of the actors on the stage was derived from the ancient Egyptian alphabet - each figure was a hieroglyphic character. This constellation of signs was used to articulate a world model with an underlying set of physical and conceptual relationships. Sixteen sound tracks - mostly spoken texts - were interactively linked to the image via the same joystick that controlled the visual movements. Functioning like an audio mixer, this modulated the relative volume of the various voices with respect to the changing spatial positions taken by the viewer. Here again an intrinsic conceptual structure was articulated in the spatial mapping of these sound tracks to four positions on four levels of a semi-sphere.

Two further installations were made using the same functional and iconographic structures as this work, but with differing contents. *Points of View II - Babel* (1983) addressed issues relating to the Falklands War. *Points of View III - A Three Dimensional Story* (1984) explored the notion of an open artwork by inviting sixteen people to make narrative contributions which were then interactively linked to the visual scenography so that the viewer could navigate between the parallel stories.

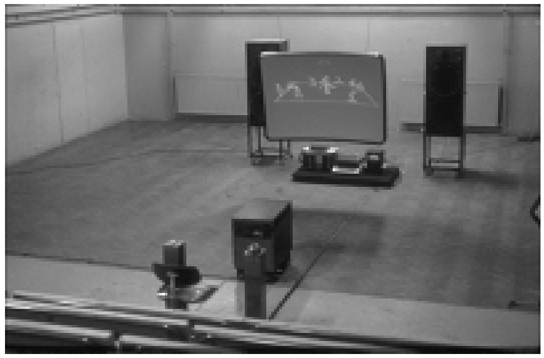

Technical implementation

Points of View was realised using an standard Apple II computer with a 2x accelerator board. The application software was developed using the SubLogic 'flight simulator' tool set (an very early example of the crossover from military simulation, to gaming and art applications). The composite video out from the Apple II was sent to a first generation monochrome video projector, projecting onto a curved silver screen. Custom joysticks were developed (connected to the Apple II paddle ports). The sound was generated by eight auto-reverse stereo cassette decks (sixteen separate mono channels) and interactively modulated (by the user) using voltage controlled amplifiers. Live radio was a seventeenth available channel.

Functional and aesthetic resolution of technical constraints.

Apple II graphics

The accelerated Apple II was capable of drawing about 100 monochrome straight lines in a 3D space at a rate of about 8 frames per second! Screen resolution was about 380 x 280 pixels. Needless to say these parameters were a severe constraint on any visualisation ambitions, but specific choices were made that allowed the integrity of the art work to succeed and even capitalise on these limitations.

Appropriate graphics design

With such restricted graphics possibilities any attempt to model real objects or personages would have been unacceptable. Therefore Egyptian hieroglyphs were chosen as a graphical language - because of their appropriately simple yet visually elegant forms and because of their strong value as 'signs' serving the articulation of the content of the work (visitors were given a leaflet on which the dictionary meaning of these signs was given).

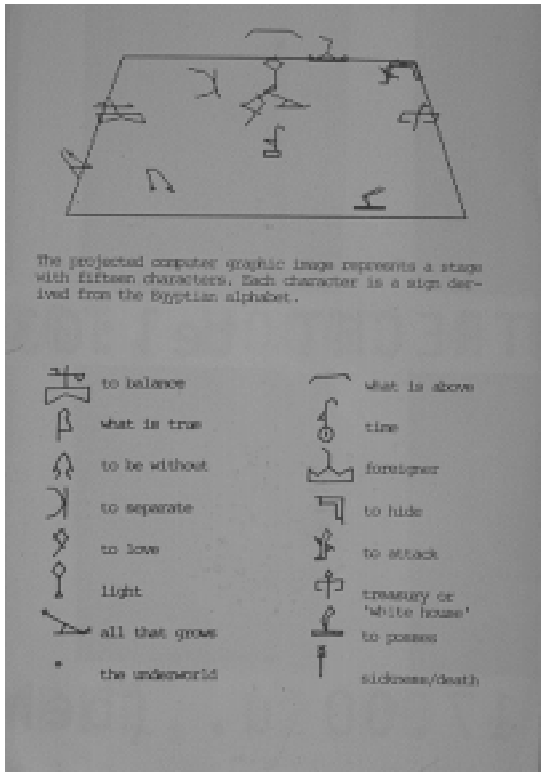

Using the 'jaggies'

The low resolution display of course resulted in obvious 'jaggies' (especially evident when projected). But instead of being objectionable the movement of these 'jaggies' seemed to actually animate the hieroglyphic signs as the viewer navigated around them - an aesthetic bonus!

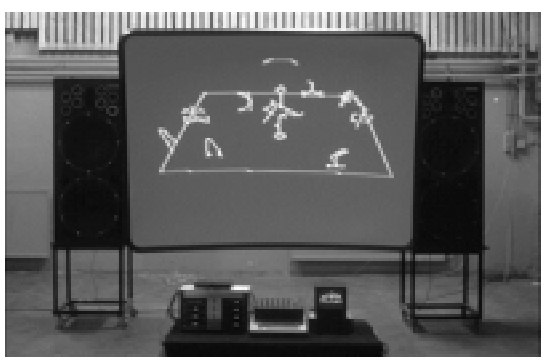

Monochrome or colour

Such a hieroglyphic universe obviously made no demands on colour - so a monochrome system was also seen as appropriate.

Specific advances made in this work

Screen space

The effectiveness of simply projecting the computer generated image onto a large curved screen lies in the fact that it immediately creates a window into a virtual space which a large group of people could share.

Operator and onlookers

Through a strong sense of identification with the operator the passive onlookers in The Legible City were able to share a real time interactive experience of the work. Thus not everyone had to be an operator to feel as if he/she had enjoyed the work.

Custom interface design

Custom designed interfaces that embody a direct visual expression of their functionality are invaluable. Of further importance are metaphorical characteristics that serve to articulate the contents of the work. A custom interface also distances the work from the undifferentiated world of computers and game machines.

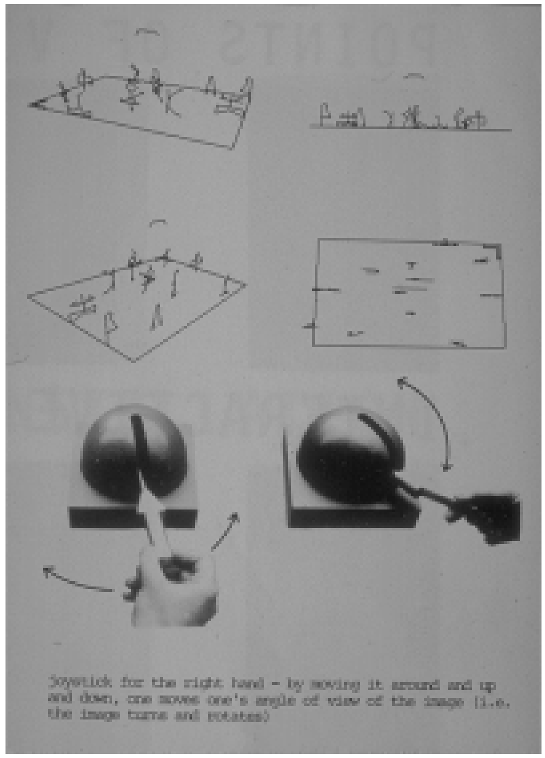

Navigation paradigm

This work is one of the first clear embodiments of the paradigm of the viewer as a navigator/explorer of a virtual space of computer generated forms. Such an interactive personal journey is both physical (travelling in the space of those forms) and narrative (experiencing an order of meanings expressed by those forms).

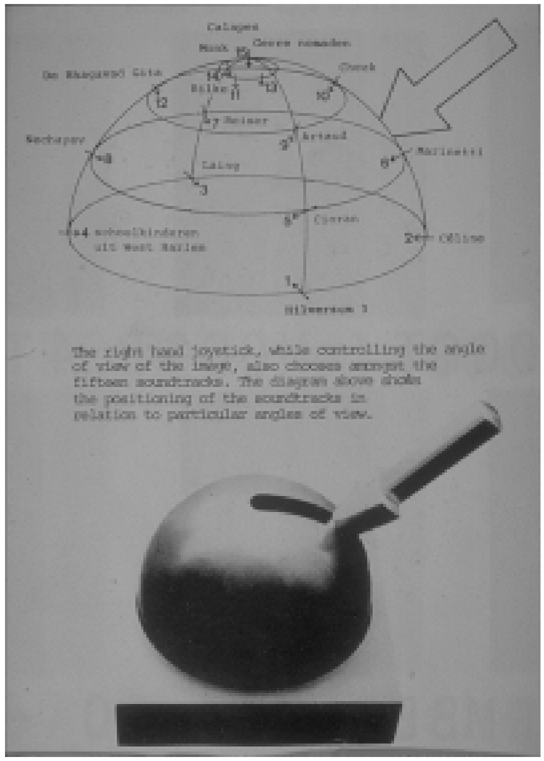

Theatre and virtuality

Presented in an actual theatre (Mickery, Amsterdam) and showing a virtual theatre space (stage) in which there are the protagonists (the hieroglyphic characters), this work first identifies the congruence (and collision) of our theatrical traditions (and understandings) with the new dramatic space of virtuality.(what's her name ? ☺).

1983-1989: Related Works

Points of View identified many of the characteristics and potentialities of a interactively navigable virtual world of visual (and auditive) forms. But the further development of such real time 3D graphics works was somewhat crippled by the limitations (and/or costs) of the then available hardware and software tools. For instance, a relatively simple project initiated by Shaw in 1986 with the French flight simulation company SOGITEC was budgeted to cost FF 3 million just to create the database! So in the interim, until a new generation of relatively affordable and efficient graphics systems arrived, Shaw pursued research in this field using different graphical solutions.

Following from the Aspen Movie Map experiences, the MIT Media Lab developed the Spatial Data Management system in which a navigable 3D space was built up out of an arrangement of flat image planes. They called this approach 2.5D, and the hardware/software system needed only to deal with real time 2D image processing.

The Narrative Landscape was created using this technical and conceptual approach, but in the process specific advances were made which also informed the development of *The Legible City*.

Description of The Narrative Landscape (Amsterdam 1985)

In this installation images are projected onto a large screen lying flat on the floor of the exhibition space. The spectators stand at least four metres above the floor, where a balcony or other structure allows them to walk all around the screen and look down at it from a height. On this balcony there is a three-directional joystick that enables any one of the viewers to interactively operate the work.

The images as well as spoken texts are digitally stored in a computer and the viewer uses the joystick to control panning in any lateral direction over the surface of these images and zooming in or out of a chosen part of an image. At the zoom extremes the joystick generates a digital transition to a new image - a process experienced by the viewer as a breakthrough from one image level to another.

The Narrative Landscape is constituted by twenty-eight images which are related to each other by a specific spatial as well as conceptual architecture. The primary image (a satellite picture of earth inscribed with a Hebraic astrological chart) is divided by a grid of red lines into nine areas which define access to nine groups of three images. By choosing any one of these nine areas to zoom through, the viewer arrives at a particular stack of three images. The three images are arranged one below the other and the viewer can move up or down through these three levels. He can also go back to the primary image and then choose another image group to zoom into.

Each group of three images has a distinct narrative formation. At the same time, all nine groups are structured as an iconographic triptych: the images on the first level represent a place; they have the scale of an aerial image of a city or landscape; the images on the second level indicate the body; they have the scale of human situations; the images on the third level show a figuration of signs that symbolically extend the themes expressed in the first two image levels of place and body. The underlying metaphor is one of emblematic places whose typologies are articulated in the fate of its denizens.

Most of the images in this work are derived from photographs that were digitised, collaged and electronically processed. Because these images are a digital raster, the action of zooming is also a process of increasing abstraction as the pixels that compose an image become progressively larger. This interactive journey through changing levels of representation induced by digital transformations is a basic formal and conceptual aspect of this work. The texts of *The Narrative Landscape* were written by Dirk Groeneveld and are conceived as nine distinct narrative poems interactively linked to the nine groups of three images.

Specific advances made in this work

Development of a new custom joystick

This joystick embodied X and Y movements for three directions of panning (X and Y and the diagonals) as well as a push-pull Z axis to control zooming in and out of the imagery. But it was noticed that without printed instruction, this unusual Z axis movement would not be discovered by the casual user.

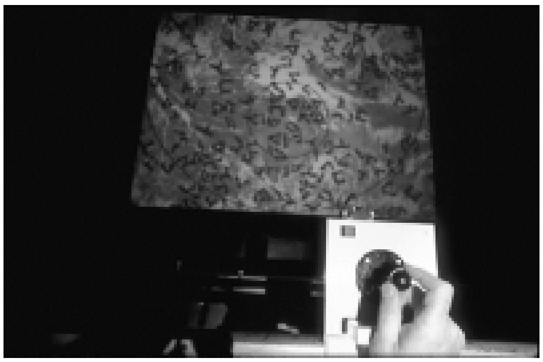

Immersive projection

With a 3m x 4m screen lying on the floor about 4m below the viewers (who could walk on a balcony all around it), an emotionally immersive situation was created that was highly differentiated from the usual experience of looking at images standing vertically.

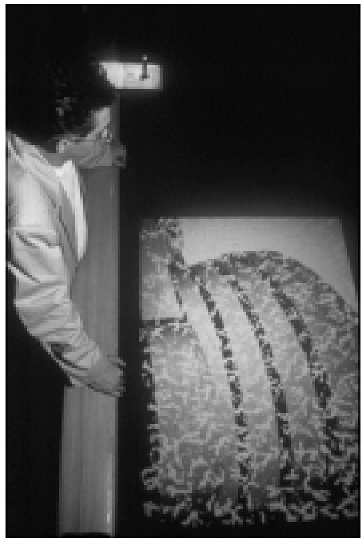

Technical constraints exploited.

The simple functionality of panning over and of zooming into the pixels of an image was here embodied as a method of interactive navigation through a narrative database of interrelated images (and spoken texts). Two metaphors were constituted - one a layered arrangement of 'pages' (and 'places') which could be travelled, explored and penetrated through from one level to another, - the other a magnifying window that could range between coherent overviews or abstracted close-ups.

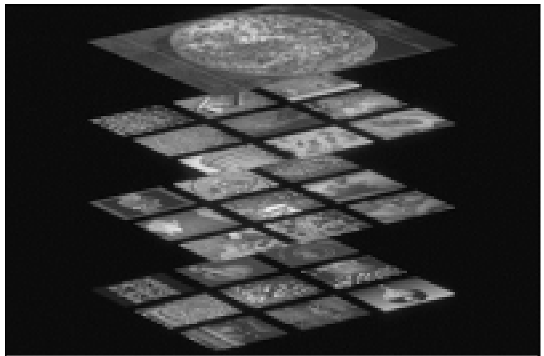

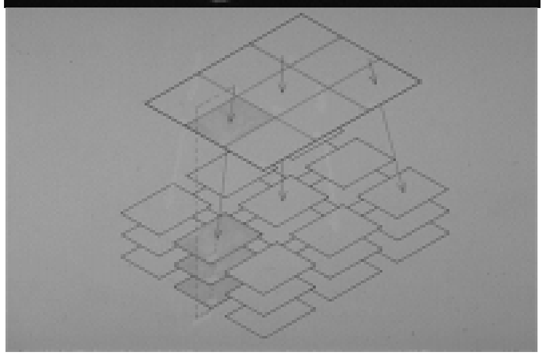

The contents of a ’virtual world’

Whereas Points of View identified the theatrical paradigm in VR, the Narrative Landscape identifies a basic geographical paradigm that is very appropriate to VR. In the first place it is simply a geographical space that is more or less identifiable with parameters of familiar space. On the other hand it attaches symbolic and metaphoric values to these spaces following in the tradition French Situationist notion of ’psycho-geography’. This means territories (or cities or buildings) that are in one way or another imbued with narrative (or other informational) value, allowing a mental construction to unfold in the process of the users’ navigations (c/f Italo Calvino’s ’Invisible Cities’).

The Legible City

Prototype (Bonnefantenmuseum Maastricht, Holland, 1988).

Description of the work

The Legible City was first presented as a wire-frame-graphics monitor based installation that was interactively operated by a joystick. The visitor was able to move through a simulated representation of a city that was constituted by computer-generated three-dimensional letters that form words and sentences along the sides of the streets. Using the ground plan of Manhattan the existing architecture of that city was completely replaced by a new architecture of letters and texts written and compiled by Dirk Groeneveld. The physical area was confined $ by 34th and 66th Streets, and Park and 11th Avenues. The texts were eight separate fictional story lines in the form of monologues, respectively by ex-Mayor Koch, Frank Lloyd Wright, Donald Trump, a tour guide, a confidence trickster, an ambassador and a taxi driver. Each story line had a specific location in the city - for instance the taxi driver following a meandering path to his destination.

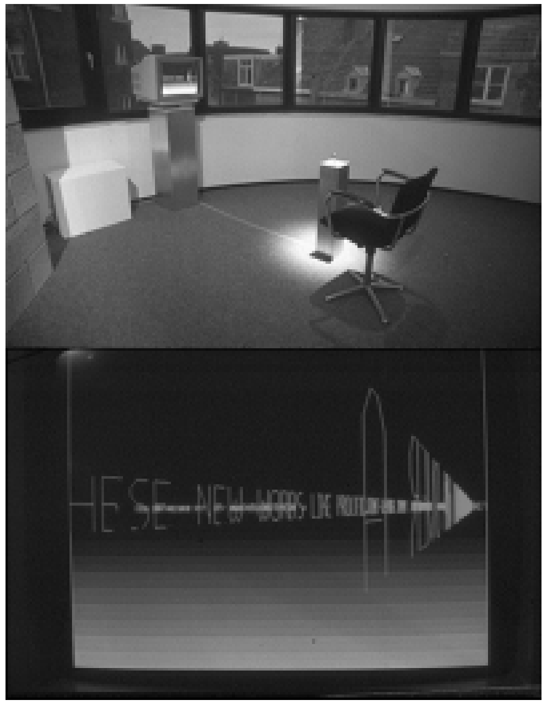

Content structure of The Legible City

In *The Legible City* the viewer rides through a virtual city made up by an architecture of letters and texts. The bicycle trip through these cities of words is consequently a journey of reading. Choosing the path one takes is a choice of certain texts and their spontaneous juxtapositions. The identity of these new cities thus becomes the conjunction of the meanings these words generate as one travels freely around in this virtual urban space. In this way one can view the city simultaneously as a tangible arrangement of forms and an immaterial pattern of experiences. Its architecture is a morphology of language, its ground plan a psycho-geographic network and its streets a labyrinth of narrative pathways.

The Legible City links this idiosyncratic architectural curriculum with earlier Dadaist, Surrealist and Lettrist subversions of the urban landscape. Andre Breton called for the Palace of Justice in Paris to be destroyed and replaced *"by a magnificent graffiti to be seen from an airplane"*. The Situationists emphasised the qualities of the *derive* - the unpremeditated drift through the city as a method of disorientating its compulsive forms and regaining the marvellous. Guy Debord exclaimed: *"That which changes our ways of seeing streets is more important than that which changes our way of seeing paintings"*

Textual structures of The Legible City

There are many interesting contextual references in the visual arts for this textual approach of the Legible City:

-- Saul Steinberg's drawing which humorously use words as objects in the landscape
-- Claes Oldenburg's proposal for the new police headquarters in Chicago was the world POLICE

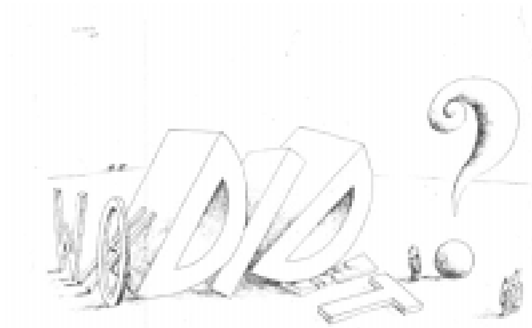

-- The traditions in art which rejected image making as 'idolatrous' and so exploited text as an alternative and visually rich expressive medium (e.g. in Islamic culture). In the 80's computer graphics was closely associated in the public mind with 'flying logos'. In effect computer graphics took letters off the surface of the pages and put them as objects into space (cf the Star Wars title sequences and Prince's clip 'Alphabet City').
-- There are also important literary references for The Legible City: - Writers and poets such a Julio Cortazar, Michel Butor and Raymond Queneau have written texts that can be read in whatever order the reader chooses. William Burrough's 'cut-up' writing techniques are also radical and pertinent.
-- A long tradition of 'visual poetry' hybridizes visual and textual expressiveness. As early as Simias of Rhodes (300 BC) who for instance wrote poems in the form of an egg and a bird. Labyrinthine arrangements of texts were also common - carrying the notion of the reading as a 'search' and

a 'journey' (e.g. Johan Kankel). The poems of Apollinaire and Tristan Tzara foreshadow the more recent 'Concrete Poetry' movement that fully concentrated itself on the conjunction of visual and literary forms.

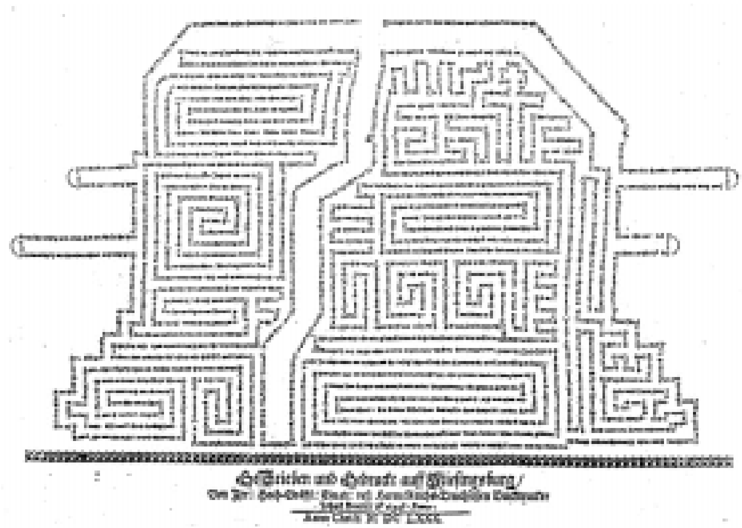

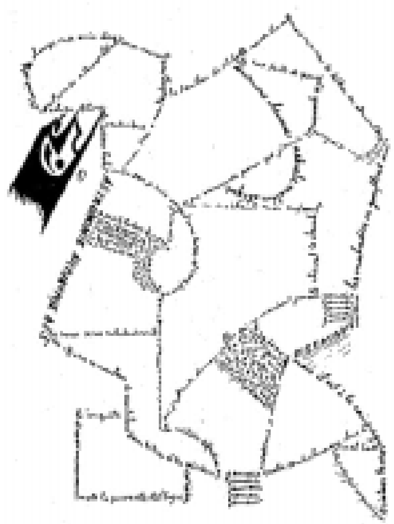

One might say that this long tradition of marrying text and image is a precursor of 'multimedia' - especially in the sense that the one does not merely 'illustrate the other' but that the convergence of the two aims to create new synthetic values.

Technical aspects of the prototype

The first version of "The Legible City", ran on a PC with an Intel 386 CPU and a hardware 3D acceleration graphics card. I don't understand this sentence: The software was written in C. and some 386 assembler for performance improvements. A simple algorithm kept the 3D drawing time within limits by clipping everything beyond a certain distance of the viewing position.

The next version was implemented on a Silicon Graphics IRIS 4D-20. This implementation had solid rendered letters, shading, better culling and colour. The culling included level of detail switching and depth fading. The core part of the software was taken over from the PC version except for the drawing module which was rewritten to use the Silicon Graphics graphics library. The joystick was replaced with a bicycle - a potentiometer attached to the steering wheel and a tachometer attached to the real wheel supplied a PC with the velocity and rotational information. The PC calculated the position of the bicycle within the city and sent this position data to the Silicon Graphics computer.

To aid the creation of the Legible City databases a program running on an MS-DOS computer was developed. This program handled the databases of the streets and allowed the user to specify each letter as well as its size and colour, and would calculate the remaining space available on the street. The resulting databases of letters and streets were read indirectly by the Silicon Graphics program.

Recent implementations of the Legible City use newer more powerful Silicon Graphics computers allowing features such as fog (instead of depth fading) and texture mapping.

Creating the database of The Legible City

Custom software tools were developed to simplify the making of the database. Mapping: in the first place a standard map of Manhattan was digitised on a digitising tablet. Each street was segmented between intersections with its crossing streets, and each segment was named and its actual length calculated in meters. Text editor: each street segment could be called up and then the texts simply typed. The program kept track of the actual letter widths and separations, and indicated when a segment was ’full’. When it was necessary to complete a word at the end of a segment, the whole segment group of letters could be simultaneously stretched or squeezed. Another software tool took all this information and automatically generated the 3D database.

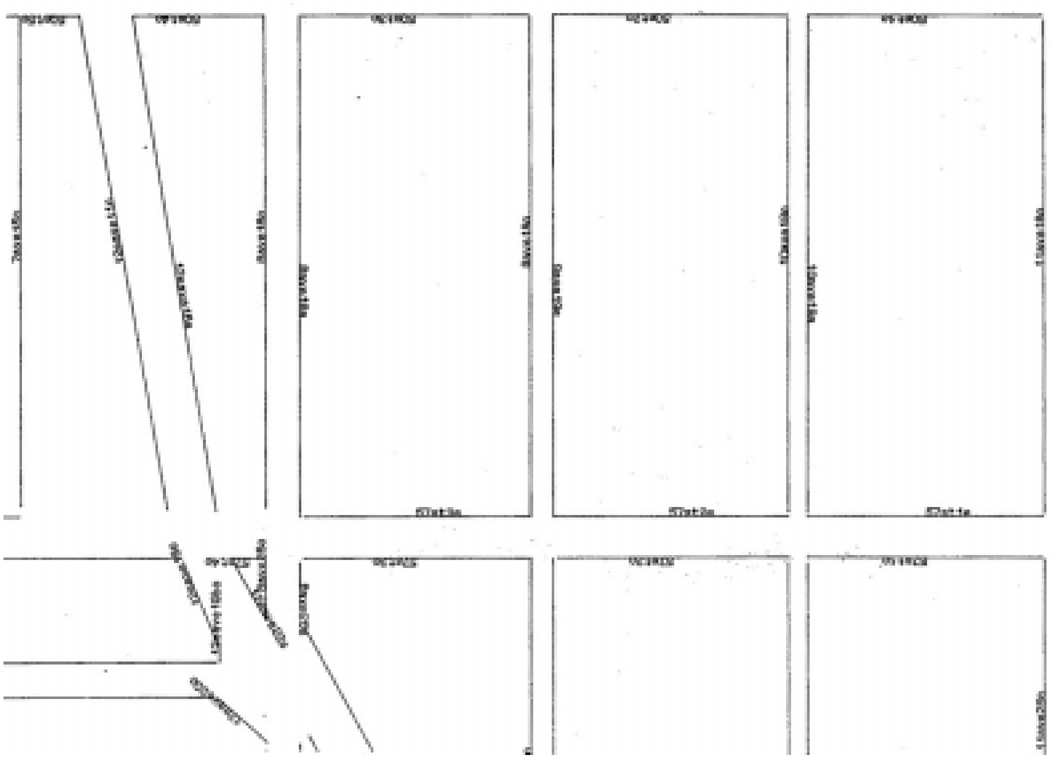

```
12eave15e          :  543 P F   HIM, HE'S A GREAT DRIVER BUT HE'LL SQUEAL). SO
 WHICH COUNTRY DO YOU COME FROM? HOLLAND? VERY NICE, I'VE BEEN THERE... (JONATHA
N'S STILL LISTENING, I THINK HE WORKS FOR
12eave22e          :   70 P F   A NEWSPAPER, BUT WHICH
7ave2o             :  600 P F   BEST WE CAN DO FOR YOU AT THIS MOMENT. WHY DON
'T YOU THINK ABOUT IT AND CALL ME BACK BY THE END OF THE WEEK... ARE YOU FEELING
 ALRIGT?
7ave10o            :  388 P F   KINGDOM OF THE NETHERLANDS BACK THE PRICE OF T
HE TICKET, IN INSTALMENTS. I'M AFRAID THAT'S THE
7ave15o            :  528 P F   IS THAT WE PAY YOUR TICKET BACK HOME. THEN, WH
EN YOU'RE FEELING BETTER, OR CURED, MAYBE EVEN WORKING AGAIN, YOU PAY THE
```

Some qualities of the prototype.

The system ran at about 20 frames per second and with a fairly high resolution (1024 x 768 pixels). The value and desirability of speed and resolution in such an interactive work became very evident.

The placement of the monitor directly in front of a large panoramic museum window was effective - the actual window behind supporting the experience that the monitor was yet another window into a virtual space.

The definitive implementations of The Legible City

Graphics developments.

In 1989, the implementation of The Legible City on a Silicon Graphics Iris 4D/20 allowed the letters to become solid models in a representational space. Custom software was developed to handle level of detail (the letters in the distance only showing their front surfaces, and getting depth when they got closer) as well as the segmentation of the overall database so that only polygons within a defined proximity would be drawn. The use of fog camouflaged the absence of these letters in the far distance, and at the same time gave a very specific aesthetic quality to the work (a sense of infinite distance). Later versions of The Legible City used more powerful Silicon Graphics machines and enabled the addition of textures (for instance cloudy sky, grass, roads, water, etc.) and faster frame rates.

Collision detection (with the letters' walls) was never felt to be necessary in The Legible City. In fact the easy permeability of surfaces seems to be an essential quality of most virtual worlds whereas efforts to simulate artificial barriers seems to betray the 'nature' of virtuality. Certainly bicycling wildly directly through the texts is one of the main pleasures of visiting The Legible City.

The bicycle interface.

In early stages a 'walking' interface for The Legible City was considered. After testing various belt type exercisers it was felt that this would not be convenient, and would not in itself resolve the issue of changing direction. The possible solutions looked highly contrived, and furthermore, the artist wanted an interface that would not cause the 'operator' any performance embarrassments when surrounded by other onlookers.

The bicycle turned out to be an ideal and natural choice in every sense. Using it is, of course, completely transparent and intuitive. It is a familiar device that belongs to everyday experience rather than a machinery that belongs to the media industry. This meant there was virtually no trepidation in its use by a broad public. The particular bicycle chosen - a Strida folding bike - was small and allowed the users (also women with skirts) to easily straddle the seat.

In technical terms, the movements of the pedals and handlebars were connected with high quality potentiometers. A dedicated PC with an A/D board read these signals and custom software then calculated the X and Z positional coordinates of the bicycle in the city groundplan and sent three values (via an RS232 connection) to the Silicon Graphics system.

The belt drive to the rear wheel was modified so that the bicyclist could pedal both forwards and backwards - useful for exact control of ones movements in the virtual space especially when one is reading the texts.

As new cities were added to The Legible City (Amsterdam and Karlsruhe), a button box was put on the bicycle's steering wheel to allow switching between these cities.

The LCD display of the ground plan

A 12" monochrome LCD monitor screen just in front of the bicycle shows a simple ground-plan of each city with a moving dot to indicate the current position of the bicyclist. Switching from one city to another changes this ground plan. Arriving in each city, the starting position of the bicyclist is predefined (for instance in the Amsterdam version it is in front of the Central Station).

This display functions a bit like a normal map - allowing the 'tourist' in this virtual city to orient himself at any time - or to set specific objectives for his journey. The PC that is dedicated to reading the bicycle outputs, is also used to generate this LCD display.

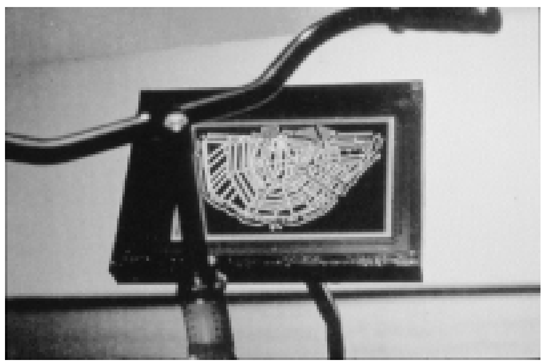

The configuration of the installation space.

After some trial and error, a seemingly 'ideal' specification of the installation space was arrived at. It is a completely black room, approximately 9 m. long, 6 m. wide and 3.5 m. high. The screen is matte white, 5m wide, 3m high, and just 10cm off the floor. It has a slight horizontal curvature (about 30cm deep at the centre). The bicycle is in the central axis of the room, about 4m from the screen, and bolted down to the floor. The high resolution (1024 x 1248 pixel) RGB three tube video projector is hung from the ceiling, approx. 6m away from the screen. A light-lock entry exit system allows people to enter and leave without doors and protects the needed darkness inside the room. Small dim spotlights illuminate the bicycle - these have the same colour temperature as the projected image. A ventilation system will be needed in the room if there are a lot of visitors. Ideally the walls should have some acoustic damping properties to keep exterior noises out.

What is achieved in this room configuration is the sense of a black void with no definable boundaries - just a big bright open window into the virtual urban landscape directly connected to the illuminated bicycle (and bicyclist) in front of it. The proportions are appropriate for a single user and small groups of onlookers, giving a very strong sense of immersion. The curved screen amplifies this immersive effect - it deprives the viewer of a uniform focal distance to the screen, and thus creates a heightened sense of depth in the projected image. The aspect ratio of the screen is wider than normal video, so that the window does not remind one of a TV screen. There is also a feeling of intimacy in this room, allowing easy conversation between the visitors, or a concentrated silence. The only other sound is the whirring of the bicycle wheels.

The Amsterdam version.

The area represented is the old inner city of Amsterdam as far as its 19th century boundary. In this version all the letters are scaled so that they have the same proportions and location as the actual buildings which they replace, resulting in a transformed representation of the real architectural contours and features of this city. The colouring of the letters matches the brick and stone tones of the real buildings. The texts are factual and are derived from archive documents which record actual events in Amsterdam from the 15th to the 19th centuries. These texts are located in those areas of Amsterdam to which they refer. Also the original vocabulary and spelling found in those texts is respected.

In creating the database for this (and the Karlsruhe) version of The Legible City, certain features were added to the text editor. As letters were entered along the street segments, height, width, and depth values were specified (in concordance with the height, width and depth data of the real buildings at those locations!). A colour palette was defined for the whole of the city, and shades of these colours were randomly distributed by the software.

This Amsterdam version very much benefited from texturing of the water in canals.

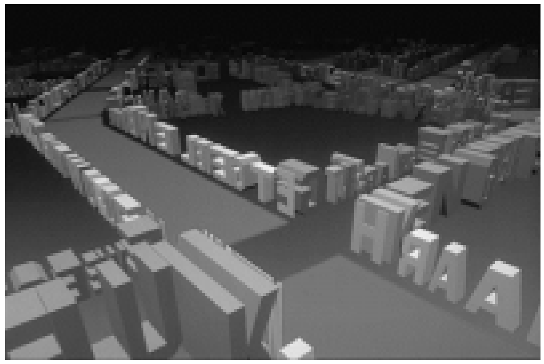

The Karlsruhe version

Here the area represented is delimited by the Karl Wilhelm Schloss, the Fritz-Erler Straße, the Karlstraße and the Kriegstraße. The text also follows the Ettlinger Straße down to the station, ending at the site of the originally planned ZKM building. The texts are largely based on existing historical accounts relating to this city. There are references to certain people who were interesting residents of this city such as Karl von Dries, the inventor of the 'walking bicycle'. Also texts have been quoted from contemporary 'promotional' brochures published by the city of Karlsruhe. In those areas of Karlsruhe where the existing architecture is fairly uniform the letter proportions have been abstracted to one size, while for those buildings which are outstanding the letters have been positioned and scaled so as to describe their actual appearance.

Whereas both the Amsterdam and Manhattan version use black fog and dark skies, setting them in a moody night-time context, the Karlsruhe version has a bright sky and white fog, setting it in a mysteriously misty daytime situation. And in all three there is a particular loneliness for the bicyclist, of being the only person in town (or at least out in the streets).

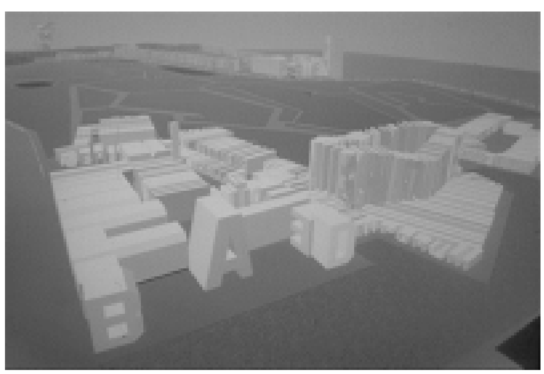

Scaling relationships

As the letters in The Legible were scaled to actual building size, a reasonable speed of movement was desirable to allow slow yet smooth reading of the text. After testing this turned out to be about 2.5x normal bicycling speed. As the ground plans of each of the cities were also full scale, constituting many kilometres of travelable streets, this 2.5x bicycle speed was also a perfect balance between achieving a realistic sense of real world expenditure of effort and time, while giving the satisfaction of being able to cover long distances in the virtual world without too much effort and time.

The lesson here seems to be that the various scaling relationships between the real and virtual not necessarily being 1:1, can still maintain a very convincing sense of experiential realism. Of course one can expect limits where this congruence will break down.

Physical effort in the virtual world

Because of the nature of the interface in The Legible City, the physical effort of bicycling in the real world is gratuitously transposed into the virtual world, affirming the absurd yet euphoric conjunction of the zealous body in the virtual domain. This is not the case with conventional interfaces - keyboard, mouse, joystick, etc. - which transpose negligible displacements of the body into media coordinates. While the bicycle cannot be considered a 'force feedback' device (though simulation of road texture and turning motions could be achieved through appropriate engineering solutions), there is a indirect experiential feedback in the tired limbs and exhaustion the bicyclist experiences after an hour or so bicycling in the virtual world. This 'realism' of this body experience is technically conjoined to the viewer's movements through the virtual scene, and thus constitute a strange level of validation of those virtual scenes. One's body insists that one has really been somewhere, even though the eyes know it was just a fantasy.

One may say that such an interface demonstrate how our presence in the virtual domain can be afforded by a type of delicate prosthesis that can give us pseudo-tactile access to what is untouchable. This is quite different to the banal simulation of literal access afforded by force-feedback devices such as the 'Phantom'.

Conjunction of real and virtual spaces.

The Legible Cities seem to be quite unreal places because of their lettered architectures. And yet they embody the groundplans of real cities, and in the case of the Amsterdam and Karlsruhe versions the letters themselves express the proportions of buildings they replace. A paradoxical 'meta-realism' is thus achieved. For instance if the installation is set up in Amsterdam, one can bicycle to the virtual spot (inside a letter) where the actual bicycle is located! Many bicyclists want specifically to visit the letters (or words) that are located in familiar places (e.g. their home). And after a few hours of bicycling in the lettered city and then stepping out into its real equivalent, one is highly sensitised to all the textual messages (names, labels, ads, directions, etc.) that a city contains in its normal state.

The realism of virtuality does not then necessarily mean a literal representational realism. Only an experiential realism needs to be achieved, and a smooth flow of meaningful information between the poetics of virtual worlds and the re-identified structures of the real world. The Legible City sets up a sort of magical black box (in the tradition of a 'camera obscura') where the bicycle is an avatar of the real world that can be ridden through a wide open window into the virtual world, only to discover there that this abstract virtual world is an imaginative (informative) mirror of the real, and that one hasn't moved from the place where one began. In matters of virtuality, this type of congruence seems to be more important than representational issues. And if this is true, it has important implications for the direction of development of appropriate hardware and software tools.

The further development of The Legible City within eSCAPE

The successes of the single user The Legible City seem to indicate it is a promising context to conduct research into shared and multi-user environments. Also the bicycle seems to be an appropriate and efficient method of locomotion for large groups of individuals in a distributed informational environment (not to mention the benefits of the physical exercise). The future proliferation of networked graphics capable set-top computers, as well as home bicycle trainers, could lead to a method of hybridisation that could make The Legible City (in all its possible forms) into a generic household utility.

The basic aim is to achieve a distributed version of The Legible City whereby many 'bicyclists' at separated sites are able to simultaneously explore this virtual world, meet there, and socialize with each other.

Objectives:

-- First two, and then many more 'bicyclists' should be able to explore the same virtual space.

-- Bicyclists should be able to see each other's actual location and movements on a 2D ground plan display. This enables a bicyclist to know that someone else is present without actually being close enough to see them, and thus enables bicyclists to move towards (or away) from each other.

-- Currently a bicyclist is represented by a blinking dot on this ground plan - in the distributed version possibly each bicyclist should have a different 'cursor' shape.

-- When bicyclists get close enough to see each other, there should be some form of simplified representation of that other person (and oneself of course). A personally articulated 'avatar' is needed for this - just a common symbolic image of a figure on a bicycle. This is a design issue that needs special care (e.g. visualizing the 3D orientation of these figures).

-- The main communicational emphasis in these meetings should be done through audio connectivity - allowing the bicyclists to easily talk to each other when they get close to each other, with the appropriate spatialisation as they move between distance and proximity. This approach permits simplification of the visual representation of the bicyclists (which would also not demand much graphics performance).

-- What have these people got to talk about? Something about the shared experience of being in this idiosyncratic virtual space - as sometimes happens to visitors to any real world event who might strike up a conversation based on a shared (or disputed) experience. Except that here there is a curious, contradictory conjunction of physical proximity and distance that will give these meetings their special flavour.

-- People who meet in The Legible City might wish to bicycle together for a while, so there should be a 'tandem' function - where one person allows himself to become the other's 'passenger'. And they can switch between these two roles (or separate) whenever they want.

-- Possibly more than two person's can join this 'bicycle train'. And possibly these 'joined' bicyclists all maintain a certain level of individual control by some kind of computational response to their various directives. Force feedback joysticks could be an appropriate embodiment in this situation (or in the deluxe version - force feedback in the bicycles).

-- The possibility of integrating speech recognition. This could lead to the bicyclists retextualising the city themselves (or alternatively reserve just one district for "graffiti" activities).

--The Legible City should be able to be concurrently operable with a variety of interface devices - ranging from a generic cheap joystick to the high-end custom made bicycle. Some interesting possibilities are: force feedback joysticks, game type joysticks in the form of a steering wheel, simple adaptations of off-the-shelf road bicycles and exercise bicycles, a deluxe custom bicycle with force feedback and motion control.

Implementing of a Multi-user Legible City

As part of the work for the coming year we envisage developing a multi-user version of the legibile city that will alllow a number of users to share the experience of engagement with the legible city. In addition, we envisage this development linking with the Cityscape described in deliverables 2 and 3 in an attempt to provide an alternative way of experiencing an e-scope. What follows is a brief outlining of the desired capabilities of a muli-user legible city and an outlining of a suitable supporting architecture.

Desired capabilities

-- approximately the same graphics as ”classic” legible city
-- implementation on consumer level computer with joystick input
-- addition of player avatars to scene (animated biker)
-- possibility of several players in one city (multiplayer capability)
-- possibility of scene-level interaction (i.e. joining of two players to form a tandem)
-- player interaction through voice which is proximity controlled (”you hear X when X is close...”)

Implementation

The system will use a standard client-server structure. The server will maintain a consistent state for the virtual city, provide networking services and is the controlling platform for the audio system. The client will provide the visual output and acquire user input. Audio interaction on client side will make use of a separate phone (or headset).

Diagram of the Client

The main design aim is to keep the cost of the client low and free of licensing restrictions to aid in placement of the client in many locations. Thus we aim to provide a “thin client” with the majority of processing taking place on the server side of the architecture.

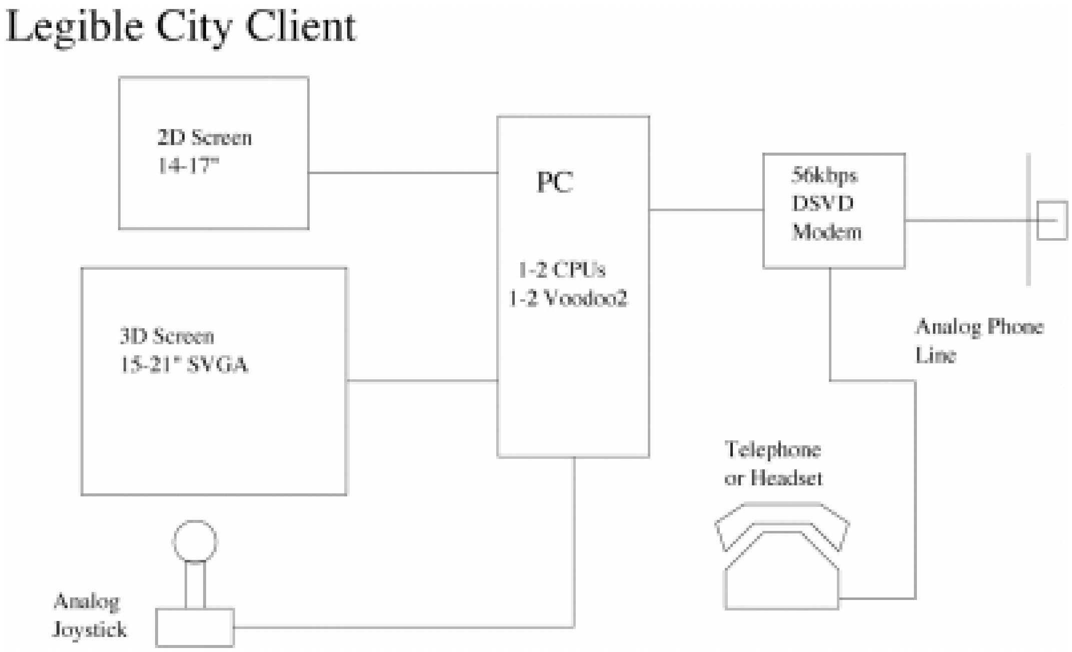

The target hardware platform consists of the following components:

-- Dual PentiumII 333MHz CPU (minimally: single PentiumII 266MHz CPU)
-- Dual Voodoo2, 12MB graphics card (minimally: single Voodoo2, 12MB graphics card)
-- Dual monitor (minimally: single monitor)
-- Force feedback joystick (minimally: standard joystick)
-- DSVD modem (minimally: internet connection)
-- Telephone and analogue phone line

The target hardware will provide a high quality visual experience on the main (3D) monitor and a map on the second (2D) monitor – emulating the original LC set-up as much as possible. The dual CPU set-up should ensure enough processing power even when many avatars are in the field of view. The fallback to a single CPU will slow down the graphics rendering somewhat. The dual Voodoo2 set-up ensures fast graphics at 1024x768 resolution. If only a single Voodoo2 card is present the possible 3D resolution is only 800x600. If only one monitor is wanted or available, a fallback in form of a 2D overlay map in a corner of the 3D screen should be implemented. The force feedback joystick is being used to physically involve the player, if present. The DSVD modem guarantees the availability of audio while the line is being used for a data connection. If such a modem is not used, a regular Internet connection can be made to the DEVA server but audio capabilities will not be available. A regular telephone set which is connected to the DSVD modem serves as a reliable interface for the player. The target software components are as follows:

-- Linux distribution (with SMP enabled) installed and configured
-- Glide/Linux (Voodoo2 driver) installed

-- Mesa (OpenGL driver) installed
-- Maverik installed
-- LegibleCity software installed

Diagram of the Server

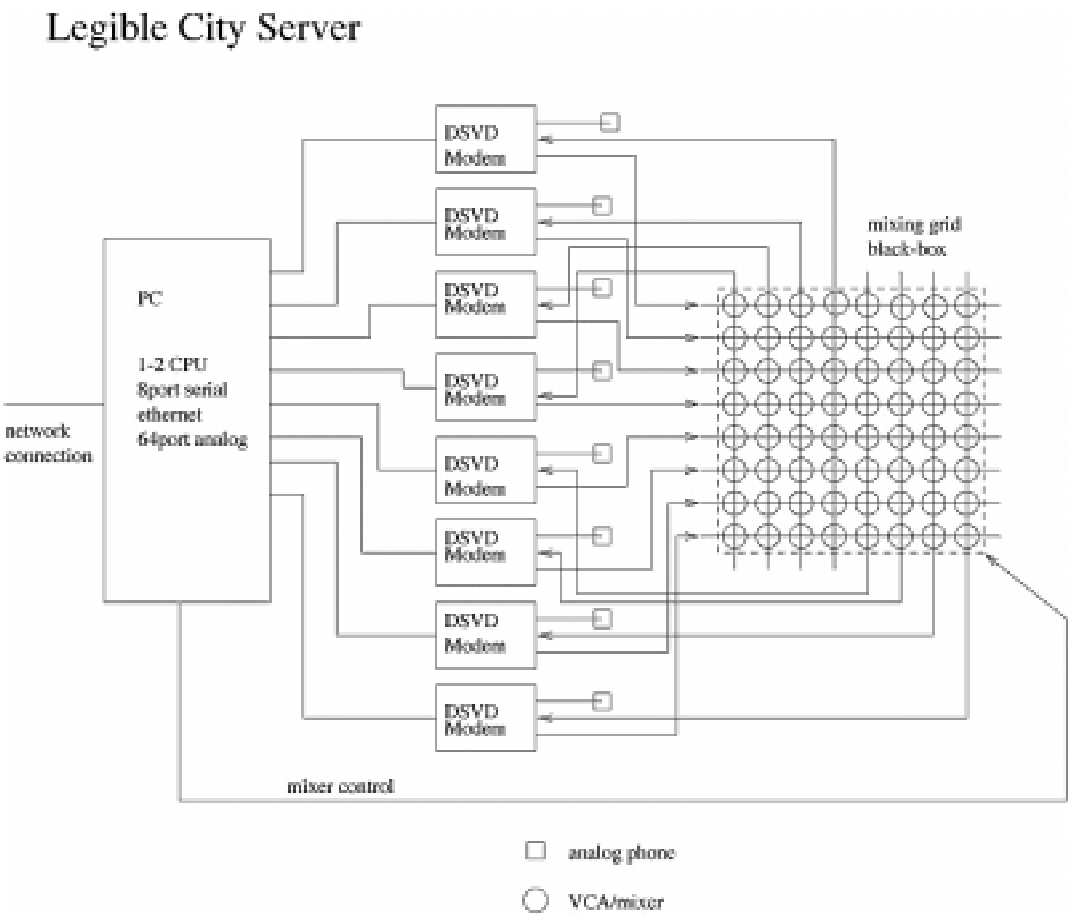

The server target platform is as follows:

-- Dual PentiumII 333MHz CPU (minimally: single PentiumII 266MHz CPU)
-- Ethernet interface
-- Multiport serial interface (8 ports)
-- Programmable digital I/O ports (i.e. printer port)
-- With software
-- Linux distribution (with SMP enabled) installed and configured

Since DEVA is not tested reliably on the Linux/PC platform (April 98) the initial DEVA server will run on an SGI. It is however necessary to move to the PC platform for implementation of the 8-modem dialup and audio mixing hardware.

Mixer implementation details

The mixer can be easily designed using special analogue chips and standard operational amplifiers. The Dallas Semiconductor chip DA 1806 provides 6 digitally controllable resistances (potentiometers, pot's). It can be programmed through a simple serial protocol with 3 wires. 11 chips can be daisy-chained to provide a total of 64 (+2 spare) pot's. 8 pot's and 9 OP-AMPs (on 3 chips) as a unit allow for controlled mixing of 8 input signals to one output signal. The input signals are the seven other players that are connected via DSVD modem plus one spare input that could be used for an ambient soundtrack or sound effects. The controlling computer has to serially clock out 528 bits which can easily be done through a PCs printer port.

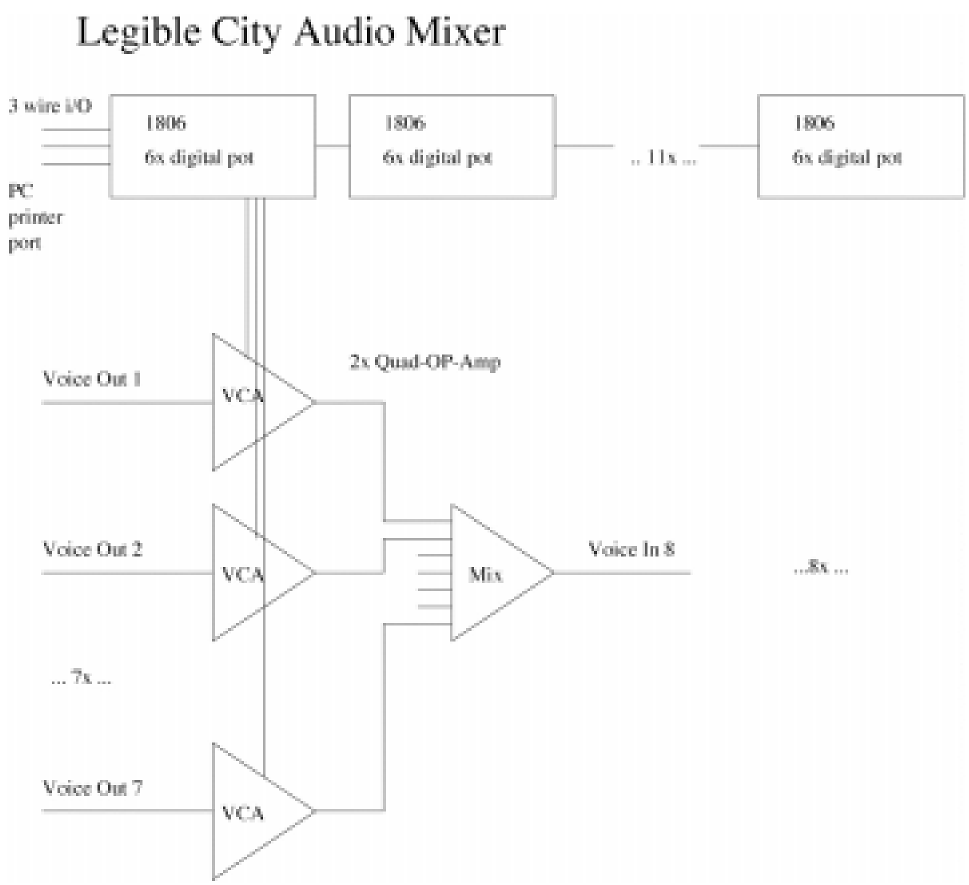

Design Tasks

-- purchase special hardware components for experimentation
-- refine existing LegibleCity (LC) port to Maverick on target platform
-- for speed (SMP support for dual-CPU PCs, refine models and culling)
-- for image quality (add textures and fogging)
-- for input device (analogue joystick, use of force feedback extension)

-- for map support (city-map and locator on 2D second monitor or as overlay to 3D screen)
-- create animated biker-avatar in DEVA and place it in LC
-- experiment with simple multi-user LC (DEVA running on Irix/SGI)
-- implement server (DEVA running on Linux/PC)
-- add dialup-support to server (1-8 analogue phone-lines required)
-- build audio mixer hardware
-- add audio mixer control to DEVA code

Special hardware:

-- Voodoo2 equipped graphics cards (i.e. Diamond Monster 3d II, 12MB)
-- 56k modem with DSVD (Digital Simultaneous Voice and Data)
-- Multiport serial card (various vendors)
-- Force-Feedback Joystick (i.e. ForceFX, SidewinderPro Force)
-- Custom mixer-hardware using Dallas Semiconductor DS1806 digital potentiometer
-- Standard Telephone

Chapter Four: Interaction and Presence in Shared Electronic Environments: fieldwork at ZKM

Monika Bűscher, Jon O'Brien, John Hughes

Lancaster University

The centre for art and media technology (ZKM) in Karlsruhe was founded as a new type of institution that brings together art and technology in an unprecedented and unrivalled way. It combines two research institutes - the Institute for Music and Acoustics, and the Institute for Visual Media -, a Media Library, and three museums: the Museum for Contemporary Art, the City Gallery, and the Media Museum. 12 years after its conception, ZKM was opened to the public in October 1997. This event was celebrated with a number of events and the 5th Multimediale (a bi-annual exhibition organised by ZKM since 1989). Until the 9th of November the Media Museum and the Multimediale5 presented some 40 installations that were arranged on three floors in the former machine and ammunition factory building that houses ZKM. The permanent exhibition at the Media Museum now on show contains around 30 works. In addition, a multitude of video installations, photographies, and paintings are presented to the public in the Museum for Contemporary Art and the City Gallery. The ethnographic study reported on in this chapter was carried out over a total of 9 weeks, covering the whole of the ZKM's opening, and 3 weeks of 'normal' opening in December 97 and March 98.

This chapter explores issues of presence and representation drawing on fieldwork undertaken in the media museum of ZKM, where a number of shared electronic environments were made available to the public. One of the key project themes has been the explorationof means of facilitating movement between 'worlds' in an e-scape cityscape or universe in a way that is socially and aesthetically rich, intuitive, and pleasurable, something that is particularly critical in view of the fact that the aim is to make e-scape environments available to the general citizen. One means of achieving such intuitive and legible design of large scale electronic landscapes has been the study of people's strategies employed in moving around real world space.

We report observations from these studies in this chapter here, but note that such observations in real world settings alone do not suffice to inform the design of CVE’s. Mediating technologies, while in many respects extending the possibilities for human interaction, inevitably alter and impoverish the richness of interactionally relevant information we take for granted in our everyday face-to-face interactions. Therefore, it is necessary to observe people’s interactions in and with such environments. However, despite the emergence of online electronic environments these still remain a specialist concern and are seldom used by general citizens - there are few opportunities to observe real world situations in which people begin to inhabit electronic spaces. To overcome this problem, we undertook a series of studies at ZKM.[1]

We put forward a set of issues that emerge from the analysis of our observations of the sense of presence experienced by museum visitors using these virtual environments within the context of the real world museum. Essentially, video, audio, and observational data suggest that people adapt their practices to the affordances of electronic environments (Gaver 1992, Gibson 1986). Crucially, everyday practices of orientation, movement, and interaction in space are drawn upon, but *transposed* rather than transplanted in order to fit in with the affordances of the environment. Moreover, a sense of presence is not ‘split off’ from presence in the ‘real world’ but, rather, *extended* from the real to the virtual space. This means that the practical engagement with electronic environments is one in which activities in the real space are relevant to the activities and events in the virtual space, and vice versa; and that in practice the interface and the screen merge into an *instrument* with which the electronic space is seen and experienced. Some of the practices involved in the management of this relationship between activities in the physical space and those in the electronic space will be addressed. However, the main issue we are concerned with in this chapter is the question of *how* people adapt to the affordances of the electronic environments. Seeing ‘ordinary’ people use different environments has allowed us to observe some of the issues involved in the successful transposition of everyday practices and the sense of presence that is thus achieved.

Presence in electronic environments

As is discussed in Appendix One of this Deliverable, a sense of presence as a sense of ‘being here’ is important in our interaction with others and with the physical world around us. In this chapter we explore interactional presence –

[1] As pointed out in the introduction, the ethnographic study made use of the fact that the works of art could be seen as ‘breaching experiments’. People often visibly experienced puzzlement with regard to some of the features of the environments. We focus on such instances, because they give us a unique opportunity to understand the intersubjective organization of the electronic spaces interactive multimedia art installations provide with a view to informing the design of eSCAPES. The point we hope to have clarified with this brief methodological note is that we are interested in the interaction *with* art, not in judgements about the aesthetic, conceptual, or interactive value of individual exhibits. We are in no position to make such judgements and have no desire to do so.

how we are present as participants in interactions with others – and a sense of presence in relation to objects and spaces in electronic environments. However, these are not two distinct forms of experience but integral parts of our experience of being in the world.

Interactional presence

Focused encounters (Goffman 1961) in face to face situations are characterised by a rich tapestry of interactional information. Through our position in relation to others, our posture, movements, the direction, intensity, and duration of our gaze, and other finely tuned embodied actions we occupy a place in the encounter, where we make available to others what actions they can reasonably expect us to take within the frame of the encounter (e.g. Simmel 1970, Ciolek 1980, Heath 1986, Kendon 1990, Robertson 1997). We dynamically and flexibly display the degree of our involvement in the interaction, how we understand others' actions and our own in its context, whether we are listening or whether we want to say something, and anything else that may be relevant at the time. Similarly, there is a large scale spatial dimension of interaction between people in spaces. Moving in public places, we routinely gauge information about who else is there and what their activities are, and weave our own actions into the flow - not least to avoid 'collisions', but also to see who we could talk to, how long we would have to wait for our turn if we were to join a queue, or where there is something to see (Sacks 1992, Sudnow 1972). In the mediated variation of such encounters, this tapestry of information is considerably impoverished. In exploring the ways in which people manage to conduct orderly encounters under such conditions we hope to inform the design of e-scapes, accountsof which are offered in the final four chapters of this deliverable.

The hunter's perspective: Labyrinthos

Labyrinthos is a collaborative environment. It is a game reminiscent of Midi-Maze, the first networked game for home computers, that allows eight players to move through a model of the exhibition space and 'shoot' each other. Each player controls a coloured sphere through a simple control (figure 3). The terminals have different colours that correspond with the colours of the avatars. If an avatar is repeatedly hit, its face changes expression. Eventually it becomes transparent and immobile. The successful 'hunter' on the other hand gets a crown. Here, three boys are hunting in a 'pack'. Y is sitting at the yellow terminal, O at the orange one to his right, and the third boy (T) is moving from one to the other (figure 1 &2)[1]:

---

[1] Most of the transcripts presented in this chapter are translated from German. Therefore, some expressions and overlaps ([xx]) do not exactly correspond with the original.

Figure 1: Y, T and O at their terminals

Figure 2: Y bending to see O's screen

O: hey I'm alive again!
T: eh cool (d'y' know)
→ [where you are right now?)]
→ Y: [eh? Where are you?] (0.) eh ((glances across)) where are you::?
O: I don't know either I'm back at (the square there)
?: (xxx face)
→ Y: WHERE?
O: there comes purple
T: ah!
O: eh!
→ Y: ((bends across and back while saying)) I I'm coming where are
you? Oh there
?: (xx)
T: blue!
Y: I've got gre- I see red I SEE RED I I'm coming to help you I've
got green under attack ((starts shooting))
O: me too I'm shooting=
Y: =eh? (xx but you said xx)
O: shit
T: he! (xx)
Y: die you (shit)
?: attention!
→ T: eh! Red is behind you
→ O: red is behind you (0.)
→ Y: where? (.) ni- attention (.)

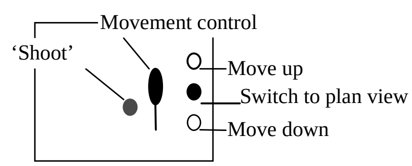

Figure 3: Labyrinthos interface

Y and O move their avatars through the two floors of a simulation of the real space and some of the exhibits of the media museum. In order to be able to play successfully, they need to know *where* they are, in relation to potential targets and each other. At the beginning of this excerpt of talk, O's avatar suddenly responds to his manipulations of the control after a period of immobility. Immediately, he makes public that he is once more an active player in the game. Now his position in the space becomes an issue. In order to 'attack' their opponents effectively, the three have decided to 'pack up' and pursue them. They need to assume strategic positions and coordinate their actions, and in their approach take advantage of the fact that they are distributed across a large distance in the electronic space, yet co-located in the real space. Y asks 'where are you' while moving forward, but the answer is too vague for him to find O. His third request 'WHERE?' is cut short by the fact that O discovers 'purple' and is put on the spot. Shots are exchanged and help seems required. Y, in one move, bends across, glances at O's screen, and comments on his activities. Having seen what O sees, Y is able to work out where he is. However, on his way he encounters first ‘red’, then ‘green’. He attacks ‘green’. What he fails to realise is that while he is busy shooting at ‘green’, ‘red’ has moved behind him and is attacking him. The others notice, because they can see ‘red’ attacking ‘yellow’ on O’s screen. They warn him. A little later on in this game, O ‘kills’ purple. He switches to plan view, where the avatars are represented as coloured arrow-heads, and notices that ‘yellow’ is being attacked by ‘green’.

Y’s move in line 11 rightly assumes that the space of the game is the same for everybody involved and thus projects a principle we routinely apply in our everyday actions into the electronic space. As a ‘hunter’ in the game Y does and, in fact, needs to take for granted that views from different positions in the game are interchangeable. The space *must* be the same for everyone, otherwise one would not be able to aim at targets. This is necessarily reflected in the visual display of one’s position in the space on the respective screens of the players. Y exploits this fact. Glancing at O’s screen he assumes that ‘if I change place with him so that his “here” becomes mine, I would be at the same distance from things and see them in the same typicality as he … does’ (Schutz 1962, p. 183). Seeing what O sees allows Y to locate him on his own screen. If the game took place in physical space, Y would not be able to put himself in O’s place so easily. Here he makes use of the affordances of the situation and transposes his knowledge of the interchangeability of perspectives without hesitation.

This ‘reciprocity of perspective’ is one aspect of a more general principle of *intersubjectivity* that underpins interactions between people and between people and their material environment. As we have outlined in Chapter One of this Deliverable, what we say, perceive, or do is part and parcel of a world known in common (Schutz 1962). In our actions, we assume that others know the world and the situation at hand in ways that are similar to how we know them, and we assume that the material world is arranged in a way that draws on and refers to such common knowledge. Thus, not only do we assume to see the same objects and events that others see, we also take for granted that we interpret them - for all practical purposes - in the same way as they do. This includes people’s distribution, appearance, and activities in space. Usually we have a whole host of clues that indicate *at a glance* ‘what sort of people’ there are and ‘what they are doing’, reaching from clothing to movements, gestures, and facial expressions. We use this information to categorise what we see and insert our own actions into the whole (Sacks 1992). This, too, is reciprocal. We see, but at the same time, we are seen, and we know it. We know *that* we are seen, but also know *what* people are most likely to see. In this instance, position, colour, ‘shooting or not shooting’ are the only clues available. However, according to the rule of relevance outlined in Chapter One this information is sufficient in the frame of the game to allow the players to monitor the state of affairs. In lines 21-23 competence is enacted in a way that mirrors everyday practices. From O’s point of view displayed on O’s screen, O and T see that Y is being attacked from behind his back. Later on in the game, similar information is used in a different way. On the plan view, O discovers that Y is being attacked. Here, O does not see the bullets hitting Y, but a semi-circular configuration of coloured arrowheads pointing at Y. Such an abstracted, animated, real time plan view is a device usually unavailable to us. It is an affordance that exceeds everyday environmental conditions. O transposes and combines competencies of categorising people (‘hunter’, ‘target’, ‘us’ and ‘them’) drawing on a minimal and abstracted set of clues, and his knowledge of real world plan views to fit in with the environmental conditions of the game.

The limits of multiple perspectives: The World Generator

The ‘World generator’, exhibited during the Multimediale5 at ZKM, combines a videolink to visitors at an exhibition in Nottingham and a networked electronic environment that invites people at both sites to choose three dimensional shapes and fragments of text, music, film, or photographs from a menu and to place them into the three dimensional electronic space. A *spacemaster* allows two people - one in each location - to navigate through the space, and to select objects in their vicinity and alter their appearance. Each of them is represented on their remote partner’s screen through a tower-shaped avatar, but there is no avatar for the local person on the local screen. At the same time, a videophone provides a visual and auditory link. The image transmitted via the videophone is mapped onto the avatar in the electronic space. This work provides an environment for a more focused encounter, where two participants can arrange objects in the electronic space and explore their joint arrangements. However, at the same time, *The World Generator* explores the experience of multiple perspectives onto a joint resource – the space and the objects placed into it by the visitors at each location. The following excerpts of talk between a visitor at ZKM (K), the artist (A), and visitors in Nottingham (N, N1) shed some light on the difficulties experienced and the strategies employed to work around them:

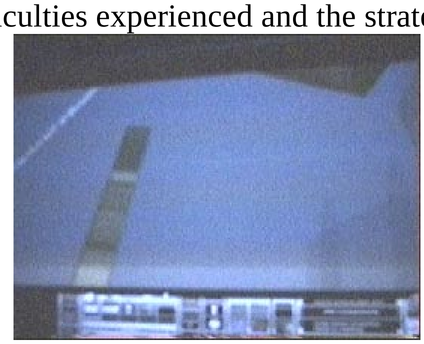

Figure 4: Chasing N’s avatar

Figure 5: The videophone

1	A:	so he’s actually looking this way
2 →	K:	I see I I didn’t- I thought he has the identic- ((speaking into telephone receiver)) OK
3	A:	so if you want you can kind of come up next to his avatar
4 →	K:	how could I how could I *get* his perspective?
5	A:	ehmm OK
6	K:	I'm trying to get your perspective to see the same thing you see (.) just that we have the same experience
7	A:	it's not so easy
8 →	K:	it's not so ea-even for the artist himself ((points to A)) it's not so err ((looking at videophone screen))
9	A:	where did he go? *There* he is. OK we’ll just chase him. OK he’s up there. ((quietly)): what does he see? He’s looking off into
10		nothingness
11 →	K:	what do you see? Tell me, describe it

12 N: (a flying object on the right side of the screen … *possibly* walking sticks (hats) (0.3)
13 K: pretty good
14 N1: (it's coming round)
15 N: that?
16 N1: yeah
17 → N: looks like a tree ((K looks at the videophone screen))
18 N1: that looks like a- a light a- a streetlight that way round
19 N: (xx) yeah a streetlight (0.4)
→ ((K turns to A, then to controls, switches to full screen menu and begins to choose a texture))
20 N: an object on the right going up and down (0.4)
→ ((K switches back to the view of the space))
21 K: I'm amazed
22 N1: (I'm off) see you later
23 N: (xx) (0.4)
24 → N: Hi here's Chris back
25 K: hi (0.4) hi how's the life?
26 N: (xx)
27 K: nice. (xx) everything's lost
28 N: have you got my viewpoint yet?
29 K: I can't see in the moment anything but I am on the way (0.8)
30 → K: what can you see now?
31 N: err (xxx) I can see a oilrigg on the right hand side of the screen going up and down
32 K: right
33 N: and the (0.3) errm (0.3) over on the left side of the screen I can see a tree
34 K: tree?
35 N: (floating around)
36 K: a tree?
37 A: [try to (xxx)]
38 N: [(xxx)] (xxxx) turning round
39 K: you said tree?
40 N: (xxx)
41 → K: a tree yeah mhm I wonder what what you see really. (0.3) I see something too but it doesn't look like a tree but maybe

The nature of N's presence is a source of interactional trouble, as it is difficult to determine his position in the space and his orientation to, and perspective of, objects in the electronic space, and his engagement in the interaction. The distortion of the principle of reciprocal perspectives that is at work here is initially not accessible to K, who assumed that N was seeing the same part of the 'world' as he does. When it does surface as one of the reasons for the difficulties the two have in coordinating their activities, K asks 'how could I get his perspective' (lines 2, 4). Cooperative activities in our everyday life that parallel this situation provide us with ample information about the task at hand and the nature, course, and potential future developments of the interaction in a way that is taken for granted. We can *see* what is going on between us and the other people, whose turn it is in the talk, where the other persons' focus of attention is, whether they see what we see or their view is obstructed. The lack of information of this kind in the 'World generator' is made explicit in the above stretch of talk. In order to find some common ground K needs to find an answer to 'where am I?' and 'where is he?'. Finding out what N sees is, again, a strategy that relies on the principle of reciprocal perspectives. But how can it be done? N could be anywhere in the large expanse of the electronic space. Two strategies are employed to repair the problem. In the absence of a plan view and given the small section of space visible on the screen, the artist helps K to 'catch' N's avatar. However, even once it is found, its shape gives no indication of what he might be looking at, and K asks ‘what do you see?’ (line 11). In order to make this a useful strategy, K would have to determine what N is seeing and bring his viewpoint into a position that is parallel to, or behind, N’s point of view. Then they could place objects in front of them and both would be able to see them. But a cooperative engagement in arranging objects making use of the *space* would require N to shadow K or vice versa, a task that is beyond K’s level of familiarity with the *spacemaster*. As another means of establishing at least some workable level of common ground, K asks for descriptions of what N sees (line 30), without trying to shadow his avatar. But this, too, turns out to be a difficult task, because the shapes and images given in the menu do not readily conform to categories of objects in our everyday experience. What shapes could be seen as is a matter of negotiation (lines 17-19, 41). A further means of dealing with this problem, observed on a different occasion, was to turn the videophone around to capture the local screen, allowing the remote partner to *see* what was seen in Karlsruhe.

The videophone is usually used as a resource for interactional coordination in a different way. K and N look at each other while talking, and they can see who else is involved in the interaction and how (see, for example, K pointing at A in line 8). A little later on, N1 gets engaged in determining what exactly a shape on the screen looks like (lines 14-19). K looks at the screen of the videophone. Having seen N and N1’s orientation towards each other on the screen, K treats this as a private conversation and turns to the people around him, then to the controls. He switches to the full screen view of the menu and begins to choose textures to place into the space. However, N’s next turn is addressed to him, and he quickly switches back to the view of the space. Such a window onto the physical surroundings in this *mixed reality* environment (Benford 1996) prevents problems such as the ‘corpsing’ experienced as a source of confusion in MASSIVE (Bowers et al 1996). There, avatars whose ‘owners’ were engaged in activities in their physical environment or for other reasons unavailable for interaction in the virtual space (they could even have left the room) showed no sign of their owner’s lack of involvement. In the *World generator* the videophone prevents such difficulties. However, the information it allows a remote partner to gauge is very limited. The camera only captures a fraction of the space that is potentially relevant to K or N’s activities, and the quality of the picture is affected by the poor display and the timelag in the transmission. The fact that it is possible, but difficult to judge what is happening from the video image alone is acknowledged through N’s announcement that he is now back ‘for real’ in line 24. But the above example illustrates an actually rather rare event of people trying to engage *cooperatively* with the work. The following extract from the field notes gives an impression of a more typical course of events:

A,B,C leave. A woman who’d been watching sits down. She doesn’t get a chance to do much, because some girls sit down at the other end and start talking. They exchange names: N, N1 (Nottingham), K (Karlsruhe) - N wants to know what the world generator does. K struggles to explain, ‘bombarded’ with other questions:

N/N1 wants to know *where* K is. Doesn't believe her when she says Karlsruhe. Is it an exhibition? Who else is there? How old are you? They are twenty. Then agree that there isn't much point in talking if they can't do anything together. K puts the phone down. Leaves shortly after. When the next person sits down, the girls come back. They wave their hands in front of the camera and talk but there is no sound. They shout 'pick up the phone', which is audible even though the receiver is on the hook. K2 picks up the receiver.

People mostly *either* focused on the screen and engaged with that part of the installation, *or* engaged in conversations with people via the videophone. While the videophone provides a link between the real spaces at each site, it thus does not generally assist in the establishment of information about the remote partners' position and orientation within the electronic space. In fact, unless it is turned around to capture the screen, the videophone tends to create a parallel space that is attended to separately.

There are four different ways of adapting everyday practices to the affordances of this environment. Three of these – trying to map one's own point of view onto that of the remote partner by physically locating his avatar, descriptively, or through sending one's own point of view ('I am here') via the videophone - attempt to access interactional information that would be relevant to the task of cooperatively arranging objects in the electronic 'world'. But none of these strategies succeed in creating a rich enough surrogate of reciprocal perspectives to allow for cooperation. Engagement with the work was one in which people explored multiple perspectives onto a shared space. The fourth strategy arises in the context of the parallel media space created through the videophone. It illustrates some of the more *instrumental* nature of the technologies involved in extending the spaces we inhabit. Waving one's hand in front of someone's head, or, more specifically their eyes, is a gesture often employed to attract the attention of a person who is looking at something else, or lost in thought taking them 'miles away'. The sequence of events where one person's waving of a hand in front of the videophone camera attracts the attention of the person at the remote site understands the camera as an extension of the eye of the one being waved at (cf. Heath and Luff 1991).

Being 'here' and 'there': The Fruitmachine

'The Fruitmachine' provides a different environment for cooperative action. A three dimensional octagonal shape is cut into three pieces and suspended in an electronic space (see figure 6). It can be put together through the coordinated actions of three players, each controlling one piece via a spacemouse. Here a family with two young children (1) and a group adolescents (2) are engaged with the work:

Figure 6: The Fruit Machine

(1)

1	V:	you're probably the one at the back (.) Laura
2	L:	yes but I'm *not doing anything* (0.4)
3	V:	that's me (.)
4	L:	the one at the front that's me
5	M:	no that's Dad (0.4)
6	K:	which one am I?
7	L:	you're tha that one (.) and I'm the [biggest one]
8	K:	[which one?]
9	L:	you're that one that's just coming to the front (.) no that's *me.* That's *me.*

(2)

1	C:	there is there is one (.) and and if we
2		[got those together now]
3 →	A:	[wait I'm coming] wait
4	B:	ah I have bananas too

The children and their parents in example (1) squabble about who 'is' who in the jigsaw. L assumes that she can just 'pick' the shape that appeals to her most ('the biggest one'), while the others are trying to determine which control, and thus whose actions, are related to which piece. K moves her piece in the course of finding out which one she 'is', prompting vehement protest from L, who has decided that this is the one she wants to 'be'. The conversational reference to the space on the screen and the objects in it resembles that of talk around a board game[1]. On the other hand, the inflexible link between control and piece makes the electronic and material parts of the installation into an instrument that

[1] Some fieldwork in a pre-school centre displays parallels between conversational practices around board games and virtual spaces like the one presented in *The fruitmachine*. Consider this stretch of talk between pre-school children and their nursery nurse over a game of 'Winnie the Pooh':

1	ChildA:	that means he's in [front of me]
2	Nurse:	[he's in front] of *you* see if you can beat him now (.) go on then, how many have you got?
3	ChildA:	four
4	Nurse:	four.
5	ChildA:	one two three: four
6	Nurse:	no: they're there where Ben is. Right James your turn.
7	ChildB:	so I chase the butterfly as well
8	Nurse:	you must have been chasing Ben I think. How many have you got?
9	ChildB:	I (wanted) to be chasing Ben

These children can be seen to use referential expressions that place them 'inside' the space of the game – a kind of grammar of playing board games. In the course of the game, they are further instructed. The way in which action is projected to the 'there' of the space of the board game does resemble the way in which this is done in the above example. However, further investigation is required.

extends one's point of action into the electronic space. Once it has been determined who is who in the electronic space, the players' responsibilities stretch from the interface to the pieces on the screen. What they are accountable for are the events on the screen. How they trigger those through the interface usually falls outside the field of attention of the other players.

In the second example, the players also employ a referential grammar reminiscent of board games: 'wait I'm coming, wait' and 'I have bananas, too', referring to the images of fruit on the sides of the pieces. In both examples people thus anchor their actions in the electronic space to that particular piece by 'identifying' with it. This strategy of placing oneself 'inside' the electronic space could also be observed in the 'Legible City', where people said 'I am on the market square' or 'I'm in Manhatten'. However, rather than experiencing presence in the sense of disembodied immersion in a new element, where ''the body', as a sense apparatus, is nothing more than excess baggage for the cyberspace traveller' (Balsamo 1995 p.229, see also Heim 1993, 1995), we would argue that people begin to *inhabit* these spaces through active involvement with the objects and people they find there. Objects that a person controlled, avatars, or simply one's position in the electronic space (determined by one's point of view) were referred to as representations of one's point of action in this place, at this moment in time. They were 'a way of making oneself present <u>there</u> as well as <u>here.</u> And, in the process, the indexical references to <u>there</u> became <u>here'</u> (Tolmie 1997). This ties people's responsibilities and accountabilities into the electronic environment and becomes a means of achieving intersubjective availability of one's position, orientation, and focus.

Interaction with objects and spaces

The examples above show that finding answers to such questions as 'where am I?' and 'where are (what kinds of) others?' is of crucial importance to interactions in electronic environments. Equally important are clues about others' and one's own activities in relation to the 'material' environment. Visual clues alone do not seem to be sufficient for a sense of presence. In the above example of large scale spatial information woven into the interactions between different players in Labyrinthos, for example, T's question in lines 2-3 indicates that despite having access to both screens, he is unable to determine the two players' positions in relation to each other and possible targets. This is puzzling. T seems to be less able to transpose everyday competencies than the other boys. One of the reasons for this may lie in the level of his engagement with the game. He is not a player, but an advisor, monitoring the state of affairs rather than being part of them. The next examples look in detail at the question of how people find their way around in the 'material' arrangements of electronic spaces.

Docking manoeuvres: The Fruit Machine

Two men (A,C) and a woman (B) (all about 20) have just put on the 3D glasses and sat down at the spacemouse controls of 'The Fruitmachine' (figure 6):

1	B:	ºeh I can't see anythingº ((said quietly, in a laughing voice))
2	A:	which one are you now
3	B:	I think the bottom one (.) yeah
4		(0.3)
5 →		do they fit into each other somehow? (yeah or this way) na:: not quite ((laughs))
6		((A&C laugh))
7	B:	those [fit]
8 →	C:	[(come] further to the front)
9 →	B:	they probably don't fit together at all (.) what'd'y mean further to the front?
10	C:	you can't you turn em can you?
11		(0.3)
12		ahh
13 →	A:	ooh now there's one really (.)
14 →	C:	ah ja
15	B:	[hey look]
16 →	C:	[heyheyheyhey]
17	A:	we're nearly there
18	B:	but that doesn't [that doesn't work for some reason]
19:	C:	[little closer little closer a:nd]

Two very different levels of engagement with the installation are documented in the talk and the activities of A,B, and C. While A and C discover the task posed by the installation through moving the pieces and take it at face value, B is more detached and critical. Here, and throughout the whole duration of their engagement with 'The Fruitmachine', she doubts the fit of the pieces. At one point she reveals one of her reasons: 'I think they're just taking the piss y'know'. She insinuates that the artist elicits a response like A and C's under false pretence, that the 'message' of the art work must be something more intricate than what can be perceived at first glance. As a result, she holds back and only half-heartedly involves herself in the task. This in turn, closes off some of the qualities of the electronic space that are discovered tacitly, through 'doing' by the others (lines 8-9). She does not perceive the space as a space and only much later concludes that it must be three dimensional: 'they [the pieces] are always in different planes somehow'. A and C, on the other hand, focus on the events on the screen. Their utterances have the character of a running commentary rather than being reflections on the nature of the game. This becomes clear inter alia through the intonation of their talk which is linked to the events on the screen. These different levels of engagement with the installation result in visibly different levels of dexterity in, and understanding of, the electronic space. A and C's grasp of the affordances of the environment is corporeal.

Increasingly, the installation – the controls and the events on the screen - become 'ready-to-hand' for them. 'Readiness to-hand' describes the tacitly and bodily known way in which tools, objects, and spaces melt into their situated use. It is different from 'presence-at-hand', a state into which these entities can switch whenever we look at them with more reflective eyes (Heidegger 1962, p. H 69-77). Focusing on 'docking' their piece onto another on the screen rather than the interface, A and C project their manual motility into the electronic space. This orientation facilitates the acquisition of a 'ready-to-hand-ness' of the space and the objects in it. People could be seen to get habituated to the thus augmented reach of their body in a way that is similar to the way we get used to the external dimensions of a new car (see also Merleau Ponty 1962, p. 143). The ways in which the interactivity of this work of art and others like it is known and experienced bodily as well as rationally suggests, again, that ‘presence’ is something that ties the real and the virtual together rather than taking people ‘into’ a virtual that is discontinuous with the real. The engagement with the electronic spaces of the works and a sense of presence in relation to them is not a case of disembodied immersion, it is one of representation, remote action, or instrumentally supported action that is rooted in the real world.

In our everyday life we handle a myriad of objects and move within spaces, mostly without being consciously aware of the factors that are involved in their being ready-to-hand. If we take a pen, for example, we touch it, suspended it between our fingers and lift it as we move our hand. If we put it down again somewhere it will stay there unless we (or somebody or something else) move it again. These activities are ordered through our perception. But how? And how do we deal with the impoverished or altered sensory affordances of electronic environments? One fairly obvious factor people could be seen to rely upon was the sequence of action and events in the environment. In the Fruitmachine, one of the pieces moves of its own accord unless it is controlled via the interface. This means, for example, that this piece cannot be placed somewhere and held static in order for the others to ‘dock’ onto it. It has to be the last piece to be ‘docked’. This also revealed that people relied on the sequence of ‘action on the controls → movement on the screen’, because the self-generated movement caused puzzlement, as one visitor states: ‘What I noticed was for example at this 3D puzzle that from one terminal it was running out of control maybe that was intended, I don’t know in order to make it more difficult’ (interview 13.12.97). Other factors are: the continuity of material arrangements – do things stay where they are and what they are?, the continuity of one’s own actions – do actions have lasting effects or are they reversible or only temporary?, and the consistency of those actions – do the same actions cause the same events every time?. But there are also more interpretative tacit means intertwined with an embodied way of knowing the world that point to some other factors involved in the perception of the electronic spaces and objects in the exhibition.

Material markers

In The Legible City a single visitor can cycle through representations of three cities (Amsterdam, Manhattan, and Karlsruhe). Each city is based on the street plan of the respective real city, while the buildings are made up of letters and words. An LCD display of a street map with a blinking dot indicating the user’s position is mounted onto the real bicycle. The following excerpts from the field notes illustrate how a succession of people moved around the Legible City:

> A is holding a leaflet and cycling. His friends B, C stood in the back, laugh when he can’t stay on the street. A: ‘where am I?’. Checks on plan. Two women behind him look over his shoulder onto display. B,C leave. A evades letters, gets off, leaves. Two women take over, one cycling, one checking with her where they are on the display. A couple come in, stand back at the wall. Then a woman (D) and boy (E) enter. She explains what can be done in this installation to the boy,

> referring to what a friend told her yesterday. A couple and a boy. The boy wants to try, but his father (F) gets to the bike first, gets on and cycles really fast, switches to Manhattan, cycling. D comments: 'there's going to be a storm' ((referring to the gloomy atmosphere over Manhattan)). E: 'it's not raining' D: 'mhm seems to just stay the same'. F cycles up closely to letters, then goes through. …. A boy gets on, goes through letters carefully, then backwards, then goes through the letters again. … a girl, audience laugh when she goes through letters. When she comes to a row of red letters, she can't get through. Turns round, goes through blue letters opposite, then tries the red ones again. Her parents come and look over her shoulder. A small child and father. Father doing the pedalling: 'No not there it doesn't go anywhere, don't go through letters, stay on the road … this is Karlsruhe now, lets' go to the castle, past the tax authorities, that's where the castle is. No, it's not there'

There are a number of issues relevant to the experience of presence that can be drawn out from this extract. Firstly, here, too, people were concerned with finding answers to the question 'where am I?', which in the course of movement through the city turned into 'where have I been and/or where am I going?'. Material markers were practically employed as orientational aids. They range from urban infrastructural elements such as roads, junctions, buildings, landmarks, to natural characteristics such as the sea, or the climate and weather, to physical laws such as the solidity of materials. In our everyday interaction with the material world, such material features are mainly understood and interpreted tacitly. Moreover, they are encountered and related to our own actions as part and parcel of a totality. Some of the features in the Legible City give rise to surprise or puzzlement and thus draw to our attention, not only what kinds of material features form part of our tacit interpretation of the physical world around us, but also point us in the direction of how such features are employed differently within electronic environments.

The combination of the bicycle as an interface and the roads and buildings on the screen that adjust to the moves of the bicycle in real-time offers the user the opportunity to travel through the cities. People readily recognise this as the function of the installation. Whether people recognised the cities as based on the ground plans of real cities or not, the urban metaphor allowed them to see at a glance what the rudimentary workings of this piece of art are. Most carried an urban interpretation further by initially trying to stay on the roads in the environment. This is an interpretation sanctioned or even enforced by the audience. The father instructing his daughter not to go through the letters but to stay on the road, makes public just how strong the thrust of such an interpretation is. Similarly, the laughter that greets the difficulties people have in staying on the roads or the first swerves into the letter-buildings indicates that there is something 'wrong' in deviating from these paths 'meant' for traffic.

A similar approach to navigation in an urban electronic environment was found in the first trials of DEVA, where people were assigned the task of finding a particular road. (see Deliverable 2.1) However, in the Legible City most people inevitably overshoot junctions and pass through a row of letters sooner or later, as the speed facilitated by the bicycle is quite high and the corners narrow. Thus the lack of collision control within the Legible City is discovered as the result of ‘accidents’ (again see the DEVA trials reported in Deliverable 3.1). It proved a source of enjoyment: after having crossed through letters once, people (especially children) visibly and audibly enjoyed cutting through the letters. This was pleasurable for its own sake, as the fact that people circled around and around through the letters laughing showed. But it also had practical consequences. It meant that the roads were no longer the only channels for traffic, and that different types of journeys through the cities were possible. While slow movement allowed people to read the text that stretched alongside the roads, they could also decide to ‘just go’ or orient with the help of the street plan and steer towards their destination regardless of the ‘buildings’. The degree to which the immaterial character of the letters is taken for granted once it has been established is illustrated through the fact that when a row of letters is discovered that is impenetrable, people react with disbelief. Several trials are made to establish that *this* row (which is the edge of the model) is not like the others.

Other features of urban environments that people used were landmarks. This is most pronounced in situations where visitors who are familiar with Karlsruhe choose this city in the Legible City. The castle, as the main architectural focal point of the city is sought out and located with the help of other landmarks in the above example. Other, more sculptural landmarks in Karlsruhe, which are set (in the real and the simulated city) in the middle of major lines of sight helped people locate themselves within the whole of the environment, and were also used as short term navigational aids: on their way to a point further down a road that was adorned by such sculptures, people cycled through them to stay right in the middle of the road. While the model of Karlsruhe had ‘real’ landmarks in the sense of such recognisable buildings or features that are visible from a distance, the other two cities had none. But people still ‘saw’ them:

> ‘I think it is Manhatten'. The boy gets on and cycles very quickly. Manages to read nevertheless: 'Central Park'. Father: 'Central Park, exactly.' Boy: 'hey sex.' Father: 'oh I knew it, you would spot that wouldn’t you? ... that's the sea. Go a bit further to the Empire State building (.) should be there somewhere, further further, it's up here ((points to map)) further further and back now. What does it say there? They stop and try to read. Boy pans left along the row of letters. Stops. Father: 'Ten?' Boy moves backwards until letters turn into rectangles, edges forward until they turn into letters again. Stops, then turns right and cycles forward. Father: 'don't go in the sea'. Boy adjust course to the left. Father turns handlebars and steers back towards letters: 'now you're looking back onto (xx)'.

In this example, the city is Amsterdam, but mistaken to be Manhattan. The fact that the writing is in Dutch is no obstacle to this, because the father and his son pick out only those words that are English: ‘Central Park’ and ‘Ten’. Since Dutch is just as much a foreign language to them as is English, the rest of the letters are taken to be unknown words. Equally, the fact that there is a waterfront in both cities supports the interpretation, while the typical Amsterdam ‘canals’ are ignored. Although the landmark is not actually 'seen' on the screen, it is assumed to be there and the journey is organised around its existence.

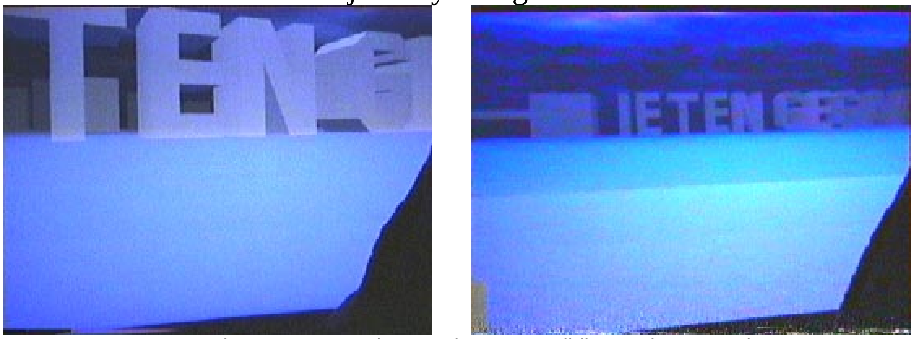

Figure 7: Next to the 'Empire State Building' … in Amsterdam

Weather or climatic conditions were another resource for orientation. In the first example from the field notes above, a woman states 'there's going to be a storm', as a cyclist is proceeding along a broad road towards a vista of dark clouds on the horizon. However, the fact that 'the weather' does not worsen even though they should be getting closer and closer to the centre of the storm, exposes it as governed by different physical laws than weather in our familiar physical environment. In this case the everyday practice of reference to large scale orientation clues in the sky fails, because the way the storm follows the cyclist makes it unusable as such a resource.

These examples illustrate that material markers are interpreted with reference to the physical 'world known in common' but submitted to review and learnt about' in light of the experience of moving around in the environment. In the final section we will provide an account of presence based on this and the previous observations and suggest some implications for the design of large scale CVEs.

Intersubjectivity and Learnability

The observations outlined above support a phenomenological explanation of the experience of presence. Recently, attempts at defining presence have turned to phenomenological and ecological accounts of human experience and perception (see, for example, Zahorik and Jenison 1998, Heim 1993). At the heart of this approach lie the following observations:

- Many aspects of both our social interactions and our interactions with the material world are known tacitly and/or bodily. This is a result of practice and lived experience. The situated nature of practice and experience allows tools, objects, and spaces to be 'ready-to-hand' (e.g. Heidegger 1962, Merleau-Ponty 1962).
- The affordances of the material world (environments, objects, tools, etc.) are the physical and interpretative constraints and possibilities we find. They, too, are situated – tied to the here and now of activity and practice (Gibson 1986).

- We do not simply ‘soak up’ stimuli from the material world through our sensory apparatus, process them, and put together a representation of the situation at hand and then act. Rather, perception and interpretation are one activity (Gibson 1986).

We offer a comprehensive account of this body of phenomenological work in Appendix One of this Delivearble, and it is in reference to this work that Jenison and Zahorik (1998) develop arguments that

> ‘Presence is tantamount to successfully supported action’, a state which can be said to have been achieved ‘when the environmental response is perceived as lawful, that is, *commensurate with the response that would be made by the real-world environment* in which our perceptual system has evolved.’ (p. 87, emphasis added).

Heim (1993, 1996) puts forward similar arguments, asking that we strive towards an ever richer and more realistic sensory interface to the electronic environment (e.g. CAVE). Equally, enriched sensory ‘input’ and realism of ‘physical laws’ and appearance also underpins much of the recommendations arising from experimental studies that have identified a number of factors that influence a sense of presence in electronic environments (Slater 1998, see also Benedikt 1992).

While not disputing these approaches we would like to offer a slightly different account of presence that is less directly connected with some sense of fidelity of external sensory inputs and visual appearance. The ethnographic observations indicate that a sense of presence is achieved through the engagement with resources provided within the environment. Perception and interpretation thus indeed are one activity. This activity is tied to ‘grammars’ of human interaction and sense making (Wittgenstein 1953). It is part and parcel of us being, continuously and inescapably, a member of a social world. It is itself social. Like any other knowledge, practice, or habit, perception is intersubjective. Whoever made the material arrangements and objects we encounter is/was part of the ‘world known in common’ (Schutz 1970) and likely to have imbued the object with features known to us all. ‘Known to all’, however, does not imply that we all know the same things. On the contrary, ‘an essential ingredient of “the world known in common” are the practical methods we have for “finding out” or adding to our “stocks of knowledge”’ (see our discussions in Chapter One). Some of these practical methods are quite explicit – in the context of art exhibitions we learn, for example, through asking or watching others, or through reading additional information in leaflets, catalogues or the like. Other ways of finding out are more tacit or embodied.

Put crudely, the more we find out, the ‘better’ we are able to interact both with respect to people and with a view to the material world. But the ways in which people do find out through the transposition of everyday practices implies that they can interpret a diversity of environments. As long as a sufficient degree of intersubjectivity is facilitated, we are flexible with regard to the visual appearance and ‘material’ structure of an environment, its physical laws, and the affordance it provides with a view to our own presence and that of other people.

Clues about what resources would be desirable and how intersubjectivity can be facilitated arise from the fieldwork. Some of the main issues are the following:

Reciprocity of perspective

In order to be able to coordinate actions, people must be able to assume that others see the same space and the same objects as they do. But this does not mean that the principle must not be relaxed under any circumstances. It can be relaxed - as long as this is made available as a feature of the environment -, or provided for in ways that are unfamiliar, such as making people's positions and activities available through, for example:

- an abstracted, animated, real-time plan view.
- a screenshot of 'this is what I see from where I am'.
- a vehicle into which others can be invited to share a particular perspective onto a joint space (see Murray 1998).

Designing for two worlds

There are two different ways in which the real and the electronic spaces are linked. Firstly, the practice of seeing and experiencing the electronic space *with* or *through* the fusion of interface and display into an instrument that becomes 'ready-to-hand' does mark out a route for the development of more sophisticated equipment. However, increased sensory richness (through for example, force feedback and proprioceptive devices) is not necessarily the only direction such development could take. Although the spaces seen are not 'real' spaces, the sense of presence in electronic environments is a version of 'telepresence' (Sheridan 1992, Benford 1996) in that it provides a sense of control over remote events. The sequence, continuity, and consistency of actions and the events they cause, if stable and intelligible, seem to provide 'enough' information for a rudimentary sense of presence. They do not have to mimic real-world forms of sensory perception.

Secondly, the context of interaction in the physical space is relevant to the activities in the electronic environment and it should be possible to make available at least some crucial information about it. At the same time as placing some demands on the designer, this fact can also be an inspiration. Chapter Two developed an account of the interactive art installations that draws on the metaphor of the theatre. Some of the observations outlined above suggest that the experience of the works is, indeed, a collective one reminiscent of some forms of theatrical performances. In the Legible City, for example, people hardly ever encountered a blank slate. Most people entered while someone else was cycling, watched, or even got engaged in the exploration of the cities, and then tried for themselves. The cyclist becomes a 'performer' and the members of the audience are often drawn into the performance as active participants. As a way of finding out about an electronic environment this practice opens up an opportunity to employ animated instructions or 'Help' functions that people can view.

Material markers

Some aspects of 'grammars' of perception in public spaces are addressed in a variety of empirical studies of architectural features in real spaces (e.g. Lynch 1960, Whyte 1980), and theoretical considerations in this area (Alexander 1977, Arnheim 1996, Hillier and Hanson 1984, Hillier, 1996, Tufte 1997). Lynch, for example, shows that cities are structured in terms of nodes, paths, districts, edges, and landmarks in people's perceptions. There have been some attempts at structuring visualisations of databases (Benford *) drawing on Lynch's work. However, one of the problems that arises from introducing such structural means into the visualisation is their potential to obstruct and clutter the view. Starting from the ethnographic work, a re-consideration of this body of work could be fruitful. There also seems to be some promise in combining this with a study of cinematography: some of the practices observed in people's approach to navigation through the electronic environments show parallels to ways of exploring the space of a movie-set (e.g. the way people 'pan' and 'zoom in' on the words in Legible City).

More generally, the observations presented in this chapter suggest that in designing and developing virtual environments we need to consider how factors that influence a sense of presence are made available to people using these environments. Uncovering and presenting these features of these environments allows us to make resources available to users that they can in turn exploit to develop a sense of presence across a wide range of different presentation techniques and devices. In addition to developing different mechanisms for presenting these properties to users we need to continue an investigation of how people make use of different properties of virtual environment to construct and generate a sense of presence in these environments.

References

Alexander, C. 1977. A Pattern Language. Towns, Buildings, Construction. New York: Oxford University Press.

Arnheim, R. 1988. Die Macht der Mitte. Eine Kompositionslehre für die bildenden Künste. Köln: Dumont.

Balsamo, A. 1995. Forms of Technological Embodiment: reading the Body in Contemporary Culture. In: Featherstone, M. (ed.), *Cyberspace, Cyberbodies, Cyberpunk.* London: Sage, pp. 215-237.

Benedikt, M. 1992. Cyberspace. Cambridge, MA: MIT Press.

Benford, S. et al 1996. Shared Spaces: Transportation, Artificiality, and Spatiality. In: Proceedings CSCW 96. Cambridge MA: ACM Press, 77-96.

Benford, S. et al. * Building Virtual Cities: applying urban planning principles to the design of virtual environments.

Bowers, J. et al. 1996. Talk and Embodiment in Collaborative Virtual Environments, in Proc. CHI'96 (Vancouver, Canada, April 13-18,1996), ACM Press, 380-389.

Ciolek, T.M. and A. Kendon 1980. Environment and the Spatial Arrangement in Conversational Encounters. In: sociological Inquiry Nos 3-4: 50.

Gaver W. 1992. The Affordances of Media Spaces for Collaboration, in *Proc. CSCW'92* (Toronto, Canada, Oct.31-Nov.4, 1992), ACM Press, pp. 17-24.

Gibson, J.J. 1986. *The Ecological Approach to Visual Perception.* London: Lawrence erlbaum Associates.

Goffman, E., 1961. Encounters: two studies in the sociology of interactions. Indianapolis: Bobbs-Merrill

Heath, C. 1986. *Body Movement and Speech in Medical Interaction.* Cambridge, UK: Cambridge University Press.

Heath, C., and P. Luff 1991. Disembodied Conduct: Communication through Video in a Multimedia Office Environment. In: *Proceedings of CHI'91*, New Orleans, Louisiana, April 28-May2 , 1991. New York: ACM Press, pp. 99-103.

Heidegger, M. 1992 (1962). *Being and Time.* Oxford: Blackwell.

Heim, M. 1993. *The Metaphysics of Virtual Reality*. Oxford: Oxford University Press.

Heim, M. 1995. The Design of Virtual Reality. In: Featherstone, M. (ed.), *Cyberspace, Cyberbodies, Cyberpunk.* London: Sage, pp. 65-95.

Hillier, B. 1996. *Space is the machine. A configurational theory of architecture*. Cambridge: Cambridge University Press.

Hillier, B. and J. Hanson 1984. *The Social Logic of Space.* Canbridge, UK: Cambridge University Press.

Kendon, A. 1990. *Conducting Interaction. Patterns of Behaviour in Focused Encounters.* Cambridge: Cambridge University Press.

Lynch, K. 1975 (1960). *The Image of the City*. Cambridge, MA: MIT Press.

Merleau-Ponty, M. 1996 (1962). *Phenomenology of Perception.* London: Routledge.

Murray, C. 1998. The Cityscape: Theory and Empirical Work. eSCAPE Working Paper.

Robertson, T. 1997. Cooperative Work And Lived Cognition: A Taxonomy Of Embodied Actions. In: Hughes, J.A., W.

Sacks, H. 1992. On Exchanging Glances. In: *Lectures on Conversation* Vol.1, pp. 81-94

Schutz, A. 1962. *On Phenomenology and Social Relations*. University of Chicago Press.

Simmel, G. 1970 (1908). Sociology of the Senses. In: Park, R.E. and E.W. Burgess (eds.) *Introduction to the Science of Sociology*. Chicago : University of Chicago Press, pp. 146-151.

Slater, M., Steed, A., McCarthy, J, Maringelli, F. 1998. The Influence of Body Movement on Presence in Virtual Environments. Forthcoming: *Human Factors*. http://www.cs.ucl.ac.uk/staff/m.slater/publications.html

Sudnow, D. 1972. Temporal Parameters of Interpersonal Observation. In: Sudnow, D. Studies in Social Interaction. New York: The Free Press.

Tolmie, P. 1998. *Representation* - Working Document for CSCW Workshop. Lancaster University.

Tufte, E.R. 1997. Visual Explanations. Images and Quantities, Evidence and Narrative. Cheshire, Conneticut: Graphics Press.

Whyte, W.H. 1980. The Social Life of Small Urban Spaces. Washington DC: The Conservation Foundation.

Wittgenstein, L. 1997 (1953). Philosophical Investigations. Oxford: Blackwell.

Zahorik, P., Jenison, R.L. 1998. Presence as Being-in-the-World. In: Presence Vol. 7 (1), pp. 78-89.

Chapter Five: Developing a framework for e-scapes

Jonathan Trevor, Monika Bűscher, Jon O'Brien, John Hughes, Tom Rodden, John Mariani

Lancaster University

Introduction

This chapter builds upon the fieldwork undertaken at ZKM in the first year of the eSCAPE project and reported in Chapter Four of this Deliverable. It presents the development of a set of techniques to allow users to understand the properties of virtual environments as they move between different environments. Such an exploration is clearly of importance in the eSCAPE context, with its focus on developing an environment within which a range of heterogeneous environments meet, a developmental agenda that requires us to consider carefully the ways in which participants in e-scapes might move within and between constitutent worlds. We do so in this chapter through a brief discussion of a range of shared cooperative systems and collaborative virtual environments, considered in the context of the spatial properties of these environments and the ways in which an understandings of these properties might be generated to inform the design of large scale environments accessible to the general citizen.

The intention of this approach is, then, to move beyond the 'single world' focus of much previous work, and we therefore re-visit the ZKM fieldwork in order to develop an understanding of the orientation of those observed towards space and spatiality in both the 'real world' of the multimedia museum and the range of electronic spaces offered by installations within the museum.

Spatial Approaches to Cooperative Systems

The development of many shared cooperative systems has exploited a spatial approach. This has ranged from the presentation of shared spaces using text in MUDs and MOOs (Curtis et al. 1994) through 2D spatial systems such as Worlds (Mansfield et al. 1997) and TeamRooms (Roseman et al. 1996) to the development of 3D virtual environments (Benford et al. 1994) The emergence of these shared spaces has raised some debate as to the nature of spatial metaphors and the development of different forms of shared environment (Harrison & Dourish 1996).

Cooperative virtual environments (CVEs) have now taken a central place in CSCW and we have seen extensive developments of applications of these environments (Benford et al. 1996) and studies of their use (Bowers et al. 1996). However, despite the growing use of these systems within CSCW there has been little consideration of the general properties of these systems. Rather users and developers have to rely on their own experiences of space when they seek to understand these environments. Little understanding has emerged as to the common properties or features of these environments.

One reason for the lack of any common properties has been the exploratory nature of the development of virtual environments. Essentially, the design of collaborative virtual environments has been driven by an exploration of research concepts. Notable examples of these concepts include a range of ‘models’ of awareness and interaction (Benford et al. 1995) and the affordances of user embodiment [0]. These research efforts are clearly driven by the need to build virtual environments that are inhabitable and shareable by multiple users. As you would expect these environments offer a potentially confusing diversity of possibilities. For a city like facsimile to an abstract data space (Benford & Mariani 1994)

As the concept of cooperative virtual environments moves beyond these early prototypes we need to explore how they may be more generally understood. This is particularly important as these environments become accessible to a diverse community of users through public networks. Given this increased access it is likely that future users will make simultaneous use of a number of virtual environments.

In short, we wish to consider the means by which virtual environments might provide large scale shared virtual worlds that are accessible to a wide range of users. To do so we need to consider the properties of these environments and how we may allow a number of these environments to be simultaneously available to users to promote their use.

In this chapter, then, we seek to understand the development of these environments from this multiple world perspective by uncovering common properties of these environments and practical techniques to support movement between these worlds. This requires us to move beyond the current isolated world perspective of these environments.

Moving beyond a single world focus

Current examples of virtual worlds (Alphaworlds, Benford & Greenhalgh 1997) consider their structure, design and development *from within*. The inherent assumption is that at any moment in time users of a virtual environment exist solely within it and they are immersed in only that world. Little or no consideration is given to the situation where users need to simultaneously inhabit a number of environments and reason between them. The support that currently exists for multiple worlds concentrates on connection and is limited to some form of portal or gateway. These allow users to move from one world to the next by "loading" a new world to completely replace the existing one. Few facilities are provided to allow users to understand the nature of the environment they are about to enter from outside the environment.

One result of this isolated focus is that virtual environments do not to allow users to reason across worlds or to present aggregated views drawn from these worlds. Essentially, these systems consider the world as existing within a single system and presented to users of this environment from this system (Carlsson & Hagstrand 1993) Thus research has considered the structure of objects within these environments and the development of appropriate avatars (Benford et al. 1994) Little consideration has been given to the overall context within which these environments are placed and the need for these environments to co-exist with other systems. However, our previous studies of environments have shown the need to allow users to manage their relationship with the environment in conjunction with other environments (Bowers et al. 1996).

One consequence of this isolated consideration is the current diversity of approaches to developing these environments. Standards such as VRML (ISO/IEC 14772-1:1997) focus on the structure and presentation of virtual environments rather than any higher level properties. Consequently considerable variability exists in the nature of these environments and the form of interaction they support. As multiple users navigate and interact between these worlds they need to understand the nature of these environments and how they interact with them.

To this end we wish to uncover an appropriate set of general properties for these worlds that can be articulated to users as a means of making them learnable. Rather than develop these properties in principle we seek to uncover them from a more practical examination of the nature of these environments. In our previous work (Benford et al. 1997) we have introduced and discussed the notion of informing the design of a single virtual environment through an ethnographic study. We now need to consider the development of techniques to allow a number of these environments to coexist and how this can be informed. This chapter considers the development of a set of properties of virtual environments and techniques to present these drawn from an ethnographic study of the use of multiple virtual worlds.

Moving from Theory to Practice

The last few years has seen a rapid growth in the accessibility of virtual environments. In the early days of their inception, virtual environments were scarce resources requiring investment in expensive machines and specialist equipment, often only accessible to members of research labs. However, this is no longer the case for virtual environments. These now tend to be targeted toward more accessible machines and are often publicly accessible across the Internet.

This shift requires us to reconsider the nature of CVEs and to take them seriously as design resources and technologies that can be used to support the work of groups. One consequence of this shift is that we need to move from our existing theoretical considerations of these environments to an understanding of these environments in practice.

The design of any particular cooperative virtual environment to support particular activities needs to be of practical utility to those undertaken the activities at hand. This demands an investment in gaining an understanding of the nature of those activities and the relationship between activities manifest in the real world and those within the virtual environment. One approach to the development of specific environments is to inform the particular Collaborative Virtual *Environments* (as well as user embodiments) through the undertaking of ethnographic studies of activities within the real world (Benford et al 1997). However, this approach will only address the needs and demands of that particular world. In this chapter we wish to consider the situation where users need to make use of more than one virtual environment.

In particular, we wish to consider people's movement between different environments. It is important to facilitate such movement in a way that allows people to find their way around both the larger 'universe' and the particular 'worlds'. Movement along these lines has already started to emerge with on-line environments and we envisage this use of electronic spaces to grow over the coming years. To achieve a more intuitive and legible design of such large scale electronic landscapes, ethnographic studies of people's strategies employed in moving around real world spaces can be used to inform the design. These need to take into account the social organisation of space with a view to both its interactional affordances and the legibility of 'material' arrangements in space.

As we stated in Chapter Four, observations in real world situations alone do not neccesarily suffice to inform the design of CVE's that fulfil these requirements. People adapt their practices to the affordances of the electronic environment. Everyday practices of orientation, movement, and interaction in space will be drawn upon, but transposed rather than transplanted in order to fit in with the affordances of the environment – this was seen as a key element of the rationale for the observation of people's interactions in and with such environments.

In essence we seek to draw parallels between the visitors to the gallery moving from one installation to the next and on-line users of an electronic landscape travelling from one virtual 'world' to another. The observations allow us to gain insight into practices employed by a heterogeneous group of people, in the context of a diversity of real world electronic environments. Moreover, while on the one hand we study the exhibition as a proxy for a large scale 'container' of different smaller environments, we also learn how aesthetic features impinge on people's ability to find their way around the whole and its parts.

Connecting Environments - Learning from ZKM

The centre for art and media technology (ZKM) in Karlsruhe was conceived as a new type of institution that brings together art and technology in an unprecedented way. It combines two research and development institutes - the Institute for Music and Acoustics, and the Institute for Visual Media -, a Media Library, and three museums - the Museum for Contemporary Art, the City Gallery, and the Media Museum. The Media Museum houses exhibitions of interactive multimedia art by local and international artists. There are around 30 works on show at present. In very different ways, these works explore the properties of multimedia environments. The nature and affordances of cyberspace, the relationship between the interface and the interactivity of the installations, and the role of electronic technologies in our society are topics that are examined at a conceptual, but also at an experiential level. Together with the fact that many installations invite visitors to interact in and with electronic spaces, this makes it an ideal candidate to study people's reactions to innovative ways of transposing familiar features, and aesthetic principles of spatial arrangements into electronic environments.

The people who visit the Media Museum form a heterogeneous audience. Although there have been no systematic studies, the internal monitoring of ticket sales allows some insight into the composition of this audience. Many visitors have an interest in art in general, as the fact that 15% of visitors buy a ticket that allows them to see all three museums shows. Almost a third of visitors are eligible for a reduction, because they are under 18 or in full time education. Interest expressed in the tours offered by the department for Museum Pedagogy show that families, schools, colleges, and businesses are particularly attracted by the Media Museum. Our own observations and those of museum staff indicate that there is a heterogeneous mix of people of all ages, of local people and people travelling some distance to visit the museum, and of people with some previous knowledge of multimedia technology and those who have little experience with such technologies.

Worlds on a string

The majority of the works are interactive multimedia art installations that project images onto a screen. In order to ensure the quality of this projection many of them are surrounded by a 'room'. Unlike the visitor to a traditional art gallery who is face-to-face with the exhibits while being within a public space, the visitor to the media art museum thus enters a succession of small, dark enclosures in order to see the works of art. The characteristics of this kind of flow of visitors are addressed at the level of the exhibition design, as one of the exhibition designers points out:

> If some kind of interior space was needed (for the sake of projection, for instance) then we were keen to give its exterior that mediamatic quality that would allow us to develop the

> Media Museum as a string of events within the architectonic space of the existing building. (29.1.98, FdO)

Through the design of the outer 'shells' of the rooms, some first clues about what expects people on the inside are to be conveyed. Visitors use these clues as a resource in their movement through the exhibition. Most effective are designs that allow people to get a glimpse of the installation from the outside, as the following extract from the field notes illustrates:

> Someone is leaning into 'Beyond Pages' through a window in the wall. A couple strolling up from the left turn to face the window, stop to peer in. He looks through the window over the other person's shoulders, she glances in the direction of the entrance and sees people coming out. She turns and walks towards the entrance, followed by her partner. The man who remained at the window is joined by his girlfriend. They watch, both leaning on the windowsill. About 30 seconds later, the couple return, They all watch the activities inside the installation through the window. Then the couple move on towards the next installation.

The couple's movement into the 'informational radius' of this installation allows them to gauge some initial information about it. It begins with the fact that someone else's curiosity has been sufficiently caught to make him stay and watch for a while. Over his shoulders the couple see a group of people gathered behind and around a table where one person is interacting with a virtual book. A window, a gap in the wall, or other structural arrangements allow people to get not only a glimpse of the content of an installation, but also of its popularity, and the nature of the experience it provides. Moreover, such permeable structures afford the visitor an at a glance availability of the 'queue' inside the installation. The queuing system that regulates access to the installations is displayed to the passer-by through the position and orientation of people in and around an installation. Individually or in groups, visitors can engage with the works of art at the level of being in control, they can watch the events in the installation as the 'next in line', they can be spectators, or they can be 'floaters' - peeping in in order to decide whether they want to stay, return, or skip this installation on their tour around the museum. This queuing system and the flow of people through the exhibition space as a whole, are part of the information visitors routinely monitor and weave into their own decision over where to go and what to do. Such at a glance visibility of other people's activities is an important resource for people's orientation in an environment that requires them to choose between different places in a string of events.

Stepping inside

Clues inherent in the architectural appearance of the installations were another resource employed. Artists attempted to make their work accessible through incorporating such clues into both the physical structures of the installation and the 'material' arrangements in the electronic spaces. Visitors could be seen to use them as orientational aids. The 'Memory VR Theatre', for example, explains its organisation through its structure. It is surrounded by a circular wooden wall, into which one entrance is cut. Stepping inside one faces a large screen and a circular plexi-glass model of the room. Between the model and the screen a pedestal is placed. The model of the room replicates this interior structure. Attached to it by a wire is a Polyhemus 3D tracker enclosed in a plastic shell. This and the model of the room constitute the interface that allows people to interactively explore the electronic spaces provided. Holding the controlling device above the plexi-glass model of the room, one sees a clock overlaid with a compass. Through diving into one of the four geographic sections, one enters one of four different electronic spaces. Each of these spaces mirrors the organisation of the physical space of the installation. There are thus six spaces that are arranged according to the same architectural principles: a circular room with a pedestal positioned in front of a screen. The logical and visible structure of the installation was, however, not easily understood by visitors to the exhibition. Rather, people frequently experienced confusion. Here, two people have just entered the installation:

> The woman picks up the control from the floor of the container and moves it upwards. On the screen, she and her friend see the interior of one of the electronic rooms, then the black and grey model of the ceiling. Briefly they also get a glimpse of the compass/clock threshold to the electronic rooms. When she moves downwards again, they enter a different room. Her friend takes the control and 'scans' his watch and his hand. The screen turns black. He then moves the control device back in, but only to the rim of the container. He moves it along. The screen displays a wobbly journey along the edge of the modelled room. She takes his hand and guides it further down, where they re-enter the room they were in previously. On one of its walls a view of a staircase is displayed. He slides the control up the stairs twice. A group of six young people enter and watch. The woman and her friend place the control on the floor of the container and leave.

Confusion, or bewilderment is visibly experienced with regard to what one can do and how the events on the screen are related to one's actions. It is clear that the control device has a camera-like function, but what and where it can 'see' is less clear. At the same time it is a navigational device, but, again, how one can move and where one can go is not immediately obvious. The clues inscribed into the overall structure of the installation as a whole are not intuitively drawn upon for purposes of orientation. Instead, people react to more generic indices of function or organisation expressed through form. In this example the stairs are tested as a means of movement through the different spaces, in other examples, doors, paths, buttons, handles, etc. were interpreted as having similar functions as their counterparts in the physical space. One of the reasons for this lies in the intersubjective foundations of our actions. What we say, perceive, or do is part and parcel of a world known in common (as outlined in Chapter One). Just as in interaction with others, where we assume that they know the world and the situation at hand in ways that are similar to how we know them, we assume that material structures are arranged in a way that draws on and refers to such common knowledge. Art makes no exception here, as its content, aesthetic, and composition must be based on intersubjective principles. Stepping into electronic environments, many people rely on familiar features in their approach to the 'material' arrangements they find there.

But in the context of a diverse and artistically motivated array of approaches to representation of function these could not be expected to mirror real world conditions. Moving from one installation to the next, visitors could be seen to explore through trial and error what worked in what ways. In some installations buttons did perform functions, in others they did not. Sometimes doors, stairs, or paths did lead somewhere at other times they were purely ornamental. The same is true with respect to the dimensions of the electronic spaces and the ‘physical’ laws that applied within the environments. In the ‘Legible City’, for example, visitors could explore representations of three cities (Amsterdam, Manhatten, and Karlsruhe). Each city is based on the real street plan of the city, while the buildings are made up of letters and words. By using a real bicycle, people could travel through these cities. Many people initially stuck to the ‘roads’ and cycled along the words. However, because it was quite difficult to control the path of the bicycle, people inevitably deviated from the roads and even penetrated the words. Once they had discovered that no negative consequences arose from cycling through letters, or across water or grass, people increased their speed and cut through the letters and across the water with no inhibition. When they did reach a row of letters that was impenetrable, often several attempts were made to ‘get through’ this row of letters until it was accepted as the ‘end of this world’. Material clues about function, and manifestations of ‘physical’ laws, and the dimensions of different environments were interpreted with great creativity but, in their variety, a source for confusion. The level of this confusion was profound enough to de-stabilise what can usually be expected of visitors to art galleries. During the first days after the opening the non-interactive exhibits in the more conventional museums were repeatedly damaged through people’s attempts to interact with them.

Cooperation

Similarly, interactional competencies are transposed into the electronic environment. Virtual meetings in CVEs have some precedence in the field of work, where it has been shown that people find ways of tuning their interactional practices to the conditions of the environment. An example would be procedures employed in selecting speakers or claiming a turn in the meeting. In the absence of visually available, non verbal clues of attention and intention to speak, people begin to rely on audible markers. They precede a turn by a succession of sounds in order to claim the floor in a polite and interactionally ‘proper’ way (Bowers et al. 1996). People who have had little or no experience with such an endeavour could be seen to be considerably bewildered as to the nature of their interaction with others in the electronic spaces exhibited at the ZKM. The ‘World generator’ is a networked electronic environment that invites people to choose three dimensional shapes and fragments of text, music, film, or photographs from a menu and to place them into a three dimensional electronic space. The installation is linked to another exhibition in Nottingham where a parallel display is located. A spaceball allows two people - one in each location - to navigate through the space, and to select objects in their vicinity and alter their appearance. Each of them is represented on their remote partner’s screen through a tower-shaped avatar, but there is no avatar for the local person on the local screen. At the same time, a videophone provides a visual and auditory link. The image transmitted via the videophone is mapped onto the avatar in the electronic space. Amongst other things, this work explores the experience of multiple perspectives onto a joint resource – the space and the objects placed into it by the visitors at each location. The following excerpts of talk between a visitor at the ZKM (K), the artist (A), and a visitor in Nottingham (N) shed some light on the difficulties experienced and the strategies employed to work around them:

A: so he's actually looking this way
K: I see I I didn't- I thought he has the identic- ((speaking into telephone receiver)) OK
A: so if you want you can kind of come up next to his avatar
K: how could I how could I *get* his perspective?
A: ehmm OK
…
A: where did he go? *There* he is. OK we'll just chase him. OK he's up there. ((quietly)): what does he see? He's looking off into nothingness
K: what do you see? Tell me, describe it
… ((P is speaking to someone behind him in Nottingham)) …
N: (a flying object on the right side of the screen … *possibly* walking sticks (that's) (0.3)
…
K: what can you see now?
N: err (xxx) I can see a oilrigg on the right hand side of the screen going up and down
K: right
N: and the (0.3) errm (0.3) over on the left side of the screen I can see a tree
K: tree?
N: (floating around)
K: a tree?
A: [try to (xxx)]
N: [(xxx)] (xxxx) turning round
K: you said tree?
N: (xxx)
K: a tree yeah mhm I wonder what you see really. (0.3) I see something blue but it doesn't look like a tree

The nature of the presence of N was a source of interactional trouble, as it was difficult to determine his position in the space and his orientation to, and perspective of, objects in the electronic space. Cooperative activities in our everyday life that parallel this situation provide us with ample information about the task at hand and the nature, course, and potential future developments of the interaction in a way that is taken for granted. We can *see* what is going on between us and the other people, whose turn it is in the conversation, where the other persons' focus of attention is, whether they see what we see or their view is obstructed. The lack of this kind of information in the 'World generator' is made explicit in the above stretch of talk. Here, two strategies are employed to repair the problem. Initially, the artist helps K to 'catch' the avatar of the person in Nottingham. However, in order to maintain this level of contact, K would have to shadow the avatar, a task that is beyond his level of familiarity with the controls of the interface. As another means of establishing common ground, K asks for descriptions of what N sees. But this, too, turns out to be a difficult task, because the shapes and images given in the menu do not readily conform to categories of objects in our everyday experience. A further means of dealing with this problem, observed on a different occasion, was to turn the videophone around to capture the local screen, allowing the remote partner to *see* what was seen in Karlsruhe.

The fact that people are eager to develop strategies that surmount the problems encountered with respect to interactional trouble or difficulties in finding one's way around in an electronic environment have important consequences. It means that electronic environments do not have to simulate real spatial or interactional resources. On the contrary, because 'an essential ingredient of 'the world known in common' are the practical methods we have for 'finding out' or adding to our 'stocks of knowledge'' (see Chapter One) there are few limits to the imagination for the designers of CVEs. However, what this does imply is that we have to provide for the *learnability* of the structure and organisation of electronic spaces.

Learnability

The visitors to the Media Museum at the ZKM had a number of ways in which they found out about the properties of the electronic space and the artistic interpretation of them. Some of these practical methods are quite explicit – in the context of art exhibitions we learn, for example, through observations of the structure of the exhibit, through following instructions, through asking or watching others, or through reading additional information in leaflets, catalogues or the like. Other ways of finding out are more tacit.

Learning from signs and symbols

As described above, people do interpret the material form of things in the physical and the electronic environments drawing on a common stock of knowledge of such things. The artists and the media museum also made some provision for people's inquiry through signs, instructions, animators, tours and written material. During the first three weeks after the opening, the policy was to rely mainly on the self explanatory power of the installations, assisted by 'animators' who provided explanations and demonstrations as and when people required them. However, the level of confusion experienced by visitors found expression in a rather forceful handling of some of the interfaces. As a result, signs with instructions were placed at the entrances of those installations most affected. Through pictorial means and short texts outlining the steps to be taken to be able to interact with the respective installation, people's actions were guided. However, these instructions only provided some rudimentary information. In order to get a richer appreciation of the works, additional information was required. This conflicts with people's desire to see all the installations (and the exhibits in the other museums), because leaflets, tours, and the explanations of animators take some time to absorb. A less time-consuming and more readily available way of finding out more about an installation was to watch others in their interaction with them.

Learning from others

People discovered the function and possibilities of the works through observing others. They can see what people do and what the installation's response is. But it is not through pure imitation alone that people learn. Finding out what can be done can be a collective endeavour (McDonnell 1994). Observation of a succession of different people interacting with the virtual book in 'Beyond Pages' revealed that people do not just learn individually, but through and in interaction with others. After watching another person for a while, B is the next in line to have a go:

> B sits down. Turns the first page, then the second. 'Eats' a piece of the apple displayed on this page. People laugh. 'Eats' the rest of the apple, until nothing is left. Then turns page back. Having seen the pages turned forwards, he now experiments and tries whether they turn back, too. They do, but the apple does not re-appear. He chases a 'stone' off the page. His eyes meet the gaze of one of the children. He stretches out his hand, tries to 'catch' the stone with his hand and laughs. Turns the virtual light switch and the lamp on the table lights up, people laugh. One of the children looks at him and says 'do it again'. So he does. Again, people laugh. The next page says 'Japanese Characters'. He looks down at it. So do the people around him. He looks up at the wall (where some events linked to actions on the surface of the book are projected), so do the others. Nothing happens. He lifts the pen and draws a question mark. At making the dot, a Japanese character appears on the page and a sound is produced. He repeats the question mark several times, then just the dot to produce the sounds. C slips into the chair very eagerly. He had been watching for some time. His use of the pen on the pages is informed by what he's seen. He 'eats' the apple, chases away the stone, … turns on the light and laughs.

B takes over from a person who had performed the actions with the book without having been able to watch someone else before him. His activities had been tentative, and although the group around him reacted to the events on the pages of the book or the screen, he paid little attention to this. As a result, people watched for a while, but then moved on. The second person involves the group in his actions through the way he reacts to their responses and the way he fulfils their requests. The engagement with the book comes to be collaborative in ways that are similar to the ways in which a story is told and received (Tolmie 1997). The 'teller' here is the 'performer'. He monitors people's responses and, as people 'warm' to his story and stay involved over the whole of the period, their actions come to be an integral part of the experience. The learning that takes place in this instance is collective. C picks up from where B left off. He has gleaned information about more than just the functions of the work. How to 'play' the audience is something that he incorporates into his performance. When he comes to the light switch in the book he already knows what kind of reaction this is likely to produce in his audience. Like a person telling a funny story they heard someone else tell, he knows where people laughed the last time. So, when he gets near this point, he prefaces it with a laugh - a display of 'here comes the funny bit'. This performer confirmed the course of events observed in the previous interaction with the work. And part of this achievement is the way in which the story is told.

Learning by doing

Those who got to interact with the installations also learned to find their way around through practical engagement. In 'The fruitmachine' by Agnes Hegedüs, for example, people had to coordinate their actions with two other people. Together, they could put together the three pieces of a virtual, three dimensional 'puzzle' (one half of an extended octagonal shape with images of fruit on each side). Each person had control over one piece via a spacemouse. Here a family with two young children (1) and a group of adolescents (2) are engaged with the work:

(1)

L: I'm not doing anything? (0.3)
V: you're probably the one at the back (.) Laura
L: yes but I'm *not doing anything* (0.4)
V: that's me (.)
L: the one at the front that's me
M: no that's Dad (0.4)
K: which one am I?
L: you're tha that one (.) and I'm the [biggest one]
K: [which one?]
L: you're that one that's just coming to the front (.) no that's *me.* That's *me*.

(2)

C: there is there is one (.) and and if we [got those together now]
A: [wait I'm coming] wait
B: ah I have bananas too
C: yeahyeyeye yes (.) eHH.
B: really that really doesn't work or does it
A: (wait)
C: yesyesyesyesyes
B: look? ((laughs))
C: no eh. what a shame
A: [(xxxxx)]
B: [they are always] in different planes somehow
(1.2) ((clicking, murmuring))
C: has anyone seen some instructions (what this is supposed to [be in the end)]
B: [nehehee no idea] (0.3) come on it works but it (.) ey it just won't fit together
A: will too
B: no
A: just turn a little how did it work just now? (.) so
C: na
((B laughs))

Conversational practices here draw on parallels between virtual spaces and the spaces of board games. Indexical reference to position places the speakers 'into' the electronic space. The children and their parents squabble about who 'is' who in the puzzle, sometimes wiggling their piece of the puzzle to illustrate that they really 'are' this one. In the second example, B states 'I have bananas, too', referring to the images of fruit on the sides of the pieces. She thus anchors her actions in the electronic space to that particular piece by identifying with it. This strategy of placing oneself 'inside' the electronic space could also be observed in the 'Legible City', where people said 'I am on the market square' or 'I'm in Manhatten'. However, this was done in a way that mirrored people's conversations around a game of Monopoly. Objects that a person controlled, avatars, or simply one's position in the electronic space were referred to as representations of one's point of action in this place, at this moment in time

They were 'a way of making oneself present there as well as here. And, in the process, the indexical references to there became here' (Tolmie 1997) This is also a means of achieving intersubjective availability of one's position, orientation, and focus of. Moreover, in the 'Fruitmachine' people could be observed to be most successful at putting the pieces together once they focused on the task rather than the interface. This orientation facilitated the acquisition of a 'ready-to-hand-ness' of the space and the objects in it. People got habituated to the augmented reach of their body in a way that is similar to the way we get used to the dimensions of a new car. The ways in which the interactivity of this work of art and others like it is known and experienced bodily as well as rationally suggests that 'immersion' is something that ties the real and the virtual together rather than taking people 'into' a virtual that is discontinuous with the real (cf, for example, Heim 1995). The engagement with the electronic spaces of the works is not a case of disembodied immersion, it is one of representation, remote action, or instrumentally supported action that is rooted in the real world.

Multiple spaces and electronic landscapes

Through the study of visitors to the ZKM moving between exhibits and their interaction with different virtual world installations, we can start to delineate the kinds of information required to aid navigation in and between different virtual environments. On the basis of our initial study one of the principle technical issues we wish to consider is the need to understand the environment within a context that makes it learnable by users, and the need to convey the manner by which users engage with these environments. In this section we wish to present an initial set of properties that allows us to convey this to users. These properties are used within an environment we have developed to allow a number of worlds to be joined together (see following chapters of this Deliverable).

One of the strengths of cooperative virtual environments is that the developer and designer are able to dramatically alter the nature of the environment. However, this variability is also problematic as it makes these environments difficult to understand. One reason for this is that users do not have a key set of concepts through which they can make the environment understandable.

To address this problem we wish to outline a key set of properties of these environments that can be presented to users in order to allow these to be interpreted and understood by users of these environments. These properties aim to serve the same indicative function of the notices used within the ZKM.

Externalising properties of virtual environments

The problem with generalising aspects of CVEs is that many of these are subjective and difficult to quantify. In addition, someone must be responsible for externalising and providing this type of meta-data about CVEs (typically the designer of the CVE), but how do their notions of the CVE correspond to others? An equally problematic factor in trying to categorise properties of CVEs is the range of possible types of CVE, from 1D text to 3D worlds.

Our approach to this essentially thorny issue relies on providing a simple small set of limited scales of externalised properties. These scales reflect different properties of the environment and are obviously subjective in nature. They explicitly do not aim to categorise environments in any absolute sense or build some form of taxonomy.

It is important to remember that users are being provided with this meta-information in order to support some form of inter-subjectivity between users and to help them learn about the CVE. In effect, users are being provided with a "rough idea" of the CVE, rather than a full description. As a result we can allow a simple indicative scoring to be used to convey the general principles of these environments to a user community.

Our initial set of CVE properties have emerged from a discussion of the variability of the environments within the ZKM in line with a consideration of existing on-line environments. These properties are not intended to be exhaustive but rather they aim to provide an initial set of properties for use in conveying the nature of cooperative virtual environments. Each property has an associated scale typically with five or less categories. These can be considered as analogous to the summary icons often used in reviews and guidebooks.

In the following sections we present some of these properties and the icons used to convey them.

Structural Indicators

Our first sets of properties consider the structure of the space and the nature of this structure. These environment properties provide the basic contents and nature of the world.

Dimensionality	How many dimensions is the environment presented in. This may be 1D, 2D or 3D
Orientation Cues	Does the CVE provide cues for how users should orient themselves inside it (like a city) or by providing a ground plane, or can it be viewed from any position, (as in planet based environments)
Physical Laws	How strongly does the CVE enforce "physical" laws – can the users collide with other objects in the CVE or can they pass through, is there gravity etc.

Abstractness and Urbanity

These properties aim to convey the extent to which the world represents a facsimile of existing physical environments or is an artificial environment. This extends the choice of representation suggested by Benedikt (1992). Rather than make this an explicit choice suggesting a world must be either wholly abstract or facsimile we allow two complementary scales to reflect different arrangements.

Abstractness	How abstract are the objects in the environment.
Urbanity	How urban is the layout of the environment.

Scale and Complexity

This set of properties seeks to convey the scale and complexity of an environment to allow users entering it to understand the extent of the environment.

Size	How big the CVE is terms of the scale used internally within the environment.
Complexity	How dense or complex is the e-scape in terms of the density of objects within it.
Connectivity	Is the CVE connected to a number of other CVEs or is it self contained with no external connections.
Population	how multi-user in the e-scape – or is it single user.
Media	does the CVE provide a rich tapestry of media that can be exploited, such as video and audio, or is it limited to only text

Persistence Indicators

This set of properties conveys the extent to which the environment remembers things and is likely to change between visits.

Stability	does the CVE grow and evolve or is a walkthrough which doesn't change very much
Remembrance (or memory)	When users make changes to objects in the CVE, are these permanent or will they be reset the next time the user visits

These simple property indicators aim to provide a summarised set of subjective measures to users on the nature of the environment and allow these to be made available at a glance.

These indicators can not in themselves convey any real sense of the aesthetics of the environment or activity within it. To allow this to be conveyed we also provide a simple snapshot of the environment updated at regular intervals (currently this snapshot is a 2D image but we could exploit some form of 3D thumbnail (Elvins et al. 1998)). The provision of this snapshot is motivated by the use of windows and glances inside exhibits reported in the ethnographic study.

Engagement Information

In addition to the presentation of the properties of these environments we also wish to reflect the means by which users can interact with them. The consideration of this information directly reflects the issues of engagement observed in our study and the need to provide cues to allow users to manage their relationship with these environments.

Essentially, we need to move beyond the current view of interaction based on an interface locked to a screen and a single point of interaction controlled by the mouse (and keyboard).

The starting point for our consideration of engagement is the work on action points and viewpoints initially undertaken by Benford and Fahlén (1994). We represent the manner by which users can engage with a virtual environment in terms of three different characteristics for determining interaction and presence within the space.

View of the environment - This conveys the particular views offered by an environment. Depending on the environment users can have multiple points of view. At any moment in time each point of view will have both *a location* and *a direction*.

Action point in the environment - This conveys the point (or points) of interaction a user may have in the environment. This is the means by which objects are selected and acted upon.

Position in the environment - This conveys the presence of users within the environment and the means by which this position depends on the other points.

The aim here is to convey the relationship between the user and the environment in terms of how they interact with it. These three engagement properties convey the affordances for engagement offered by an environment. For each of these different properties we also encode the control details and they can also be linked together to show different restrictions in engagement enforced by the nature of the environment.

As an example consider the arrangement offered by a simple virtual world provided by DIVE. Users are represented in the world as simple embodiments and these embodiments convey position. In addition, a user's viewpoint is linked to the orientation of this embodiment while their action point is independent of this. This provides the engagement properties suggested in Figure 1.

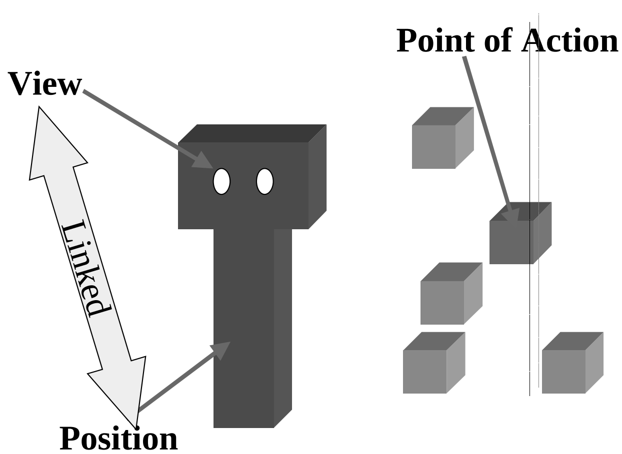

View	
Control details	Vehicle
Linking	Location -Position Direction - Position
Action	
Initiator	Select
Control details	Mouse
Linking	
Position	
Control details	Vehicle
Linking	View

Figure 1 Engagement properties for a DIVE world

These properties can also be used to describe more radical worlds with a different range of interaction affordances. Consider for example, the fruit machine exhibit described in the study. Here three different users are offered independent action points to rearrange pieces in the environment. The environment is presented using a single fixed viewpoint and no means is provided to represent users in the environment. Rather they are considered external to the shared environment. This arrangement is shown in Figure 2.

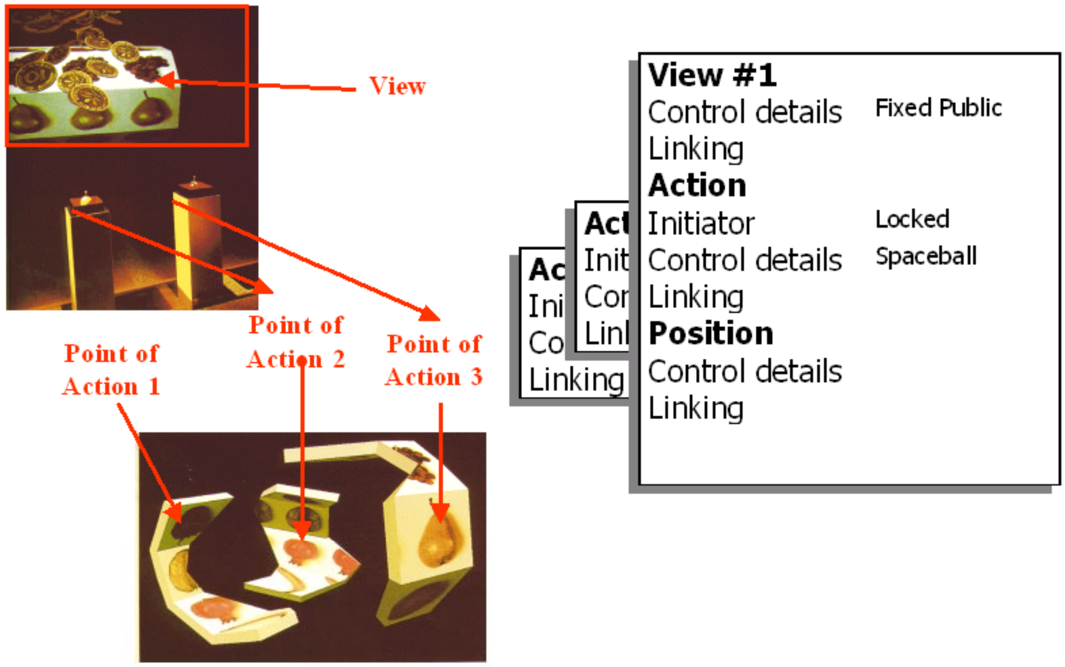

Figure 2 Engagement properties for the fruit machine

Where appropriate these different engagement properties can be made available to users of environments using some form of iconic representation showing the values associated with these three properties. However, to date we have focused on understanding interaction with these environments more generally prior to presenting these different engagement properties.

Representing interactional affordances

The engagement information provided by the environment represents only a small part of the interaction affordances offered. In fact, we can consider a user's interactional affordances with a given virtual environment as a combination of three things:

Engagement properties of the CVE –. This is essentially represented in terms of the action point, view point and position offered by the environment.

Capabilities of the interaction device – what the user's current interaction hardware/software allows the user to do. For example, a user interacting at a desktop computer may have multiple windows (allowing many views on the same CVE simultaneously) but can only physically interact in two dimensions using a keyboard or a 2D mouse. However a user within the same world may be wearing VR-goggles, which only allows a single view but can interact with the CVE in three dimensions using a 3D mouse.

Capabilities of the users vehicle – Users within a virtual environment often exploit some form of *vehicle* that provides a mapping from the capabilities of an interaction device to movement in the environment. Vehicles allow a range of different navigation profiles to coexist in the one environment. For example, the default avatar may restrict a user to movement in a 2D plane, whereas a *helicopter vehicle* may also allows the user to move up and down.

What types of interaction a user can perform within a virtual environment at any given moment is a combination of these three capabilities. There are at least two advantages gained by externalising these capabilities to the user (or making them accessible to other applications and CVEs).

-- The current interactional capabilities of the user in any CVE can be presented to them, helping them to understand what is possible and how the CVE may be interacted with.

-- The interactional capabilities of other users can be made available to others allowing inferences to be made about their experience of the environment and some sense of intersubjectivity to be build up across the environment.

This arrangement also allows some “policing” of users travelling between CVEs. For example, one particular CVE may require several simultaneous views to be presented to the user. If the users interaction device is an immersive headset then the CVE may either signal the problem to the user, who can then experience some denuded form of the CVE, or may deny access altogether.

External access to environment properties

The previous sections have considered how we may wish to represent the nature of cooperative virtual environments. To make these properties most useful to users we need to make these properties dynamically accessible from outside the environment. Three possible solutions can be used to make these properties available to external applications to present them to users:

1. Adding meta-information nodes to CVE scenes

Most CVEs or VR systems can support nodes that are not seen in the world. Each platform can use these nodes to provide a specialised node that encodes the environments properties. A problem with this approach is that the CVE must be fully loaded to access the meta-information.

2. Adding meta-tags to HTML pages

Many CVEs use URLs to define world locations. The META tag provided by HTML allows miscellaneous information concerning the contents of the main body of the URL to be recorded which can be fetched without getting the whole page. Such information can be the external properties the CVE world contained within the HTML page.

3. Providing a CVE registry service

A CVE registry service may either supplement or supersede the previous two mechanisms. The CVE registry records meta-information about a CVE. When it registers itself as a known CVE, the registry creates a new entry. The entry will contain the properties above along with additional management information. Applications wishing to access these properties can contact the registry service and query it for the appropriate information.

None of these mechanisms are mutually exclusive. Indeed, certain approaches are more suited to different occasions. For example, where a system does not provide a URL access point to a world, mechanisms 1 and 3 are suitable. Where URLs are used, 2 or 3 would be better. In addition, it would be difficult to directly store the current “snapshot” of the CVE using (2), although a reference to a location (e.g. a URL) where the snapshots can be found could be provided. The use of these techniques to manage the representation and presentation of these properties is central to an extended CVE session control model and system developed to support the joining together of CVEs to form electronic landscapes.

Presenting and using CVE properties as Points of Presence

The mechanisms presented above allow us to maintain information about the properties of CVEs and allow users of these environments to access these properties to support their activities within these environments. The information represented by these environments can be presented to users in a variety of ways. Each of these different techniques we term a point of presence where the e-scope makes its properties present within other escapes. In this section we briefly consider some of the different presentation mechanisms that we have developed to convey the information about the nature of these environments.

User profiles

When we talk about users moving from one CVE to another, we mean the actual engagement of the end-user themselves, rather than any particular embodiment, avatar or vehicle. Indeed in all likelihood nothing will actually “move” or be transferred between CVEs supported by different applications. In such a situation, the old information about the user is lost (or must somehow be recorded by the CVE when the user leaves) before the user enters the new CVE.

However, the emergence of CVEs that are inter-connection implies some common “thing” travelling between these worlds.

A user profile is state information associated with the end-user, which can be accessed (read from or written to) by any CVE. This allows each CVE to record information about a user as they leave the environment that can be recovered when the user returns at some later time. For example, if a user picks up several objects in a CVE and then leaves, the CVE could put these objects into that user’s profile. When the user returns, the CVE can query the profile to get the objects back. Where different types of CVE understand the format or rules of another, the user profile would allow objects to actually move between environments.

Further a single user-profile provides a logical point for information about the user’s current interaction capabilities (mouse, monitor, keyboard, or 3d-mouse and immersed-headset etc.) to be stored. This information can subsequently be exploited by the CVE (or gateways) to warn the user of possible discrepancies between the CVEs interactional capabilities and the users current capabilities.

Annotated gateways

In normal virtual environments, gateways transport the user from one environment or world to another. This transportation is performed when the user interacts with the gateway in some fashion (clicking on it, walking through it etc.). In order to heighten a user’s awareness of exactly what the new world is like, to lower their potential bewilderment and aid learnability, annotated gateways present meta-information about the world on the “other-side” of the gateway. An example of an annotated gateway in a 3D CVE may show externalised properties of the world as icons around the portal. Figure 3 shows such a gateway.

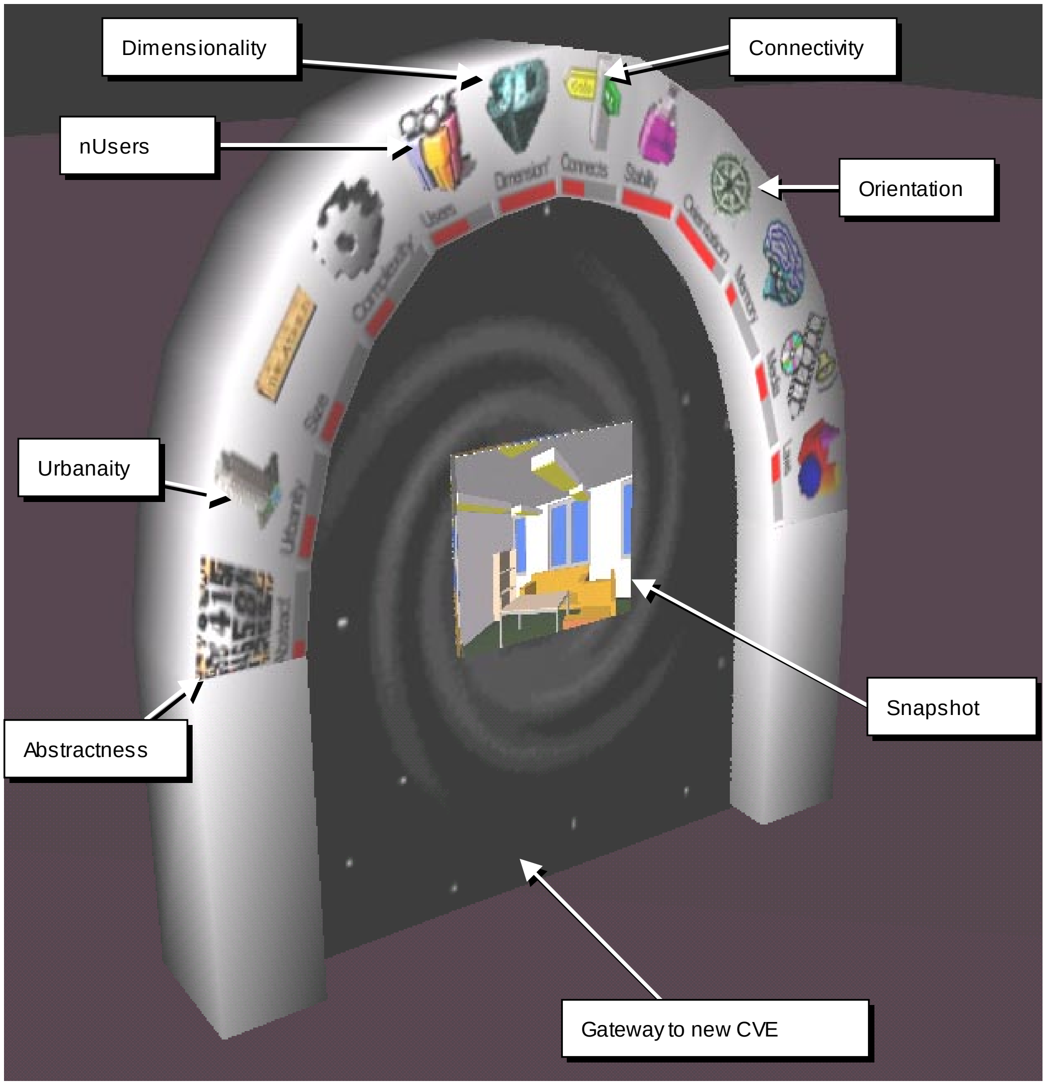

Figure 3 An annotated gateway Point of Presence

An annotated gateway could lead directly to a new CVE, or for more novice users, could lead to a more elaborate reality lock.

Reality locks

Reality locks are inhabited spaces that link CVEs together. These can be considered as drawing directly from the use of the spaces surrounding the exhibits within the study. These locks provide access conduits that highlight and teach users travelling between the two CVEs about the changes and differences between the source CVE (they have just left) and the destination CVE (they are travelling to). For example, consider moving from a virtual world using a city metaphor to an environment exploiting an abstract visualisation of information. In this case the lock would attempt to show that the new CVE no longer presents a facsimile with navigational clues that mimic the real world, but instead contains large amounts of abstract data in a three dimensional space.

There is one reality lock associated with each pairing of CVEs. For example, all users travelling from a CVE ‘A’ to CVE ‘B’ will enter a particular reality lock, whereas all users travelling from CVE ‘C’ to CVE ‘B’ will enter a different one. The reason for different locks connecting to the same CVE (e.g. ‘B’) is that users may be travelling from very different environments to reach the CVE, where each starting CVE may require very different property changes to be imparted.

The inhabitants of locks may be real-users that happen to be travelling between the same CVEs at the same time, or automatons that interact with users to provide particular information. For example, “tourist” guides who provide information about the environment In addition to presenting the differences between the CVEs, in an iconic form, reality locks can also provide tools which teach users useful skills in the new CVE or convey these using some form of animation. A simple example of this is an animation showing how users navigate through the environment or initiate some environment specific action.

Finally, the ability to access periodic snapshots of the destination CVE allows the reality lock to present users with a history of the CVE, showing how it has changed and evolved. Figure 2 shows a “history tunnel” which lays out these snapshots, from the earliest snapshot to the latest. Users pass through each snapshot to reach the next although they are free to move backwards and forwards in the tunnel as desired.

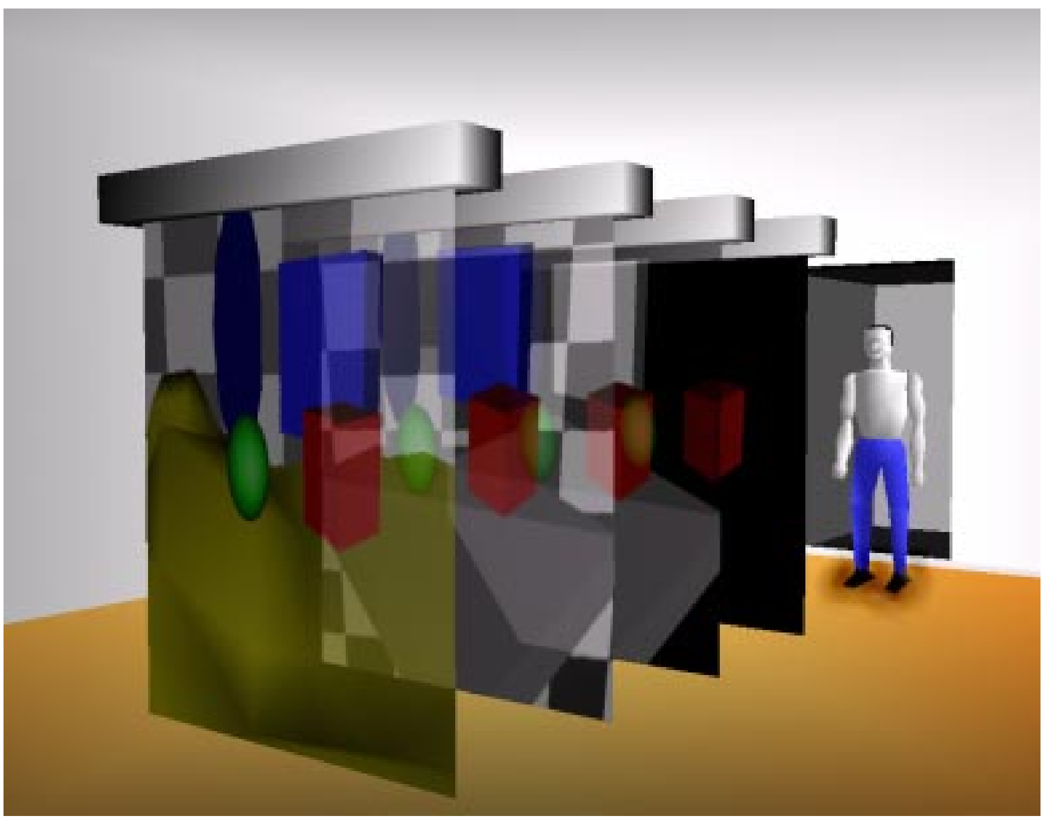

Figure 4 The history tunnel Point of Presence

The annotated gateways and reality locks (such as the history tunnel) are intended to provide users of virtual environments with cues that promote learnability and allow users to develop a sense of intersubjectivity as they develop an understanding of what is known in common within these virtual environments.

Summary and Conclusions

In this chapter we have presented the development of a set of facilities that support users co-operation across different virtual worlds. These mechanisms have been informed by an ethnographic study of visitors to a multimedia museum containing a number of different virtual worlds.

The core of our approach is the representation and externalisation of properties of virtual environments. In this chapter we have suggested an initial set of properties that provide a "rough guide" to the virtual worlds. We have also briefly reviewed a number of ways in which these properties can be accessed and presented to users.

The use of these properties as external cues aims to support users as they traverse a collection of virtual environments. We have developed a supporting infrastructure that allows a number of electronic virtual environments to make these properties externally available to allow the construction of electronic environments that join different virtual worlds together. We plan to augment this work by developing specialised browsers that support and promote the users perusing a collection of virtual environments.

References

Alphaworlds, http://www.alphaworlds.com

Benedikt, M., *Cyberspace*, Cambridge, Matt, MIT Press, 1992

Benford, S. D. , Fahlén, L. E. Viewpoints, Actionpoints and Spatial Frames for Collaborative User Interfaces, in *Proc. HCI'94* (Glasgow, August 1994).

Benford, S. D., Bowers, J. M et al.., Supporting Co-operative Work in Virtual Environments, *The Computer Journal*, 37 (8), Oxford University Press, 1994.

Benford, S., Bowers, J., Fahlén, L et al.., User Embodiment in Collaborative Virtual Environments, in *Proc. CHI'95*, (1995). ACM Press

Benford, S., Greenhalgh, C.et al.. Staging a Public Poetry Performance in a Collaborative Virtual Environment, in *Proc. ECSCW'97* (Lancaster, UK, Sept.7-11. 1997), Kluwer. 125-140

Benford, S., Mariani, J. Virtual Environments for Data Sharing and Visualisation – Populated Information Terrains, in *Proc. ICDCS'94* (Lancaster, UK, July 13-15, 1994), Springer-Verlag. 168-182

Benford, S., Snowdon, D., et al. Informing the Design of Collaborative Environments, in *Proc. GROUP'97* (Phoenix, Arizona, Nov.16-19, 1997), ACM Press. 71-79.

Bowers, J. et al. Talk and Embodiment in Collaborative Virtual Environments, in *Proc. CHI'96* (Vancouver, Canada, April 13-18, 1996), ACM Press, 380-389.

Carlsson, C., Hagsand, O. DIVE: A Multi User Virtual Reality System, in *Proc. IEEE VRAIS*, (September 1993), 394-400

Curtis, P., Dixon, M, MUDS grow up: Social Virtual Reality in the Real World, *in Proc. of IEEE Computer Conference*, (January 1994).

Elvins, T., Nadeau, D., et al.. Worldlets: 3D Thumbnails for 3D Browsing, *Proc. CHI'98* (Los Angeles, CA, April 18-23, 1998), ACM Press.

Harrison, S., Dourish, P. Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems, in *Proc. CSCW'96* (Boston, MA, Nov.16-20, 1996),. 67-76

Heim, M. 1995. The Design of Virtual Reality, In: Featherstone, M. and R. Burrows. Cyberspace, Cyberbodies, Cyberpunk. London: Routledge. 65-77

Mansfield, T., Kaplan, S., et al. Evolving Orbit: a progress report on building locales, in *Proc. GROUP'97* (Phoenix, Arizona, Nov.16-19, 1997),. 241-250.

McConnell, D. *Implementing Computer Supported Cooperative Learning*, London, Kogan Page, 1994

Roseman, M., Greenberg, S, TeamRooms: Network Places for Collaboration, in *Proc. CSCW'96* (Boston, MA, Nov.16-20, 1996), ACM Press. 325-333

Tolmie, P.. Representation, CSCW Workshop Working Document. Working Document. Internal report available from Dept. Sociology, Lancaster University, 1997.

VRML97, International Standard ISO/IEC 14772-1:1997, "The Virtual Reality Modeling Language "

Chapter Six:
Session support for e-scapes

Jonathan Trevor, Gareth Smith, Tom Rodden

Lancaster University

Introduction

In this chapter we consider the growth in interest in virtual environments and the genral focus on co-operation within these environments offered by the majority of these approaches – little consideration has been given to users' management of these environments and their movement between them in the manner outlined in the previous chapter of this Deliverable. An integral element of first year work for the eSCAPE project has been an exploration of the means by which support might be provided for bringin electronic environments together. An important part of this work has been the development of an architecture built upon the HTTP protocol that supports session management for virtual environments. Accordingly in this chapter we present the architecture and its use to support both virtual environments and more generic cooperative applications.

Support for on-line virtual worlds

Over the last few years a large number of systems to support on-line virtual worlds and collaborative virtual environments have started to emerge. These include 3D virtual environments such as: MASSIVE (Greenhalgh & Benford forthcoming); DIVE (Dix 1997); AlphaWorlds; Blaxxun; 2D graphical systems such as the Palace; and textual virtual environments such as MUDs (Edwards 1994). The large number of these systems and their rapid development is impressive. However, the diversity of these systems is potentially problematic - particularly given the generally insular philosophy adopted by these systems.

Users seeking to exploit these cooperative environments are currently faced with worlds that are simultaneously massively interconnected *within* each system and yet isolated *between* systems. It is not unusual for users to face problems of managing the interconnection between many different worlds within virtual environments supported by a single system. Users are equally likely to become frustrated by the closed nature of these systems and the inability to seamlessly move between environments supported in different worlds. This is despite the development of standards such as VRML (ISO/IEC 14772-1:1997) which focus primarily on the visualisation of these virtual environments.

This chapter seeks to directly address two frustrations for the users of online virtual environments:

-- The problem of managing the growing scale and complexity of virtual environments.

-- The problem of heterogeneity of supporting environments and the inability to provide links between virtual environments supported by different classes of system.

We have addressed these needs by undertaking a ethnographic study of users moving between virtual environments within a multimedia museum, studies which highlighted the need to allow users access to information about environments from outside them. In this chapter we present a supporting architecture that allows this information to be made public.

In the rest of this chapter we describe an architecture to allow an extensible approach to managing on-line virtual environments sessions. We present the developed architecture and demonstrate its use to provide access and connectivity between different virtual environments. In doing so we aim to provide a set of management services that exist *independently* of any single virtual environment and allow users to exploit information about the nature of these environments when they use them.

Core to the development of our supporting infrastructure is recognition of the current trend towards the formation of large-scale *electronic landscapes* formed from the interconnection of virtual environments. The notion of electronic landscapes is predicated on a vision of socially oriented information systems for the citizen. The assumption is that the increased growth of accessible distributed infrastructures (such as the World Wide Web) will result in distributed accessible information spaces capable of supporting social forms of interaction. What is required however is a shift from our current considerations of these virtual environments as closed entities to a more open arrangement of these worlds. In particular, we need to consider how these environments make key features about their nature more publicly available.

From Space to Landscape

The formation of electronic landscapes builds upon work in HCI and CSCW on the use of spatial metaphors and techniques to represent information and action in electronic systems. This work is founded in the use of a ‘room' based metaphor to allow the presentation of information (Henderson & Card 1985) From these early approaches we have seen concepts of spatial arrangement exploited in the development of desktop conferencing systems such as Cruiser (Root 1988) and more generally in the work of Mediaspaces (Gaver 1992).

The development of co-operative systems within CSCW has also seen a growing application of concepts drawn from spatial arrangements. These include the development of *groupkit* to form *teamrooms* (Roseman & Greenberg 1996) the emergence of the *worlds* system (Fitzpatrick 1996) and the use of a notion of places to support infrastructure (Patterson 1996). This notion of using concepts of space has considered equally the augmenting of existing physical spaces to form augmented spaces populated by electronically sensitive physical artefacts (Ishii & Ulmer 1997) and virtual spaces into which users project their presence (Benford et al. 1995). Our particular interest is in the formation of on-line electronic landscapes that is both large scale and distributed in nature.

The metaphor of ’landscape’ is deliberate in our consideration of these electronic environments. We seek to develop more general concepts that extend existing considerations of the use of space as a metaphor within virtual environments to consider its more general application across these environments. We wish to directly consider the issues of diversity and scale inherent to the more general use of spatial approaches as we move beyond virtual environments.

Extending virtual environments

The work reported in this chapter seeks to extend the work currently being undertaken through a number of initiatives that have attempted to address the development of shared online virtual environments. This work has its history in Muds and Moos (Edwards et al. 1997) but has more recently turned to the use of VR technology to visualise a space shared by multiple individuals.

Virtual environments have been used to directly support a number of different forms of cooperation. For example, there exist in the research literature several descriptions of virtual conferencing (Greenhalgh & Benford forthcoming) and collaborative information visualisation and retrieval systems (Brodbeck et al. 1997). The work of these systems to promote and support co-operation within them has become a major area of CSCW research with the emergence of a cooperative virtual environments research community.

Proponents of shared virtual environments often claim that they support social interaction in ways which go beyond the limitations of using more familiar technologies such as videoconferences, ’mediaspaces’ or shared desktops. Crucially, shared virtual environments permit users to become *embodied* within a shared space by means of an embodiment or ’avatar’. It is often argued that avatars permit a degree of self-expression and it has also been claimed that appropriately designed shared virtual environments enable users to sustain mutual awareness about each other’s activities (Bowers et al. 1996). Commercial trends certainly suggest that these environments will become increasingly accessible to the citizen. This pressure towards growing accessibility motivates the need for a consideration of the general structure and support of these environments.

This research work is complemented by the continuing efforts of a number of organisations that provide access to shared virtual environments via the Internet (The Contact Consortium). Of these AlphaWorld, Worlds Away, The Palace and Blaxxun are some of the best known. While the basic VRML (Virtual Reality Modelling Language) standard for distributing models of virtual environments over the Internet does not provide explicit support for simultaneously shareable worlds, although already suggested for VRML 3.0 ("Living Worlds"). In all these respects, there are a number of existing and imminent developments that offer early versions of the emergence a future large-scale electronic landscape.

Overcoming limits of existing systems

While a number of different systems postulate the view of large-scale virtual environments populated by many users, their view of the world tends to be closed within the worlds supported by a single system. While VRML (ISO/IEC 14772-1:1997) offers some sense of standardisation it still does so with a focus on issues of visualisation and presentation rather than offering a wider consideration of virtual environments. It is clear that a number of systems will continue to coexist, each exploring particular approaches in supporting virtual worlds.

One model common to most virtual environments is the partitioning of space into discrete components linked by some form of porthole or gateway. Users traverse the environment by using these gateways to move between worlds. However, while this partitioning allows large scale virtual environments to be constructed it also hides the details of these different worlds from users as these separate worlds do not provide any external information about their nature. This partitioning in isolated closed worlds causes a number of real problems for users.

-- **Support for heterogeneity** is limited in that each virtual environment is closed and does not provide external access to information about the environment or how to access it.

-- **Browsing collaborative virtual environments** is difficult as no information on the nature of these environments is provided without entering these using the particular application that supports them.

-- **Seeing how currently inhabited virtual environments are** is difficult as few provide on-line access to information about who is currently connected. Again those that do provide this information require users to enter the environment.

-- **Configuring and launching environments** is difficult to do and it currently requires some considerable knowledge of the detailed nature of the application.

-- **Understanding the nature of the environment** is difficult as users are seldom provided with cues that allow them to orient to the particular nature of the environment.

Currently users are left to manage the information needed to address these problems and there are no techniques that allow different forms of virtual environments to be knitted together in a seamless manner. This chapter presents the development of an architecture that allows such information to be managed externally to the environment, and provides facilities that allow both users and applications to access and exploit meta- information about environments. The core to the development and support of this meta information is the notion of an e-scape.

The general architecture

The development of an architecture that allows a range of virtual environments to coexist and be reasoned about is dependent on the notion of an electronic landscape (or e-scape). An e-scape is any electronic world where sets of properties that provide an indication of the nature of the environment are made public. Thus an e-scape provides support for the development of applications and virtual environments that join these virtual environments together.

The current architecture provides facilities for an extensible set of environments' properties to be maintained. In addition to these user-defined properties each e-scape has a core set of properties that aim to convey the nature of the virtual environment. These properties build directly upon those listed in the previous chapter.

The e-scape session architecture draws upon three separate services to provide support.

-- **A User Profile Service** that maintains information about users and their relation to e-scapes.

-- **An E-Scape Registry Service** that allows different environments to register as e-scapes by making meta-information about the environment publicly available.

-- **A Personal E-Scape Service** that runs at each client that interprets and locally manages e-scape information.

The main aim of the e-scape session service is to provide support for the management of large-scale electronic landscapes that are both open and accessible by a wide community of users. In addition, the principal pattern of communication needed for establishing e-scape sessions is a client-server. As a result of these two requirements we have adopted the HTTP protocol as the supporting mechanism for transporting requests and replies. In addition, since most of the interaction itself does not require any immediate active feedback between the services we can also exploit HTML to provide general management interfaces for the different services. Our use of these two standard and widely adopted protocols has many advantages (Dix 1997) not the least of which is that of deployment and access to a large established community of users. The general arrangement of these services is shown in Figure 5

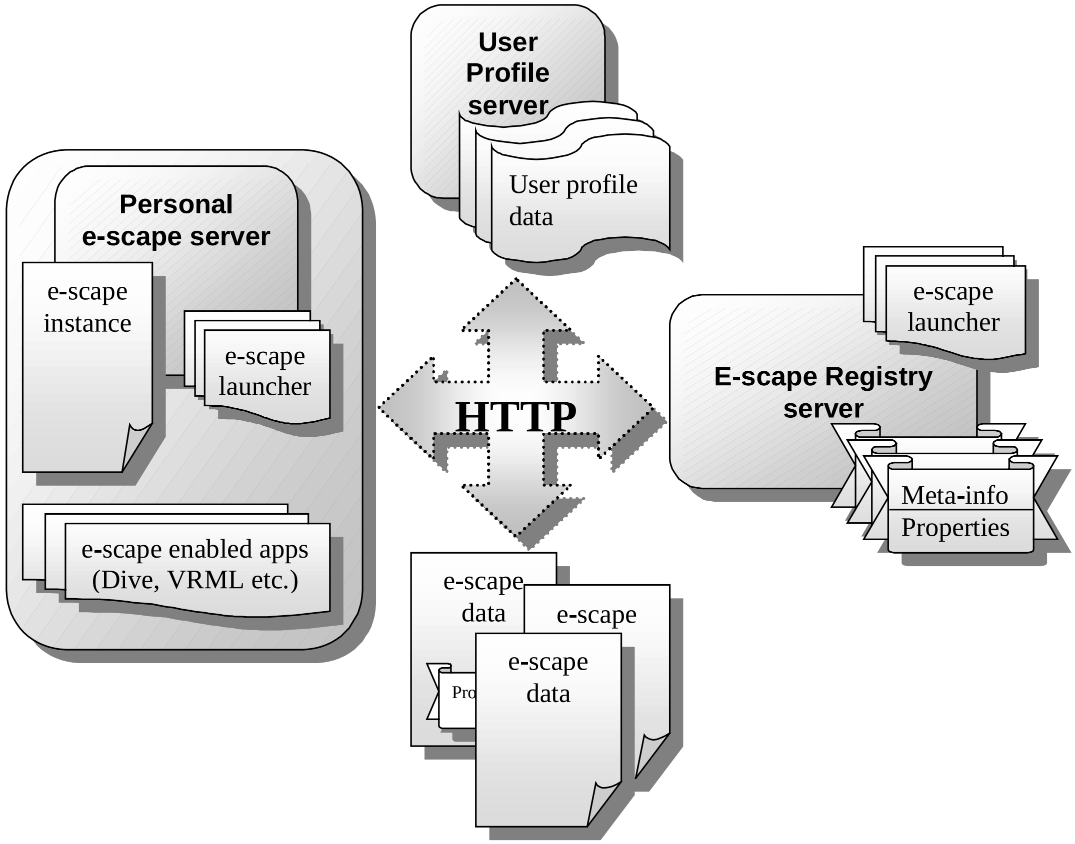

Figure 5 The e-scape session architecture

The general arrangement shown in Figure 5 exploits a distinction between the model of the e-scape and the means by which this is launched. The architecture of most virtual environments systems maintains a separation between a centralised world model that is shared and a local client that manages interaction. The e-scape architecture exploits this separation by providing appropriate launchers for different types of virtual environment.

User Profile Server

The User Profile Server maintains information about the users of e-scapes and allows this information to be shared with other users when necessary. When first created a user profile contains only a few values:

-- The *name* of the user, e.g. Jonathan Trevor.

-- An *alias* for the user, e.g. Jonathan.

-- *Personal e-scape server* URL.

-- An Authentication *password.*

The e-scape service allows these profiles to be extensible to allow additional information to be maintained as needed. Once created the profile can be extended by a user or anyone with sufficient access (see below). Each profile supports any number of additional name-value pairs. To aid access these pairs are stored in a similar way to the Windows95/NT registry within a hierarchy of folders, where each folder can have many name-value pairs and other folders. The server allows access to information in these folders through normal HTML and can be modified through HTTP GET and POST requests. The amount of information made accessible for each user profile depends on the access rights of the application or user accessing the server. Figure 6 shows the details of Jonathan's user profile.

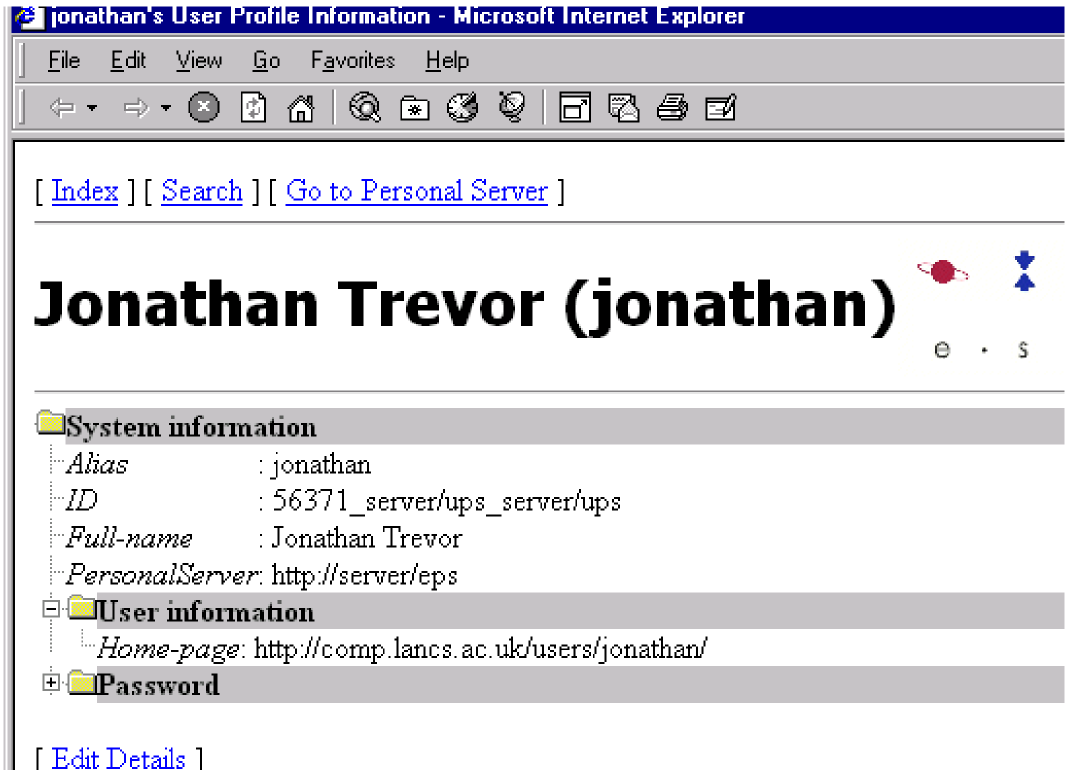

Figure 6 A user profile

The architecture allows many different instances of the user profile server to be running simultaneously. Each server maintains zero or more user profiles and each e-scape user can only have one profile at one server (although the profile may move during its lifetime between profile servers). For example, both Tom and Jonathan may have their profiles stored at the URL http://server/ups, but Jam has his profile at http://anotherserver/ups. The general location transparent access provided by the HTTP protocol allows for large scale global federation of these servers and means we can exploit the general scalability provided by the internet.

E-scape Registry Server

The registry service is responsible for maintaining all the meta-information for an e-scape. This information is added to a server when users and/or developers register the e-scape with the server. The registration of the meta information associated with the URL combines both management information for the e-scape and a set of general properties of the e-scape. Registration is simply a case of filling in an HTML form that allows the specification of:

-- the *location* of the actual e-scape data

This is typically a URL giving the location of an access point to the shared world.

-- a *description* of the e-scape

This provides a textual overview of the nature of the world.

-- a *snapshot* of the e-scape

This provides a view of the e-scape and can be either a 2D image or a sample 3D VRML world. This could make use of the some form of 3D thumbnails[0].

-- Current *users* of an e-scape

The registry maintains a list of users currently connected to an e-scape.

-- the *type* of e-scape, (e.g. Dive, VRML2, Quake etc.)

These are normally specified in terms of mime-types allowing the appropriate launcher to be initiated by the personal e-scape server held locally.

-- the hardware and software *requirements*.

What requirements does the e-scape need in order to run.

-- any other *name-value pairs*.

In a similar manner to the user profile service, any number of name-value paired parameters may also be added to the meta-information and structured into a folder hierarchy.

-- the e-scape *properties*.

These provide a general indication of the nature of the environment. These are held as a series of name value pairs for the general properties of the environment.

Similarly to the profile servers, the e-scape session architecture allows any number of e-scape registry servers where each server can be accessed through single URL, e.g. http://server/ers.

Presenting e-scapes

Once the meta-information about an e-scape is on the server, users are free to search and browse it. This can be done using a range of different interfaces. The simplest of these is the HTML interface provided by the server, as shown in Figure 7. Searching can be based on matching arbitrary name-value pairs in an e-scapes registration information, or may rely on the fixed fields like the e-scapes name, type or properties.

In this simple interface users are presented with a link at the bottom of each page of information about an e-scape that allows the user to start and join the e-scape on their local machine. This is done using a URL that causes the browser to execute a HTTP request URL (with appropriate parameters) to the users Personal E-scape server.

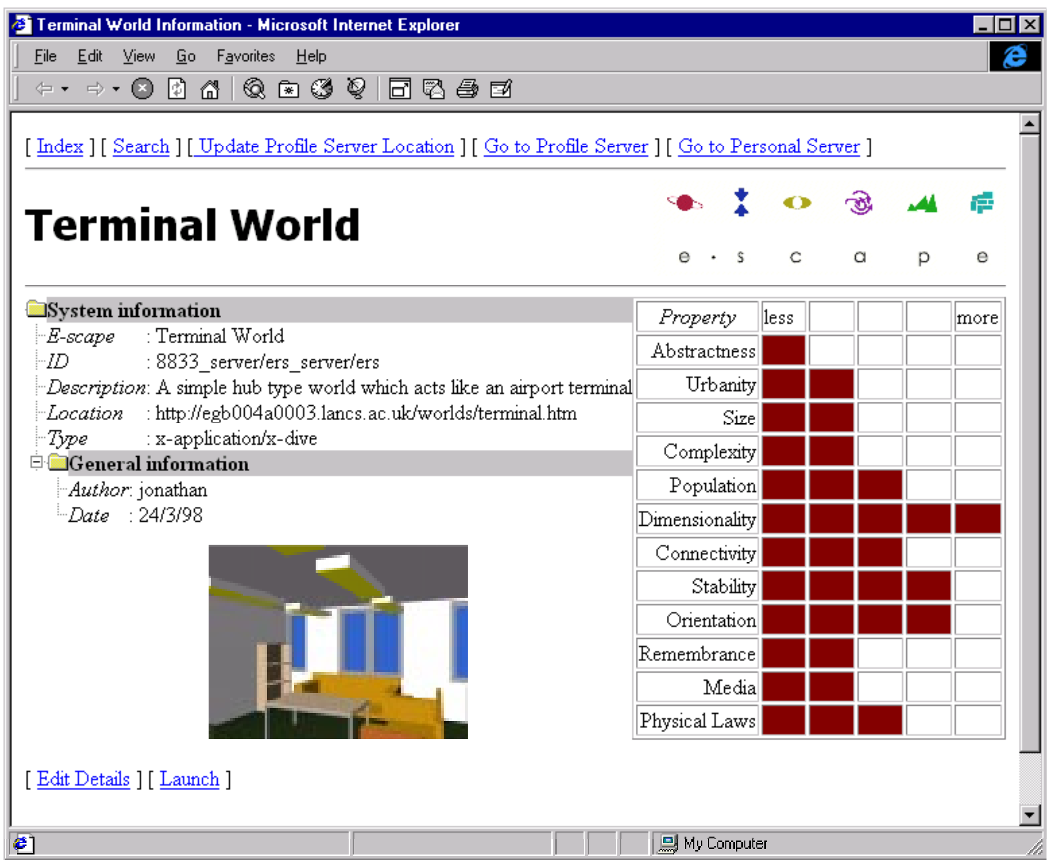

Figure 7 A registered e-scape

Other presentations of the information in the registry are possible and alternative displays can be constructed from the e-scape properties. This includes the generation of appropriate virtual environments that provide access to the local information. Environments and e-scapes can also access registered information to present it as annotations on gateways to these worlds allowing users to gain an idea of what to expect before actually travelling to the new world (figure 4).

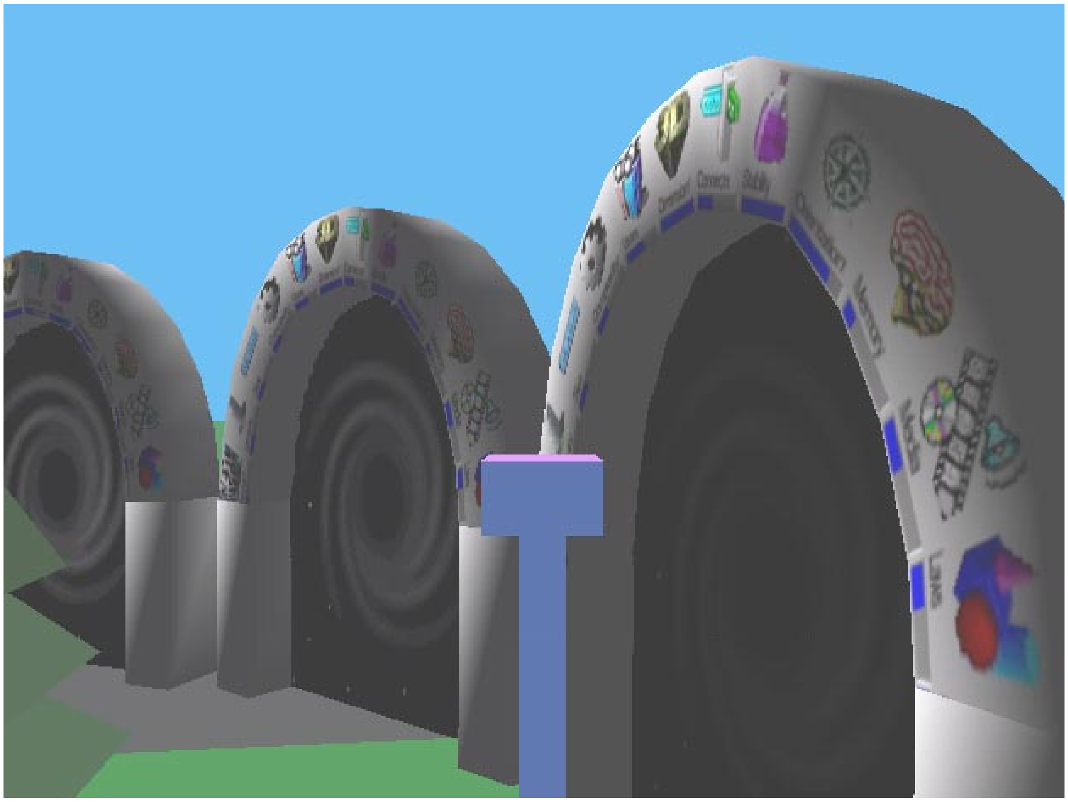

Figure 8 Annotated gateways to other e-scapes

In addition to providing a source of meta-information concerning available e-scapes, each registry server also contains a set of "*Launcher*" Java classes. These classes can be transported from the registry server to the user's local machine in a similar way to applets being fetched to a local browser. The launcher classes provide wrappers that allow e-scapes to be set-up at started.

Finally, each e-scape registry server contains an editable list of URLs to other known e-scape registry servers (on other machines). This linking of registries allows users to perform searches across a uniform federation of servers. Because URLs are location independent, browsing users can easily find themselves seamlessly moving from registry server to registry server. Each e-scape may only be registered with one server at once and facilities are provided to migrate the registration information between servers.

Personal E-scape Server

Each e-scape user requires a personal e-scape server to be running on their local machine. Again, the server is manipulated through HTML pages it provides and can be accessed from a single URL, e.g. http://server/pes. This standard interface allows users to configure, access and manage the information in the Personal e-scape server.

The personal server also presents a more active interface that provides a graphic display of the properties of e-scapes the user inhabits. This interface is written in Java to overcome the static nature of HTML pages. Without such an interface remote servers would have to be repeatedly polled for the information to see if it has changed. Figure 9 shows the Personal E-scape Server Java GUI. This interface is analogous to the use of awareness widgets within other cooperative systems in that it keeps the users aware of the nature of the environment they are currently using.

Figure 9 Java Personal E-scape Server Interface.

Accessing E-scapes

In addition to providing a dynamic indication of the nature of the environment being inhabited the personal e-scape server manages the local launching of applications that start and interface to virtual environments. In order to use the different types of e-scape on their local machine, the applications which can provide access to those e-scapes must obviously also be installed. For example, a Dive world requires a local Dive viewer and supporting processes, a VRML97 world requires a VRML97 capable browser, a Quake map requires the Quake game and so on.

These applications need to be started in different ways, take different parameters, require the data about their e-scape locations to be present locally or remotely, and so on. To overcome this problem we use the notion of *e-scape launchers* which can be fetched from remote launcher repositories (i.e. the e-scape registry server) as required. Once made available locally, each launcher may need to be configured further (with the location of the e-scape application for example). Such configuration is performed again through standard HTML via the personal e-scape server. Once suitably configured, each launcher can store this information locally to be reused later. Figure 10 shows an example of the configuration interface presented by the Dive launcher on NT4.

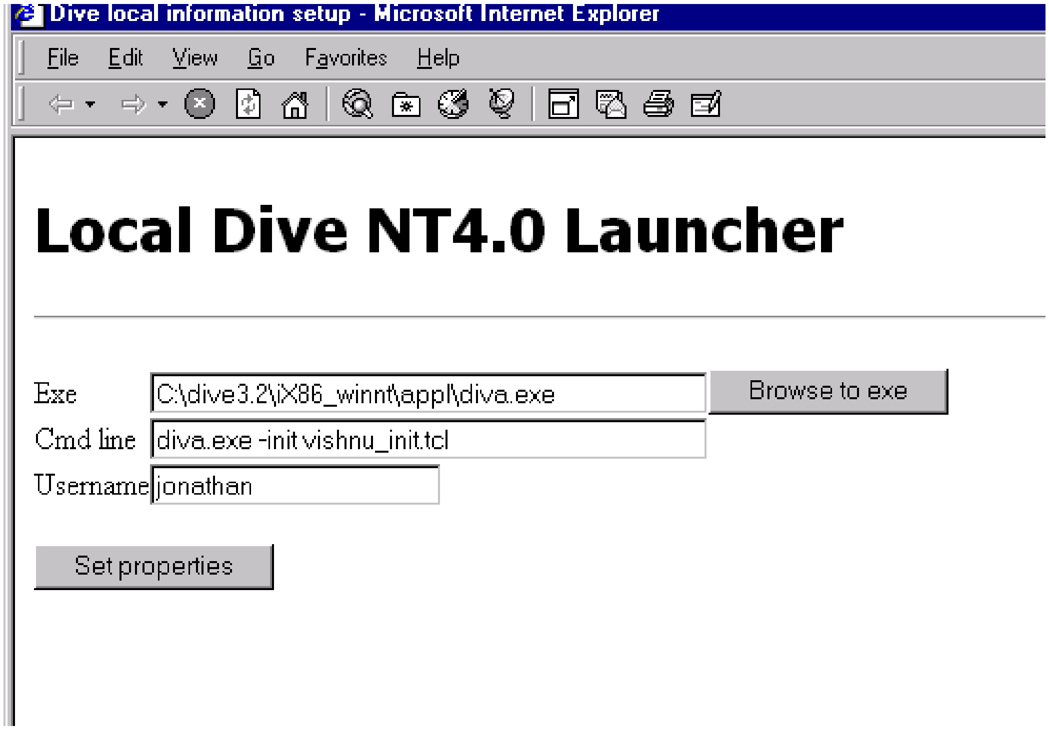

Figure 10 A launcher configuration screen

In addition to launching particular types of e-scape on the local machine, a launcher is responsible for augmenting the basic virtual environment application with e-scape aware features. E-scape aware applications can make public certain users activities. For example it can detect a user leaving the e-scape (e.g. by quitting the application); entering a new e-scape. It can also extract the e-scape properties from the e-scape itself. Each of these actions will cause the launcher to perform a callback to the personal e-scape server which instantiated it, which can then act on the action (by fetching the properties of the new e-scape etc.).

Access control

The e-scape session architecture employs a relatively simple hierarchical access control scheme for data in the user profile server and the e-scape registry service where information in these servers is protected with an access control list. Each entry in the access list contains an e-scape or user identifier and the permissions they have.

Each user and e-scape in the session architecture is allocated a unique identifier by the user profile server / e-scape registry service which first registers it. This identifier is a combination of an increasing value (maintained locally at the server in question), the URL of the server which first registered the entity, and the URL of the current server the user / e-scape information can be found at. In addition to creating this identifier, the registering server also creates a password for the user / e-scape and a locally associated alias which can be used instead of the full identifier in conjunction with the servers URL. For example, Jonathan could be registered at the user profile server with an identifier of “56371_server/ups_server/ups”, with a password of “secret” and an alias of “jonathan”. Jonathan can then authenticate himself at any when accessing any instance of the three types of server by providing either his identifier and password, or his alias, password and location of his current user profile server.

Each entry in a folders access control list dictates what operations can be performed on information in the folder by a particular identity. The set of allowable operations which can be specified are:

-- Read (R) – the names and values of the name-value pairs can be read

-- Write (W) – the values of the name-value pairs can be modified

-- Create/Delete (C) – new name-value pairs can be created or removed

-- Access (A) – these access rights can be modified

An e-scape obtains its current password and identifier through its e-scape launcher, which in turn receives this information from the e-scape registry which executes the launch command on the users personal server.

Examples of the default access rights for new user profile and a new e-scape registration can be seen in Figure 11, where “8833_…” is the identifier for e-scape XYZ and “56371_…” is the identifier for Jonathan. Note that the creator of the e-scape registration, in this case Jonathan, is automatically inserted into the e-scape details access list. This allows the registering user to modify the details of the e-scape at a later time if necessary.

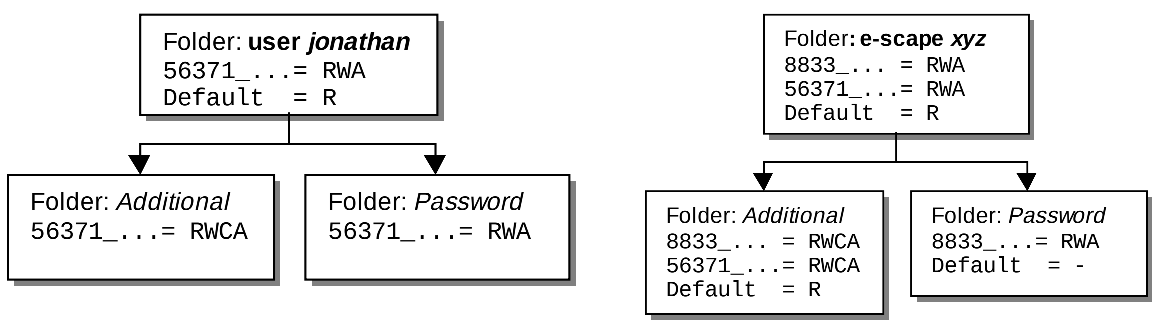

Figure 11 Default access rights for a newly created profile and e-scape

Top level folders have a special property which means that name-value pairs can never be created or deleted from them, i.e. the ‘C’ permission can never be set on a user. This stops users from deleting the main “system” allocated pairs, like the e-scape type. In addition, if Read access is not given on the contents of a folder then the folder itself is hidden.

Implementation

The e-scape session architecture has been implemented using Java servlet technology, which communicates through HTTP requests and responses. Each of the servers described above is a distinct servlet. Java servlets provide CGI-like functionality and can be embedded behind any appropriate Web server but offer several advantages over the more traditional CGI approach. First, servlets are loaded and embedded in the actual server itself. This means that each servlet is persistent and state is stored between requests. Secondly, each request does not require any forking of processes to deal with the request, making servlets very fast.

In addition, the combined use of Java, HTTP and HTML to support the architecture allows each server to run on a variety of operating systems and hardware, but also to be accessed from any computer with a standard web browser.

Finally, because Java inherently supports transportable code, its is relatively simple to provide for the notion of launchers which are transparently fetched and executed locally from remote servers - which allows the types of e-scape supported by the session architecture to grow and expand as required.

Example: Launching an e-scape

In order to demonstrate the responsibilities and inter-relation of the various e-scape services, consider the following example where a user locates and enters a Dive e-scape. If it not already running the user starts their personal e-scape server running on their machine. The user opens a web browser and enters the URL for a registry server. The registry server asks the user for the URL of their user profile before prompting for their username and password. Users do not need to keep providing the name of their profile server as once the information has been provided it is stored within a cookie, which can be passed automatically by the web browser the next time the user wants to log onto the registry server. The profile server then authenticates the identity of the user and then provides an overview of the registered e-scapes. In the default HTML display it provides a page listing the e-scapes registered on the server. The user refines the list by entering a search for only e-scapes of type “x-escape/x-dive” and clicks on one of the returned e-scape links to display the information about the e-scape (see

Figure 7 above).

The e-scape is then started by clicking on the “Launch” link at the bottom of the page. This executes a GET request to the URL of the user’s personal server, whose address will have previously been fetched by the registry server using the profile cookie. The GET request carries with it several pieces of information, the:

-- *name* of the e-scape, e.g. “Patio area”

-- URL of this registry server, e.g. http://server/ers

-- type of the e-scape, e.g. x-escape/x-dive

Upon receiving the request, the personal e-scape server checks its local launchers for any which can handle the Dive e-scape type. If more than one launcher could potentially deal with the e-scape then the user is offered a choice, and the process is repeated. In this example however, no launchers are able to start the e-scape. Therefore the personal server contacts the registry which sent this request and asks for any launchers that have been registered which can handle this type. Again, if more than one registered launcher could deal with the type the user would make a selection through an HTML form. Once a remote launcher has been identified, the Dive launcher, the personal server downloads the Java class to the local machine and fetches all of the e-scape information from the originating registry (the e-scape properties, location URL etc.). It then instantiates the launcher with this information and registers itself to receive the launcher callbacks relating to the status of the e-scape. In addition, the launcher is added to the personal server's list of currently executing launchers.

As this is the first time the Dive launcher has been instantiated locally on this machine, it doesn't know if or where Dive is installed. Therefore, the launcher returns an HTML page to the personal server requesting this information (Figure 10 above). Once the user has filled in the form the launcher records it locally to use the next time the user launches Dive e-scapes, and proceeds to start the e-scape. The final HTML reply by the launcher is actually an HTTP re-locate request which returns the users browser to the initial information screen on the registry server with the "Launch" button.

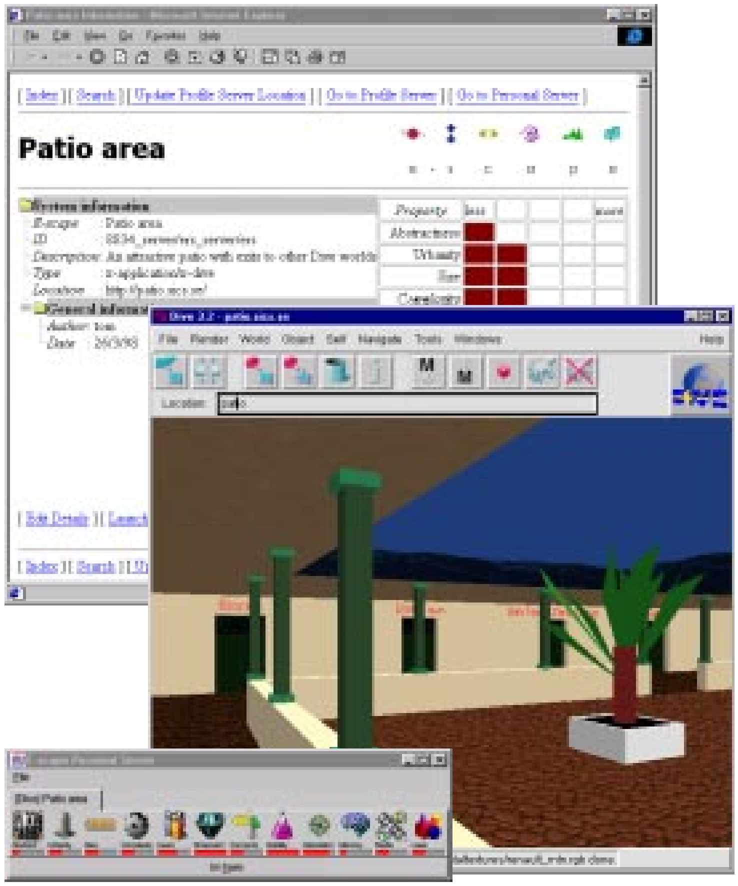

Figure 12 A Dive world with the properties display

The Dive launcher starts Dive by executing the Dive executable in a separate process with several different command line arguments which tell Dive where the world is it needs to load (the URL), the name of the current user and so on. Once the process executes, the "started" callback is executed on the personal server that displays the e-scape properties (figure 8) and contacts the originating registry server to say the user is currently in the e-scape. The launcher enters a wait loop for the Dive process terminating. The Dive application will then run with the user inside the e-scape. When the user finally quits Dive, the launcher instance will signal the "stopped" callback on the personal server. The server then removes the properties for the e-scape from the display and contacts the e-scape registry to remove the user from that e-scape.

In this simplified example, the Dive launcher provides the most rudimentary e-scape aware wrapping of the application it starts (Dive). This type of "process" wrapping is available in all cases, whether the e-scape is a VR system or a text editor. However, the real Dive launcher connects to the Dive Client Interface and monitors the users current Dive world. If the user exits the original e-scape to another Dive world, the Dive launcher can detect and signal this. If the new world is also a known e-scape on the registry server then the personal server can move the user appropriately and display a new set of property icons.

A broader consideration of e-scapes

Although this chapter has considered e-scapes as 3D collaborative virtual environments, e-scapes are by no means limited to that domain. Any application that has some sense of user engagement and which externalises meta-information about itself can be considered as an e-scape. The ability to externalise properties of the application allows our architecture to integrate a diverse set of applications particularly those that exploit some form of spatial metaphor. These include the use of rooms by Teamrooms (Roseman & Greenberg 1992) and locales by Worlds (Fitzpatrick 1996).

The architecture and protocol is intended to be sufficiently open to allow developers to integrate most CSCW applications. This builds upon the approach to session management suggested by Intermezzo (Edwards 1994) where properties of the application are separated from the policy issues controlling sessions. In our case the ubiquitous nature of the HTTP protocol set allows us to move beyond the language specific approach suggested by Intermezzo. In the following section we consider how an existing session based CSCW application that does not adopt a spatial approach could be integrated as an e-scape.

Integrating Existing Systems

Many existing shared user interface systems, such as GroupKit (Roseman & Greenberg), SOL (Smith & Rodden 1995), Rendezvous (Hill et al. 1994), MMM (Bier & Freeman 1991) and others, implement a session management environment to: control invocation of new sessions; provide a selection of possible applications to be launched; indicate the number of users in each session; and so on. For example, MMM provides a 'Home Area', which is a control panel, which links the user to the underlying system, which is external to any applications. This is similar to the *Access Point* of Rendezvous, the *Registrar* (and *Registrar Client*) of GroupKit, and the *User Panel* of SOL. The underlying architecture of each of these session management systems operates in a similar manner.

Each of the collaborative tools supported by each of the CSCW systems mentioned above is only supported by the registration service of that system. For example, to use the GroupKit shared drawing tool alongside the network management tool of SOL, each user must use both GroupKit's Registrar *and* SOL's User Panel to manage each of the two tools independently. If a number of shared applications from a range of different CSCW systems are used simultaneously, the number of session management tools apparent in a user's display can become distracting and tedious. A preferable solution in this scenario is to use a unified registration service. This would provide users with a single range of possible applications, independently of their underlying implementation details. It would also provide a single interface presenting the sessions a user is currently participating in, and the other participants of a session. These existing session management systems can be replaced by the e-scape session management system, to provide such a solution. For example, consider the case of GroupKit. Its architecture is depicted in Figure 13.

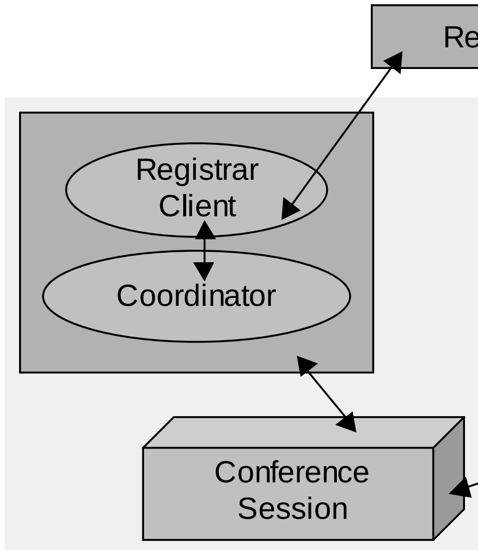

Figure 13 The GroupKit Architecture

The Registrar Client and Coordinator provide the user with the GroupKit session management user interface. Upon start-up the Registrar Client is provided with the location of the Registrar itself. It is the Registrar Client/Coordinator process that enables users to view current conference sessions (such as the collaborative drawing tool) and their participants, or initiate new sessions of a tool. The Registrar is simply a common process that coordinates the information held by the registrar clients. Additionally, conference sessions communicate directly with each other, to maintain consistent application specific state. Hence, a normal configuration is that a user's display will contain of a single Registrar-Client/Coordinator window and one or more Conference session windows.

The entire Registrar and Coordinator service can be replaced with the e-scape session architecture. Here users would register with the e-scape server as discussed earlier, then access the e-scape Registry server through a web browser. All e-scape sessions will then be available to that user, including GroupKit applications. On selecting a GroupKit application, their e-scape personal server will identify that application as being of type "x-escape/x-groupkit", and execute the e-scape launcher for that class of application. As with the Dive example above, the launcher updates the participant list for the new session by adding the new user to it. Other users wishing to join this GroupKit application may locate and join it using the e-scape Registry server. Here the active session is displayed, along with the current participants. Selecting 'launch' will invoke a new GroupKit session for them. The e-scape launcher for that class will provide the correct parameters so that it can communicate with other user's applications in that session.

In essence, the use of the e-scape session architecture provides all individual collaborative applications with a unified session manager. However, in addition to providing a common session manager, use of the enhanced e-scape launchers to start the GroupKit tools enables extra meta-information to be gathered from each of these applications. This information can then be relayed and externalised at the e-scape Registry server. For example, if the GroupKit launcher was able to communicate with the GroupKit application in some way, information such as ‘idle time’ or the current type of tool being used (pen, eraser, rectangle etc.) for each user may be copied to the e-scape Registry server, which could then display it. In addition to allowing this information to be browsed, supplementary more powerful services can easily access and process this information through simple HTTP commands.

Future Work

The basic session architecture supports the migration of code to provide appropriate functionality required for launching different types of e-scape. This can be extended to promote the personal e-scape server as a framework for “plug-in services” to be downloaded and incorporated. Because these plug-ins are connected to the personal server they automatically gain access to e-scape session data. Such information can be exploited to enable cooperative services, the simplest of which would be chat facilities between the personal-servers.

Since each launcher makes any application e-scape aware, they can provide information concerning the implicit connections between e-scapes. For example, if a user moves through several Dive worlds, connected via Dive portals, the Dive launcher can record the paths and connections traversed by a user. This information can be used later to provide an overview of the connections between different e-scapes, or a history of a users progress.

More generally, the session architecture is extendable to allow a more diverse set of usage information to be kept. The use of simple name-value pairs means that any application specific information can be made externally available and maintained between sessions. This allows us to support a range of different temporal operations. What these are and how they might be supported remains a matter of some further investigation based on the usage of the architecture.

Conclusion

In this chapter we have further defined the notion of an “e-scape”. An e-scape is any electronic world where sets of properties are made public that provide an indication of the nature of the environment. This publicity allows the development of applications and virtual environments that join these virtual environments together.

A session management system was presented that aims to support users accessing a range of different virtual environments or e-scapes. To promote scalability and accessibility the developed architecture is built upon the HTTP protocol. This allows us to realise a simple yet robust architecture that allows people to encode meta-information about environments, virtual or otherwise. This arrangement provides support for a heterogeneous set of connected environments through the notion of environment launchers.

The architecture makes a clear separation between the management of these environments and the users' connections with these environments, and the activities of users within them. Such activities and interactions of users are the responsibility of the environment itself. The architecture focuses on the issues of support and management needed to move these environments from isolated research prototypes to accessible everyday systems.

References

Alphaworlds, http://www.alphaworlds.com

Benford, S., Bowers, J., Fahlen, L., Mariani, J, Rodden. T, Supporting Cooperative Work in Virtual Environments. *The Computer Journal*, 1995. 38(1).

Bier, E., Freeman, S., MMM: A User Interface Architecture for shared editors on a single screen, *Symposium on UI software and Technology*, (Hilton Head, South Carolina, 1991), ACM Press, 79-86

Blaxxun, http://www.blacksun.com/

Bowers, J., Pycock, J., O'Brien, J., Talk and Embodiment in Collaborative Virtual Environments, in *Proc. CHI'96*, ACM Press

Brodbeck D., Chalmers M., Lunzer A., Cotture P.: Domesticating Bead: Adapting an Information Visualisation System to a Financial Institution. in *Proc. Information Visualization '97 (*Phoenix, AZ, October 20-21, 1997), 73-80.

Carlsson, C., Hagsand, O. DIVE: A Multi User Virtual Reality System, in *Proc. IEEE VRAIS*, (September 1993), 394-400

Curtis, P., Dixon, M, MUDS grow up: Social Virtual Reality in the Real World, *in Proc. of IEEE Computer Conference*, (January 1994).

Dix, A., Challenges for Cooperative Work on the Web: An Analytical Approach, *CSCW: The Journal of Collaborative Computing*, 6, 2-3, 1997, 135-156.

Edwards, K. Session Management for Collaborative Applications, in *Proc. CSCW'94* (Chapel Hill, North Carolina, October 22-26, 1994). ACM press. 323-330

Elvins, T., Nadeau, D., Schul, R., Kirsh, D. Worldlets: 3D Thumbnails for 3D Browsing, to appear in *Proc. CHI'98* (Los Angeles, CA, April 18-23, 1998), ACM Press.

Fitzpatrick, G., et al, Physical Spaces, Virtual Places and Social Worlds: A study of work in the virtual, in *Proc. CSCW'9*6 (Boston, MA, Nov.16-20, 1996), ACM Press. 334-343.

Gaver W., The Affordances of Media Spaces for Collaboration, in *Proc. CSCW'92* (Toronto, Canada, Oct.31-Nov.4, 1992), ACM Press. 17-24.

Greenhalgh, C., Benford, S. MASSIVE: A Virtual Reality System for Tele-conferencing, to appear in *ACM Trans. on Computer Human Interfaces*

Henderson, A.J., and Card, S.A., Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention, *ACM Transactions on Graphics*, Vol. 5, No. 3, (July 1985), ACM Press.

Hill, R., Brink, T., Rohall S., Patterson, J., Wilner, W., The Rendezvous architecture and language for constructing multiuser applications, *ACM Trans. on Computer-human Interaction*, Vol. 1 No. 2, (June 1994), ACM Press, 83-125

Ishii, H., Ullmer, B., "Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms," in *Proc. CHI '97*, ACM press, 234-241

Living Worlds, http://www.livingworlds.com

Patterson, J.F et al., Notification Servers for Synchronous Groupware, in *Proc. CSCW'96* (Boston, MA, Nov.16-20, 1996), ACM Press. 122-129

Root, R.W., Design of a Multi-Media Vehicle for Social Browsing, in *Proc. CSCW'88* (Portland, Oregon, September 26-28, 1988), ACM Press. 25-38

Roseman, M., Greenberg, S, TeamRooms: Network Places for Collaboration, in *Proc. CSCW'96* (Boston, MA, Nov.16-20, 1996), ACM Press. 325-333

Roseman, M., Greenberg, S., GroupKit. A groupware Toolkit for building Real-Time Conference Applications, in *Proc. CSCW'92* (Toronto, Canada, 1992), ACM Press, 43-50

Smith G., Rodden, T. SOL: a toolkit for cooperative interfaces, *International Journal of Human Computer Studies (IJHCS)*. No 42. 1995. Academic Press Limited. 207-234

The Contact Consortium, http://www.ccon.org/

The Palace, http://www.thepalace.com

VRML97, International Standard ISO/IEC 14772-1:1997, "the Virtual Reality Modeling Language"

Chapter Seven: Open Support for shared spaces based on e-scapes

Jonathan Trevor, Kevin Palfreyman, Tom Rodden

Lancaster University

Introduction

This chapter presents an infrastructure to support a range of real time cooperative environments by allowing information to be shared across environments. The infrastructure builds upon the use of shared common spaces by using a distributed tuple space to provide information sharing at its base level. The use of a tuple space moves away from previous models of distribution in cooperative systems that focus on the propagation of events. The chapter presents the infrastructure and shows how it can be used to support information sharing across a number of different forms of cooperative system.

The CSCW research community has seen the development of a number of systems to support real time cooperation. These systems have their roots in shared drawing and shared screen systems. The issues surrounding the management of these systems have focused on a consideration of the management of these interfaces and the development of interaction techniques such as WYSWIS (Foster & Stefik 1986) and interface coupling (Dewan & Choudhary 1991).

These shared interface systems initially focused on the development of techniques to support a small number of users. For example, the development of the ClearBoard system (Ishii & Kobayashi 1992) considered the issues surrounding the use of shared interfaces by two users. Similarly work on shared drawing (Tang 1991) focused on a small number of users.

The application of shared interface systems in real world settings saw some re-examination of the basic principles of sharing. Perhaps the most notable of these is the re-examination of WYSIWIS (Foster & Stefik 1986) and the subsequent development of a number of techniques to manage the separation between public and private interfaces. These techniques included the development of coupling techniques (Dewan & Choudhary 1991) and the construction of techniques to support the tailoring of cooperative user interfaces (Greenberg 1991).

As technology matured and the development of systems that allow real time shared interaction became more stable researchers considered the development of a variety of different applications that supported this form of interaction. The need to allow the rapid development and construction of cooperative applications saw a growing shift toward the development of toolkits. These included the development of Groupkit (Roseman & Greenberg 1992), the modification of Conversation Builder (Kaplan et al. 1992) and the emergence of Rendezvous (Patterson et al. 1990).

Underpinning the development of these shared interfaces was a growing acceptance of a common model. The general strategy agreed was to consider multi user interfaces in terms of different views on a pool of shared data (Bentley et al. 1992, Dewan & Choudhary 1991, Hill et al. 1993, Lee et al. 1996, Mansfield et al. 1997). Work continued to focus on the development of appropriate models of distribution and how the issues of shared information may best be supported. This focus led to a general consideration of the architecture of shared interface systems and how best to provide responsive interaction while sharing information between remote users (Roseman & Greenberg 1997).

The development of the World Wide Web and the subsequent establishment of a large user base for distributed multi-user applications saw some shift in focus in shared interface systems and CSCW systems more generally. Rather than consider the need to develop interfaces for a small number of users the growth of the Internet allowed the rapid deployment of systems within organisations (Dix 1997). As a result researchers needed to consider the issues of developing systems that supported a large number of users and provided facilities to allow users to manage this interaction.

The need to consider the management of interaction has tended to focus on the development of techniques to support the process of users coming together and techniques to allow different contexts of interaction to be supported. Many real time cooperative systems have turned to some form on spatial metaphor to support this.

The Emergence of Shared Spaces

The need to provide techniques that allow users to manage multiple simultaneous interactions has seen the rapid development of spatial approaches to CSCW. Nearly all real time cooperative systems exploit some form of spatial metaphor as a means of presenting the system to users. As was the case in the initial development of the Rooms metaphor (Henderson & Card 1985) the use of a spatial technique is often aimed at providing a set of simple easy to use techniques to manage interaction.

The application of spatial techniques and arrangements is also a focus for study for those interested in MediaSpaces (Gaver 1992) and has seen the development of a range of different physical shared spaces augmented with multimedia connections. These spaces have tended to focus on the physical arrangement of spaces and the inclusion of different media within these spaces. However, in doing so they have also often exploited spatial metaphors as a means of controlling the media. Perhaps the most notable example of this is the

Cruiser system (Root 1998). A similar use of a spatial metaphor can be seen in the Portholes system (Dourish & Bly 1992) with its structuring of space.

The use of a shared space has also been exploited in MUDs and MOOs (Dourish 1998) to support a disparate community of users. These simple text based environments often exploit some form of spatial metaphor to support communities of users. The general model of interaction is to allow a number of users to enter some shared "world" and allow their interaction to be seen by other users. These systems have recently been extended to provide 3D interfaces to these virtual worlds and a number of systems exist that support these virtual worlds (The Contact Consortium).

The development of shared virtual environments within the CSCW community has seen the development of a number of cooperative virtual environments. These systems place users within a 3D virtual environment and provide facilities that allow them to navigate these environments and interact with each other (Carlsson & Hagsand 1993, Greenhalgh & Benford forthcoming). These environments have been used to support various forms of teleconferencing and cooperative applications such as the cooperative browsing of information (Benford & Mariani 1994).

The emergence of 3D virtual environments has been balanced by a growing application of spatial techniques within 2D shared interfaces. Many of these are based on the notion of shared workspaces and the partitioning of these workspaces. These include the ORBIT system (Mansfield et al. 1997) with its use of locales. The use of rooms within the Teamwave system (Roseman & Greenberg 1992). The use of workspaces within CBE (Lee et al. 1996). A spatial approach is also exploited within the Piazza system (Isaacs et al. 1996) with its focus on users sharing information as a means of supporting work.

The need for common support

It is clear from this that some common agreement on the development of applications to support real time cooperative interaction has started to emerge. The core of this agreement is an exploitation of the concept of some form of shared space that provides a focus for interaction.

This shared space provides a number of significant features that are provided to users to allow them to manage the interaction:

- Activities can be partitioned in a number of distinct shared spaces (often the use of rooms, worlds or locales are used to describe these spaces).
- The shared spaces allow resources associated with the cooperative interaction to be gathered together.
- Information can be shared between users by placing them within shared spaces.

- Users can be considered present within these shared spaces when they access the information and resources within them and this presence is reflected to others.

Essentially the majority of shared interface systems agree on the need to allow resources and information to be shared, the provision of techniques to structure this sharing and the need to allow users to be aware of the activities of others. However, despite this broad agreement the majority of systems do not provide any mechanism for interoperability with other cooperative systems. In essence, each of these cooperative systems behave as closed worlds with limited access to other forms of cooperative systems.

One reason for the lack of cooperation between systems is that a mismatch exists between the need for sharing within these systems and the provision of supporting distributed infrastructures. The current generation of cooperative support systems focuses on the propagation of events to promote awareness. While the notion of awareness is important we feel that this focus has led to imbalance within the provision of supporting services and we need to consider the development of sharing services that allow a variety of different forms of heterogeneous information to be shared between these cooperative systems.

The Development of Supporting Services

The majority of supporting services and protocols have focused on the propagation of awareness information in the form of events. This has included the development of protocols to augment the World Wide Web (Palfreyman & Rodden 1996) and a number of general event mechanisms.

Although they focus on propagation of awareness some of these protocols already incorporate some model of shared space. For example, the Corona communication service incorporates a model of shared spaces (Hall et al. 1996). A similar model of shared spaces is built into the services supporting Teamwave (Roseman & Greenberg 1997). Like many other applications both of these approaches focus on the configuration of applications from a number of components that each provide a common shared API. This component based architecture mirrors the general architectural arrangement of applets supported by a common protocol found in Internet development.

The use of space is also mirrored within the development of the NSTP protocol that uses places and things (Patterson et al. 1996). The use of places within the NSTP allows a generalisation away from the concept of session as a means of supporting shared interaction. This use of place exploits shared state as a means of allowing the interaction to be managed. The notification server maintains the state of things placed within the places on the server. The NTSP protocol explicitly excludes support for persistence and streamed data (e.g video and audio..)

In providing support for state sharing these supporting services each offer some form of model of sharing. The supporting services in Teamwave promote rooms, while Corona offers shared spaces and ORBIT offers locales. Although not linked to a particular systems NTSP provides places as structuring devices. While each of these concepts are very similar the services are committed to supporting a particular model of shared spaces.

Rather than develop another model of shared spaces and support for this model we wish to provide support to bring these different spaces together. In order to do so we focus on the provision of a very general model of shared state that allows users to make information available to others and to allow linkages between different forms of shared state. The service we have developed aims to provide a number of key features

-- A lightweight data model that allows support for a number of different forms of shared space

-- Extensible support for sharing allowing the addition of different persistence mechanisms.

-- Support for a wide variety of forms of data (including support for multimedia streamed data)

--Support for flexible forms of sharing with varied patterns of use.

In order to provide these features we have moved away from the traditional communication based model of distribution. This move reflects the changing demands from cooperative applications as they have moved from previous session based models of control to focus on the use of shared spaces.

The distribution model we have exploited centres on the provision of a distributed tuple space. The tuple space provides distributed access to name value tuples and manages the issues of distribution to ensure that changes to these tuples are available across the community of users.

The tuple space concept

A *tuple space* is a well known distributed systems mechanism originating from the work on the Linda project (Gelernter). Tuple spaces are traditionally used for (and best suited to) supporting systems with producer-consumer flows of data - where one part of the system is responsible for providing data that another processes.

A tuple space contains zero or more *tuples*, as well as references to other tuple spaces. Each tuple consists of an ordered list of typed fields. For example, a tuple that would hold a name-value pair would be (<*String*>, <*Data*>), where <> identifies the type of the field. Note that fields are not named, making the ordering important. Each field in a tuple is also said to be an "*actual*" or a "*formal*". An "actual" field is one that already contains real data, whereas a "formal" field is a placeholder containing no data and is used in matching (see below).

Traditional tuple spaces support three key operations:

RD (query tuple)

Read any tuple from the tuple space that matches certain fields in the query tuple. Matching is performed by checking the “actual” fields have the same types and values as the tuple in the space. Once all the fields have been matched, the remaining (if any) formal fields in the tuple are filled with the corresponding values from the matching tuple. Importantly, the operation is not deterministic and when more than one match is possible, there is no way to guarantee you have read all matching tuples using consecutive RD operations. Indeed, a hundred RD operations may match the same tuple again and again even if there are a thousand possible matches. In addition, the operation is blocking so that the execution of the program making the RD call will stop until a tuple successfully matches the RD.

IN (query tuple)

This operation is similar to RD but removes the matched tuple from the space. Again, there are no guarantees as to which tuple will be consumed if more than one would match the query tuple. The operation is also blocking.

OUT (tuple)

Write the tuple <*tuple*> from this application out to the tuple space. Once the tuple has been written it becomes available to all applications connected to the space. All the fields in a tuple being OUTed must be “actuals”.

The tuple space we have used, Limbo2 (Davies et al. 1997), requires that the tuples be typed in a type hierarchy. Thus, applications must first export **Type** tuples to the space in order to define which fields each particular type of tuple will contain. The implementation of a distributed tuple space is described in detail elsewhere (Davies et al. 1997) and relies on multicast groups as the underlying transport mechanism.

Tuple spaces provide us with two distinct advantages in terms of an infrastructure for real-time cooperative applications.

-- The mechanism is based on fast access to *decentralised* shared state and the infrastructure is optimised to support this arrangement. Applications sharing data with the space should not have to worry about how changes to their shared state are propagated to other applications.

-- The use of tuples as a representation of information does not impose a particular data model and applications can externalise diverse forms of structured information.

The shared universe platform

The Shared Universe platform views applications as isolated “worlds” populated with information and data. These applications may choose to share some or all of this data with the rest of the “universe”, so that other applications may subsequently access this information, regardless of machine, operating system or network. The Shared Universe Platform manages and provides support for applications to insert and extract information from a universe of shared data, to allow the population of many worlds to communicate.

What makes this approach novel and advantageous over previous approaches is that the shared space metaphor is being modelled at the underlying distributed system level in a completely semantic-less manner. This is important for several reasons. First, it enables the programmer to reflect changes in the applications model of cooperation within the underlying communications structure of the distributed system (Bentley & Dourish 1995). Secondly, by moving the data sharing into the distributed system, the natural boundaries of sharing between applications can be transcended, allowing the data to be accessed and shared across applications, operating system and network.

The Universe architecture is divided into three separate layers, as shown in Figure 14 below.

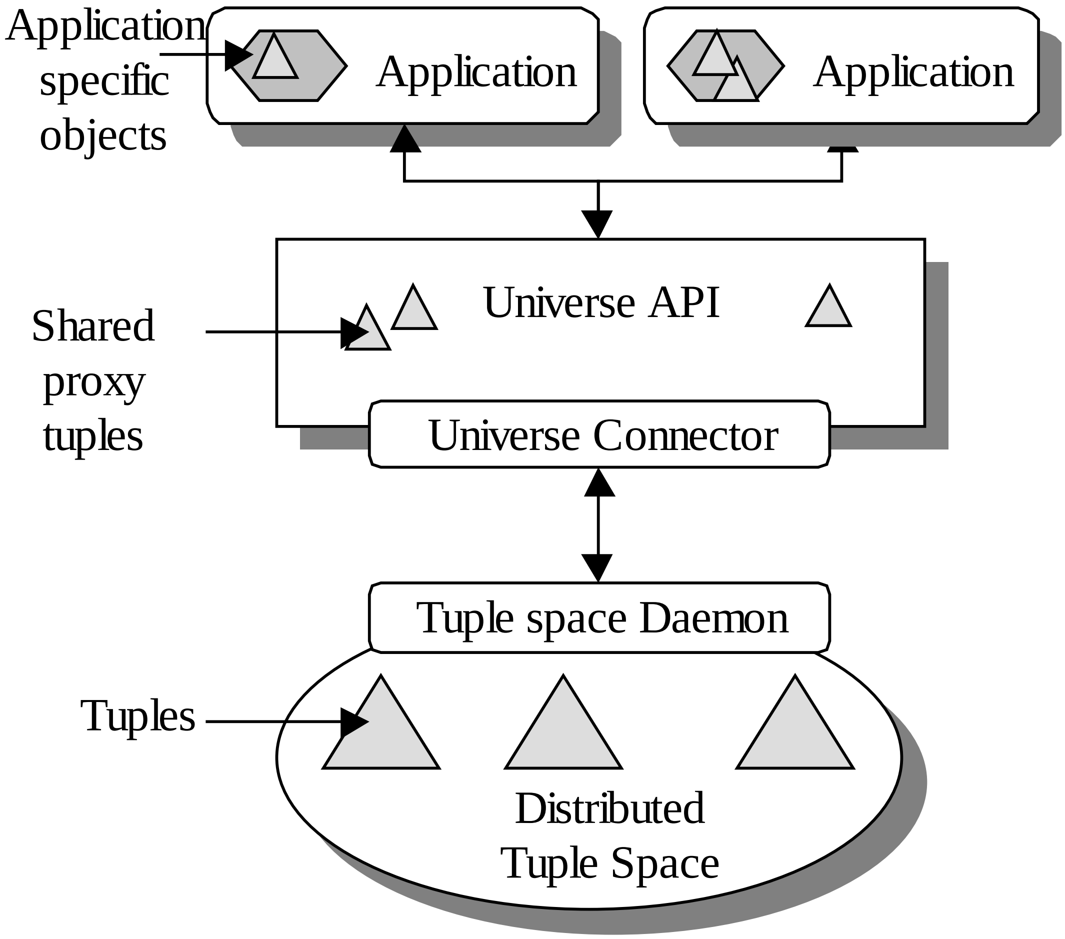

Figure 14 Shared Universe Architecture

The features of the architecture can be best described in four separate pieces:

I. The *application layer* contains the applications themselves. The application contains *application specific objects* that are native to that application. The application decides which of these objects are exported to the universe of shared information. The application communicates with the shared universe through the Universe API.

II. The *Universe API* maintains and decomposes the application specific data into *shared proxy tuples*. These are objects that reflect the contents of tuples in the underlying *tuple space*. They change as their tuple space counterparts change and v.v.

III. The *universe connector* provides the main coordination point for applications and the Universe API, facilitating access to information in the universe. This includes tasks such as adding and removing shared state from the tuple space and monitoring the tuple space for changes.

IV. The *tuple space deamon* provides access to the tuple space and allows applications (such as the universe connector) to add, query and remove data from the tuple space, as well as maintaining the same tuple spaces in a replicated fashion across any number of machines.

Extending traditional tuple spaces

Although the tuple space metaphor provides natural and advantageous notions of a shared space which map directly onto the underlying communication used to maintain it, the traditional view of a tuple space is limited to the producer-consumer model of information flow. That is to say information is provided by one or more sources and put into the tuple space. This information is read or consumed some time later by one or more consumers. This model is problematic because the consumers are *expecting* the information to be produced.

The stability of this arrangement has led to implementations of tuple spaces as *passive* entities. No facilities are provided for actively notifying applications of events occurring in a tuple space (such as a new tuple arriving) and only very limited facilities are available for querying, rather than consuming or extracting, the current contents of the tuple space.

To promote the more dynamic arrangements involved in cooperative settings we extended the Limbo2 tuple space with a simple awareness facility offering two operations:

REGISTER (callback, ARRIVED/DELETED, tuple-type)

This operation registers a callback <*callback*> in the current application which will be executed whenever a tuple of a particular type <tuple type> (or subtypes) is either added to the tuple space (OUTed), or removed (INed), depending on the passed flag. The callback itself receives the identifier of the tuple space in which the operation occurred, as well as a copy of the actual tuple the operation occurred on (the one that was added or deleted).

UNREGISTER (registration-id)

Remove a previously registered callback. In addition, several other extensions to the tuple space operations have been added:

- copy_collect; rd all tuples matching a pattern to a new space
- collect; in all tuples matching a pattern to a new space
- inp; non-blocking version of in
- rdp non-blocking version of rd

These extensions assist in transforming the traditional tuple space to make it a fast and simple distributed shared space. Each tuple can also take an explicit “expires” value which determines how long the tuple can exist in the tuple space before being automatically removed. This feature is necessary to support the transient nature of events sent between applications

The universe API

The Universe API provides a simple common interface to allow applications to share and manipulate data through the tuple space. The API is written in Java to provide as much cross-platform support and integration as possible. The API defines two types of tuple that are exported into the shared tuple space and carry application data*: data tuples* and *relation tuples*.

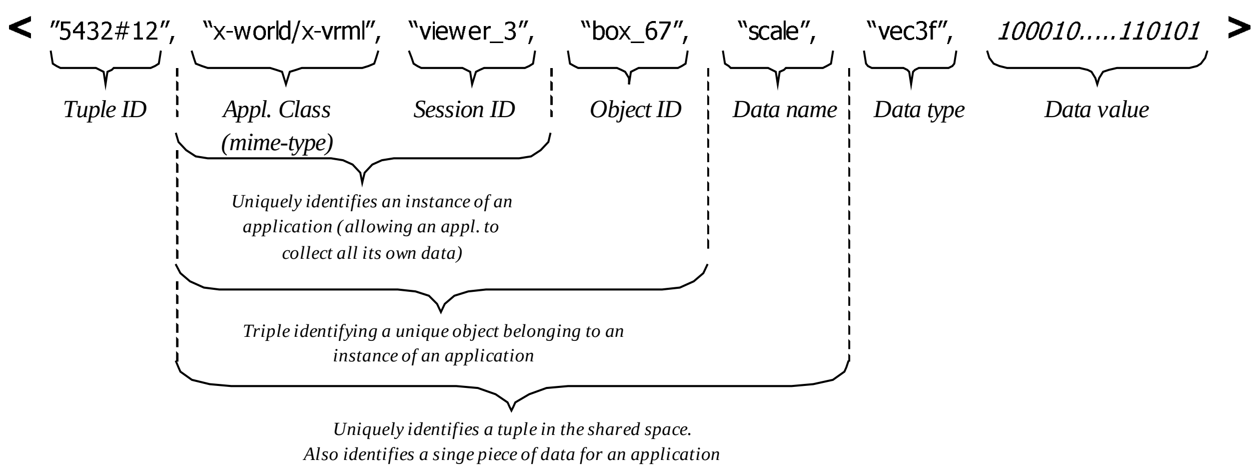

Figure 15 Composition of an example Data tuple

Data tuples

The Data tuple represents a single atomic piece of shared data. A Data tuple consists of the fields shown in Table 1.

Field	Type	Description
Tuple ID	String	Some unique identifier that identifies this particular tuple in the tuple space. It contains a tuple sequence identifier.
Application	String	Identifier of the application sharing the shared object (typically its mime-type)
Session ID	String	Some unique identifier for this session.
Object ID	String	Some unique identifier for the shared object to which this data belongs too.
Data name	String	The name of the data value
Data Type	String	The type of the data value
Data value	Data	The data itself

Table 1 Fields in a Data tuple

In order to extract tuples relevant to a particular application from the shared space, we must be able to identify distinct pieces of application shared data. Aside from the *TupleID* field, an application may uniquely identify a particular piece of data using the following quad:

{Application, Session ID, Object ID, Data name}

The uniqueness or scope of each identifying field in the quad is the responsibility of the previous field. Thus, each application with a given application class is responsible for maintaining and allocating the different session ID's in that application. Each session ID allocates the object ID's for objects in that session and so on.

The *Application* field and the *Session-ID* field are given equal weighting. This means that the platform supports both session-based and room-based views of the data, even simultaneously. If an application were to query the tuple space using a specified *Application* class (resource), then it would be possible to present to the user a view of application sessions. In contrast, the *Session-ID* can be associated with a particular room and if the space is queried using a specified *Session-ID* it would present a view of all of the resources available to that particular room.

The value for the *Application* field of a data tuple is specific for each class of application and would normally be compiled into the application code (or its external interface). For any given application class, there are a number of *Object-ID* (and their nested arbitrary *Data-name*) values reserved for management functions. This enables an application to display standard properties for each item in a list of sessions or resources.

Once data is made available in the tuple space, it can be retrieved using the normal tuple space operations. By combining different actual and formal values in a data tuple, applications can execute template queries for different data. Some useful combinations are shown in Figure 2. For example, by providing an application class of "x-world/x-vrml" and leaving all other fields as formals in a data tuple, the application can obtain a list of all session ID's for that type of application in the universe.

Data tuples can be cached as local proxy objects within the Universe API which are maintained by the universe connector.

Relation tuples

Arbitrary M-N relations can also be shared within the universe by using multiple 1-1 relationship tuples, whose definition is shown in Table 2.

Field	Type	Description
Application	String	Identifier of the application defining this relationship (typically its mime-type)
Session ID	String	Some unique identifier for this session.
Object ID	String	Unique identifier for the shared object which this relationship concerns.
Relation	String	Name of the relationship this 1-1 mapping belongs too.
Left	String	The LHS of the relationship
Right	String	The RHS of the relationship

Table 2 Fields in a Relation tuple

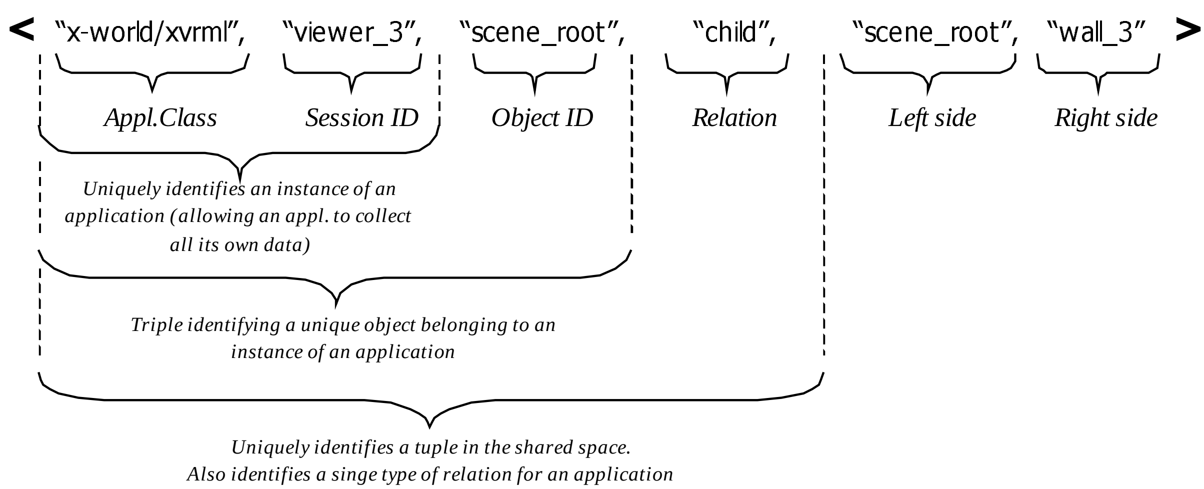

Figure 16 Composition of an example Relation tuple

A particular instance of a relation is therefore defined by the quad

{Application, Session ID Object ID, Relation}.

An entire M-N relationship is broken down by the API into 'M x N' 1-1 tuples, to allow queries to be executed on the universe over the shared relationships. To extract an entire shared relationship, the API need only perform a single "collect" operation on the tuples with the complete quad.

An example use of a relation triple would be to store a hierarchy such as those used in 3D geometries. The geometry would typically be broken into several smaller simple parts and grouping nodes used to tie them together. To share this state in the space we need data tuples for each part of the geometry and a data tuple for the group node. We then need to create a relation tuple to link the parent to each of its children. An example of this style 1-N use of relation tuples can be seen in Figure 16 above.

The Universe Connector

The Universe Connector object provides the point of communication between the Universe API and the underlying tuple space daemon, manifesting changes to the local shared data proxies as the actual values are modified in the tuple space, and visa versa. The connector provides methods that can be invoked through the API to query and manipulate the shared universe.

Two important methods are *shareData* and *shareRelation* which instruct the connector to share some application-defined data tuple or MxN relationship with the space. If the universe contains a matching data tuple or set of 1x1 relation tuples (for the MxN relationship) then the connector makes local proxies for the matching tuples and returns these values to the application. If no such data or relation exists, then new tuples are put into the universe and local proxies are created. Using these calls it is a trivial matter for an application to either initially share new data with the universe or to get existing shared values back.

State update and propagation

Propagation of state and notification of state changes between the underlying tuple space and the application is performed in two distinct ways depending on the nature of the update. Consider a telepointer object whose position is shared in the tuple space. Another application is displaying the co-ordinates of the point in a simple text field. If the pointer moves quickly from point A to point B several hundred potential state changes occur. However, the most important changes are the starting position of the pointer and the final resting position of the pointer. Those in the middle tend to be less important. In essence, we can categorise state updates to data tuples as either **transient** (if we miss an update it doesn't matter) or **persistent** (everyone sharing this state should see the same value).

Persistent updates

Because persistent events change the fundamental state of the data in the universe, no explicit event tuple is created for these updates. Rather, the sender (the source) of the event (typically the local proxy) IN's the state to be changed, modifies it, and then re-OUTs it back to the space. Two fields of the tuple are modified: the new state value itself and the tuple ID. The sequence part of the tuple ID is simply incremented and put back.

This change to the tuple space is then distributed to the other copies of the tuple space, where the daemon will trigger all the registered callbacks in the universe connector for the recently OUTed tuple. The universe connector then extracts the new state from the tuple and updates the local proxy for the data in question. Any objects, such as the applications, who have registered callbacks on the shared proxy subsequently get called. Note there are two levels of callback, first from the tuple space daemon to the proxies via the universe connector, and secondly from those proxies to any interested applications.

Transient events

Transient events are carried by special event tuples in a specific event tuple space. The format is shown in Table 3.

Field	Type	Description
Tuple ID	String	Identifies the data tuple containing the same Tuple ID on which this transient event applies.
Event Sequence#	Int	Current sequence number for this transient update
Data name	String	The name of the data value
Data value	Data	The data itself

Table 3 Fields in a Transient Event tuple

The tuple ID identifies the data tuple over which this transient update is occurring while the event sequence number provides ordering between each transient event. This is maintained by each universe connector which updates its last known value whenever a new transient event is added to the event tuple space or arrives through a callback.

To generate a new transient event, the source of the event (typically the proxy of the shared data tuple) creates a new transient event tuple (incrementing the local sequence number and using the last known tuple ID of the state being changed) and adds it to the tuple space. This tuple is subsequently propagated by the underlying tuple space system to all connected machines, which then issue callbacks to connected applications via the universe connector. Each of these universe connectors forwards the change in state to the relevant universe object and records the last event sequence number of the tuple.

Transient events are never IN'ed by the Universe API, only RD. This means that no negiotation has to occur in the distributed system for the ownership of the tuple, which means using transient events to send information is very fast. However, because they are never IN'ed transient events are never explicitly removed from the tuple space by the Universe API or applications. Therefore, each transient event has an *expires* field to enable the tuple space to remove old events during garbage collection.

Update Consistency

Consistency between the persistent events and transient events on the same data tuple is guaranteed because of the incremented tuple ID sequence value on the shared state whenever a persistent change is OUTed to the space. This means if the transient events "tuple ID" starts referring to a data tuple with a tuple ID which no longer exists, or whose sequence number is less than the last known one, these events can be safely discarded as they are “behind” some other part of the system (and are transient anyway).

Consistency between the transient events themselves is fairly loose as the sequence ordering is only guaranteed on each machine, rather than across the whole tuple space. This is because no IN’s are ever needed for transient events and consequently no serialisation is ever performed at the tuple space level. The advantage of doing RDs rather than IN’s is a vast improvement in performance as no checking of ownership exchange needs to occur in the tuple space. Consequently, the sequence numbering of two contradicting “actions” in a distributed VR system could be the same, since the sequence number is added by the local universe connector, and the object may get wildly different transient state updates in quick succession. However, the first machine to commit a persistent value for the data tuple to the tuple space will “win” as it will change the actual tuple ID these transient events are applying to (and persistent updates are serialised because they rely on INs).

Examples of the shared universe platform

To demonstrate various aspects of the platform, we present some examples of the platform in use chosen to illustrate some of the key features of the platform.

Extending shared space applications

One of the principle aims of the developed platform is to provide facilities to allow shared space based applications to extend the use of shared spaces beyond application boundaries. This means that information shared in these applications can be made available externally to the application and used by other cooperative systems that exploit shared space-based approaches. To demonstrate this we extend an existing single user VR system to allow the information in the world model to be shared across multiple users versions of the application.

There are three different stages involved in extending a single-user VR system to make it multi-user. In this example, we consider making a VR application, such as one supporting VRML scenes, multi-user. The original VR application is considered to be a “black-box” that provides a rich API to access its state. We need an additional universe-aware application to mediate the VR applications access with the Universe API.

The initial stage of allowing the application to access and share data with the universe is to allocate it a unique application class identifier. In this case we assign the class identifier, “x-application/x-shared-vr”.

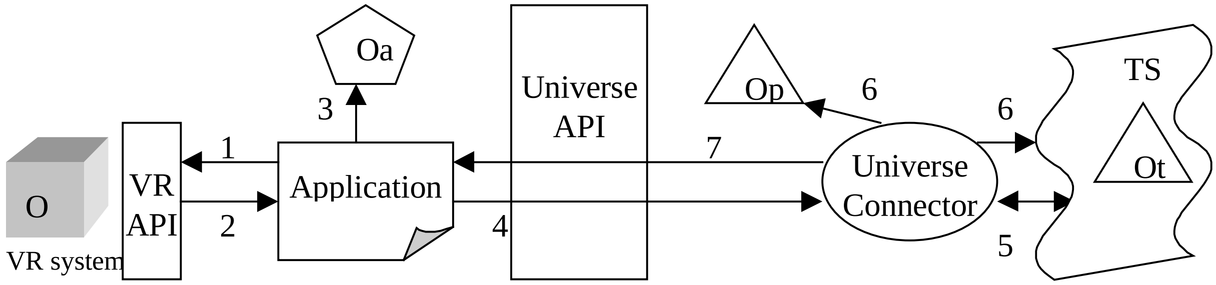

Figure 4 Exporting a VR object to the universe

The second stage of making the application multi-user requires that the actual VR scene data is made available to other instances of the application. This involves taking the data contained in the scene and sharing it with the universe. Figure 4 shows the steps involved in sharing a single object with the universe. First, the application requests VR system specific scene description through the VR API (such as the EAI[1]) (1). The application proceeds to create any external representations of the scene objects (in this case only Oa is created) which will consist of the objects' geometry, scale, rotation, translation etc. (3). Once created the application performs a "shareData" call on the connector, passing its application class "x-application/x-shared-vr", a default session identifier "session1", the name of the object "Oa" and initial value of the data it wants to share (4). The connector tries to locate a current tuple matching these (5). This will fail and the connector creates a new local proxy for the data (Op) and OUTs the new data tuple to the tuple space (Ot) (6). A reference to the local proxy Op is then returned to application by the connector (7). The connector also registers itself and the application as objects to be notified whenever the proxy (Op) changes. Therefore whenever the local proxy is updated by either the connector (in response to a change in the universe data tuple) or the application (in response to a change in the local VR world) the application and connector will receive notification of the update. Once notified the change can be manifest in the tuple space or VR world.

The third stage in making the application multi-user is to allow the exported shared data for the VR world to be imported into another application. Figure 5 shows the steps involved in obtaining an existing shared object from the universe.

---

[1] VRML External Authoring Interface, http://vrml.sgi.com/moving-worlds/spec/ExternalInterface.html.

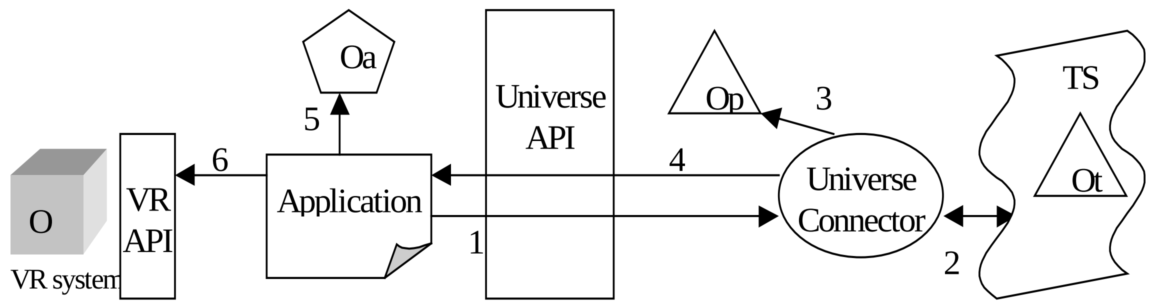

Figure 5: Importing a VR object from the universe

The application asks the Universe Connector for the object with the same application class, session ID, and object ID as was previously exported above (1). If these were not commonly known the application instance could issue various RD queries through the API to get a list of the current sessions for its application class in order to resolve which objects and session it wished to share. The connector queries the tuple space for objects with this matching application, session, object and data name. This time the query succeeds and the tuple is RD in (2). The connector proceeds to create a local proxy for the shared tuple (Op) and as before registers itself and the application to be notified whenever the data in the object is modified (3). A reference to the newly created proxy object is returned to the application (4) which then creates a new application specific object containing a reference to data in the proxy tuple (Oa) (5). Finally the application uses the VR API to add the new object (Oa) into the VR world (6).

The final stage in making the multi-user application is to make sure that actions occurring in the VR application are received by the universe aware application component which updates the shared data tuples containing related data. The connector then receives these local proxy updates and makes the appropriate changes to the tuple space.

Although this example has used a VR system as the application being made multi-user, any application which provides a sufficiently comprehensive API could be extended in the same way. Further, the use of the tuple space to mediate access to the applications shared data allows heterogeneous VR systems (providing they can interpret the data types and structure correctly) to share the *same* world *simultaneously* and is currently the subject of much further work.

Sharing across data models

One of the advantages offered by the use of a tuple space is the provision of a lightweight data model that allows information in different forms to be easily shared. To demonstrate this our second example considers sharing information between applications with different data models. This example also demonstrates how the separation between the application concerns and the tuple space can be exploited to allow standard applications to be linked with external systems.

As before, the Excel application does not provide any access to the universe platform but does provide a comprehensive API for accessing cell elements in a spreadsheet (in this case through Visual Basic). Therefore an additional application add-on is provided which provides the communication between Excel and the Universe API.

Initially the add-on, identified by the application class “x-application/x-excel-addon”, performs a sweep of the current state of the current excel spreadsheet. Each cell is shared as a distinct data tuple, where the session ID is the name of the spreadsheet the cell belongs to and the data is the contents of the cell itself. As changes are made to the spreadsheet the add-on reflects the changes in appropriate local proxy tuples, which the connector forwards to the shared universe.

Once the tuples are in the space, any application may get and share the current contents of the spreadsheet simply by performing a RD on all tuples with the application ID of “x-application/x-excel-addon” and the name of the spreadsheet as the session ID.

One application that may then use such data is a VR information visualisation called Q-pit (Benford & Mariani 1994)which maps data into a 3D space according to properties of that data. Clearly it is a simple task for the Q-pit to read out the spreadsheet from the universe and lay the data out in 3D. An obvious mapping could be to lay the cells out in two dimensions according to their X,Y coordinates and then represent their value in the Z axis. Interesting alternatives include laying the data out by mapping the actual data type of the value onto one axis, the value onto another and its two dimensional position onto the third.

While this example has used a spreadsheet to provide the data into the space, any source of data can provide information that can not only be visualised in alternative ways, but also shared between very different types of application.

Sharing multimedia data

Our final example focuses on the sharing of streamed multimedia data between applications. This final example shows how the tuple space can be used to stream video (MPEG) from one source to multiple recipients.

The original Limbo2 tuple space we use was initially designed to support streamed media, such as audio and video, across networks with variable performance and capabilities (e.g. between a wireless connection and a local area network). Consequently, the Shared Universe of data is not limited to facilitating the sharing of persistent (and fairly static) pieces of data. For example, Data tuples can be used to stream video between a video capture application and a textured video wall in a VR application

An MPEG video stream consists of several different types of frames. I-frames are considered key frames, storing the whole of an image. P and B-frames contain the differences between images in the stream. By defining a video stream to a single data tuple we can separate each type of frame into persistent updates (the I-frames) and transient updates (the B and P-frames). Streaming video from a single MPEG source (for example a video camera equipped with an MPEG encoder) to multiple recipients is then a case of putting each I, B or P-frame into a local proxy of the data tuple and generating the appropriate update on the tuple space. I-frames will cause the data tuple to be replaced, causing a persistent update, while B and P-frames will be manifest as quickly expiring transient update tuples. Recipients receive the frames through the normal state update mechanism.

It is worth noting that although the tuple space can support this type of streaming, a more pragmatic alternative would be to use the tuple space to support the setting up of the stream and multicast group, and then use multicast to actually send the data. However, extensions and further work on the actual tuple space API should alleviate many of the performance bottle-necks - such using a single “replace” call on the tuple space which would remove the need to IN the existing data tuple (and I-frame) in order to update it and re-OUT it.

Other Universe Services

The open design of the Universe API and tuple spaces makes it easy to add additional services and functionality to the Universe – an application or service needs only to connect to the tuple space and can then query, add or remove the tuples within it. One of the most common types of service is the persistence service. The service simply registers a callback for all tuples being added or removed from a particular tuple space. This callback then copies or removes tuples it receives from a local database. If the tuple space crashes or fails for some reason the backup application can restore the lost information by retrieving the original persistent data and relationship tuples from the database and merging them with the current state of the tuple space.

Future work

Our current work with the platform has focused on developing a uniform model for the common distribution of shared data between heterogeneous systems and applications. Two of the areas which must be addressed in future work include the nesting of tuple spaces and access control.

The current model uses a single flat tuple space to form a universe of shared data (although the tuple space itself can consist of a network of different tuple spaces). However because each tuple space is replicated on each machine, this potentially does not scale well. For example, streaming video through the same tuple space as all the other data will be detrimental to all of the connected systems, regardless of whether they are actually using or even aware of the streaming media. To overcome this we propose to allow applications to define how the tuple spaces can be nested into sub-tuple spaces.

Another important area is access control. Clearly applications may want to maintain some degree of control over data once it is shared with the universe, to enable it to be accessed by only certain people and so on. Again we could exploit the ability to nest tuple spaces to place particular restrictions on access to particular spaces. This could exploit special “access” tuples to enforce the access model.

The first two examples highlighted a common approach for applications to share data with the universe – to use a special add-on which knows about the universe but which can drive the legacy application. We see this as a common requirement of many applications and are examining techniques to extend the API to provide more automatic conversion facilities between shared proxy tuples and the end-application data they reflect. In particular we consider the development and integration of two services: the mapper and binder. The Mapper will get, modify and route events both to and from the application state being shared with the universe (e.g. the VR system) and the shared proxy objects of the universe itself (and visa versa). In essence, when an update is made on the application state, it is automatically converted by the mapper into the universe equivalent (if any or more than one) event and dispatched into the universe. Conversely, when an event is received from the universe concerning a change to some shared object state, this update event is converted into none or more specific updates on the application object.

The binder mediates access to the underlying universe tuple space for the application by creating instances of generic universe objects from specific system objects and visa versa. As such, the binder is able to convert between the different data-types and methods in these objects.

Conclusions

In this chapter we have presented the development of a platform that allows applications to exploit a shared data space that spans application boundaries. The developed platform exploits the concept of a shared tuple space to allow a number of disparate applications to make the information they share more widely available. The use of a tuple space provides a lightweight information model that allows the platform to support a wide range of information structuring models.

The existing view of tuple spaces considers them as passive repositories of information. To realise a cooperative platform that needed to support a wide range of different patterns of sharing we needed to extend the tuple space by providing active mechanisms that made applications aware of the data being shared.

The platform allows a wide range of shared space applications to share information between them. To demonstrate this we have extended an existing single user application to make the world model shared. To demonstrate the lightweight nature of the tuple data model we have also shared information between applications that had radically different information models.

References

Benford, S., Mariani, J. "Virtual Environments for Data Sharing and Visualisation - Populated Information Terrains". In Proc. ICDCS'94 (Lancaster, UK, July 13-15, 1994), Springer-Verlag. pp168-182.

Bentley, R., Dourish, P. "Medium versus Mechanism: Supporting Collaboration through Customisation" In Proc. ECSCW'95, Stockholm, Sweden, Sept., 1995, pp 133-148, Kluwer.

Bentley, R., Rodden, T., Sawyer, P., Sommerville, I., "An architecture for tailoring cooperative multi-user displays", In Proc. CSCW'92, Toronto, Canada, Oct., 1992, pp 187-194, ACM Press.

Carlsson, C., Hagsand, O. "DIVE: A Multi User Virtual Reality System". In Proc. IEEE VRAIS, (Sept. 1993), pp394-400

Davies, N., Wade, S., Friday, A., Blair, G. "Limbo: A Tuple Space Based Platform for Adaptive Mobile Applications". In Proc. ICODP'97, Toronto, Canada

Dewan P., Choudhary R., Flexible User Interface Coupling in a Collaborative System, Proceedings of ACM CHI'91, New Orleans, April 1991, ACM Press, 41-48

Dewan, P., Choudhary, R., "Flexible User Interface Coupling in a Collaborative System", In Proc. CHI'91, New Orleans, LA, April, 1991, pp 41-48, ACM Press.

Dix, A., Challenges for Cooperative Work on the Web: An Analytical Approach, CSCW: The Journal of Collaborative Computing, 6, 2-3, 1997, pp135-156.

Dourish P., Bly S., "Portholes: Supporting Awareness in a Distributed Work Group", Proc. CHI'92, Monterey, CA, May 1992, ACM Press, pp 541 547

Dourish, P., "Introduction: The State of Play", Special Issue of Computer Supported Cooperative Work, 7(1), 1998, Kluwer.

Foster, G., Stefik, M. "Cognoter: Theory and practice of a Colab-orative tool", In Proc. CSCW'86, Austin, TX, December, 1986, pp 7-15, ACM Press.

Gaver, W., et al. "Realizing a video environment: EuroPARC's RAVE system." In Proc. CHI'92. 1992. Monterey, CA, USA: ACM

Gelernter, D. "Generative Communication in Linda." ACM Transactions on Programming Langauges and Systems, 7(1), pp80-112.

Greenberg, S. "Personalizable groupware: accommodating individual roles and group differences", In Proc. ECSCW'91, Amsterdam, September 1991, pp 17-31.

Greenberg, S. Marwood, D., "Real Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface", In Proc. CSCW'94, Chapel Hill, NC, Oct., 1994, pp 207-217, ACM Press.

Greenhalgh, C., Benford, S. MASSIVE: "A Virtual Reality System for Tele-conferencing". To appear in ACM Trans. on Computer Human Interfaces

Hall, R.W., Mathur, A., Jahanian, F., Prakash, A. Rasmussen, C., "Corona: A Communication Service for Scalable, Reliable Group Collaboration Systems", In Proc. CSCW'96, Boston, MA, Nov. 1996, pp140-149, ACM Press.

Henderson, A.J., Card, S.A., "Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention", ACM Trans. on Graphics, 5(3), 1985.

Hill, R. D., Brinck, T., Patterson, J. F., Rohall S. L., Wilner, W. T. "The Rendezvous language and architecture", In CACM, 36(1), Jan.1993, pp 62-67.

Isaacs, E., Tang, J. C., Morris, T. "Piazza: A Desktop Environment Supporting Impromptu and Planned Interactions", In Proc. CSCW'96, Boston, MA, Nov. 1996, pp 315-324, ACM Press.

Ishii, H., and Kobayashi, M. "Clearboard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact". In Proc. CHI '92, Monterey, CA, May 3-7, 1992, pp. 525-532. ACM Press.

Kaplan, S. M., Tolone, W. J., Bogia, D. P., and Bignoli C. "Flexible, active support for collaborative work with ConversationBuilder", In Proc. CSCW'92, Toronto, Canada, Oct., 1992, pp 378-385, ACM Press.

Lee, J. H., Prakash, A., Jaeger, T. "A Software Architecture to Support Open Distributed Collaboratories", In Proc. CSCW'96, Boston, MA, Nov., 1996, pp 344-353, ACM Press.

Mansfield, T., Kaplan, S., Fitzpatrick, G., Phelps, T., Fitzpatrick, M., Taylor, R. "Evolving Orbit: a progress report on building locales". In Proc. GROUP_97 (Phoenix, Arizona, Nov. 16-19, 1997), ACM, pp241-250.

Palfreyman K. A. Rodden T. "A Protocol for User Awareness on the World Wide Web". In Proc. CSCW'96, Boston, MS, USA, Nov. 1996, ACM Press.

Patterson, J. F., Day, M. Kucan, J., "Notification Servers for Synchronous Groupware", In Proc. CSCW'96, Boston, MA, Nov. 1996, pp 122-129, ACM

Patterson, J. F., Hill, R. D., Rohall, S. L., Meeks, W. S., "Rendezvous: An Architecture for Synchronous Multi-User Applications" In Proc. CSCW'90, Los Angeles, CA, Oct., 1990, pp 317-328, ACM Press.

Root, R. W. "Design of a Multi-Media Vehicle for Social Browsing" In Proc. CSCW'88, Portland, OR, Sept., 1988, pp 25-38, ACM Press.

Roseman, M., Greenberg, S, "TeamRooms: Network Places for Collaboration". In Proc. CSCW_96 (Boston, MA, Nov.16-20, 1996), pp 325-333, ACM Press.

Roseman, M., Greenberg, S. "GroupKit: A Groupware Toolkit for Building Real-Time Conferencing Applications", In Proc. CSCW'92, Toronto, Canada, October, 1992, pp. 43-50, ACM Press.

Roseman, M., Greenberg, S., "Simplifying Component Development in an Integrated Groupware Environment", In Proc. UIST'97, Alberta, Canada, Oct. 1997 pp 65-72, ACM Press.

Stefik, M., Bobrow, D. G., Foster, G., Lanning, S., and Tatar, D., "WYSIWIS Revised: Early Experiences with Multiuser Interfaces" In ACM Trans. Office Information Systems, 5(2), 1987, pp 147-167.

Tang, J. C., "Findings from observational studies of collaborative work", In International Journal of Man-Machine Studies, 34(2), 1991, pp 143-160.

The Contact Consortium, http://www.ccon.org/

Chapter Eight: Using presence to manage shared e-scape displays

Gareth Smith, Jon O'Brien, Tom Rodden

Lancaster University

Introduction

The development of shared environments and displays has also seen the emergence of facilities to allow some form of subjective tailoring of shared interfaces. Different researchers have investigated techniques and architectures to support the real time management of these interfaces. Research has focused on the development of facilities to coordinate and manage interaction across the different interfaces. This chapter considers the need to dynamically re-couple tailored interfaces as users become increasingly aware of each other. We present a general model to support awareness based re-coupling of shared interfaces and show its implementation in cooperative virtual environments and shared graphical displays. This general model builds upon the issues of intersubjectivty discussed previously in this deliverable and the considerations of subjectivity in Chapter 10 of Deliverable 3.1.

The recoupling model presented here drawns upon the need to conisder the commonaly shared by multiple inhabitants of a shared world. In essence, the model considers the development of techniques that allow the presentation of an e-scope to dynamically alter to reflect the manner in which it is been shared and the need to promote sufficient commonality of display to allow a world known in common to emerge. This work builds upon previous research within the CSCW community into the manner in which shared displays can be managed. In addition, to considering the application of these mechanisms within shared 3D virtual environments we also present a consideration of how the general model may be applied within 2D interfaces.

Early multi user interface systems supported shared interfaces by presenting exactly the same image of the application to all users. This simple replication of the system's image secured a founding abstraction for multi user interfaces: What You See is What I See (WYSIWIS). Experiences with these systems in the CoLab environment highlighted problems with the WYSIWIS approach (Stefik et al. 1987), Stefik concluded that "*WYSIWIS (What You See is What I See) is too inflexible, if strictly interpreted, and must be relaxed to better accommodate important interactions in meetings*". Stefik suggested that the shared interfaces which conformed to the WYSIWIS abstraction should be relaxed along a number of dimensions. The notion of relaxed WYSIWIS provides each individual user with the ability to configure their shared user interface to best suit their working needs. User interface tailoring is now generally accepted in the areas of all user interface systems (Greenberg 1991).

A similar approach is emerging in the field of Cooperative Virtual Environments (CVEs). CVEs may be considered as shared *three dimensional* interfaces which provide a shared 3D space that may be populated by users and any type of three dimensional artifact. Generally, CVEs present the same shared data to each of the cooperating users, albeit from a different viewpoint, which is analogous to the use of strict-WYSIWIS in early 2D interfaces. That is, they lack a number of cooperative features such as interface tailoring. But, the lessons from work in the field of shared 2D user interfaces are being introduced into the area of 3D shared user interfaces. For example, a model to support individual *Subjective Views* of a shared 3D world has been presented (Smith 1996), and implemented (Smith & Mariani 1997) within a publicly available CVE, the *Dive* system (Carlsson & Hagsand 1993).

Shared user interface research has formed a solid theme of supporting the ability for users to tailor their user interfaces on an individual and group basis. However, different presentations of shared user interfaces, whether 2D or 3D, can cause problems in collaboration due to lack of 'common knowledge' and awareness between the collaborating parties and their activities (Susan et al. 1992). This problem is also highlighted by (Gutwin & Greenberg 1996) but is termed *workspace awareness* and defined as "*the up-to-the-minute knowledge a person holds about another's interaction with the workspace*". Therefore there is an operational trade-off between the benefits of interface tailoring and the amount of 'common knowledge' required to facilitate collaboration (Ellis et al.). Generally, this fine tuning of user interface coupling between users is the responsibility of the users themselves. This activity is problematic as it is difficult for individual users to know the customised details of other users' interfaces.

This chapter argues that user interface tailoring mechanisms provide individual users with different *static* user interfaces, which must be explicitly re-configured to aid in certain collaborative tasks. Additionally, the process of *user interface re-coupling* must be an automatic process without placing additional responsibility on the user. Rather, it should operate seamlessly, and re-configure the different shared user interfaces on the users' behalf.

Motivation

The motivation for this work comes from a consideration of empirical studies of action within a CVE constructed in MASSIVE (Greenhalgh & Benford 1995) as part of a previous project (Bowers et al. 1996) These ethnographic studies were carried out with the intention of exploring the ways in which users might make use of the CVE for interactional purposes, and subsequently making these findings available as a resource for the refinement of both user embodiments and the CVE itself. Much of this work focused on the interactional affordances of the user embodiments, and some mention was made of the use of alternative ‘views’ of the virtual world by users in order to achieve certain activity-specific practices most effectively. We include one particular extract here, as an example of the exploitation of ‘freedom of viewpoint’ along with a brief section of discussion of the action within the CVE, taken from (Bowers et al. 1996).

> CT suggests gathering in a circle. Some complaints... notoriously difficult to manage. Problems of orienting one’s own embodiments to those of all others. Some attempts at first, people’s embodiments run through each other.
>
> Someone from Site B explains that you can switch to a bird’s eye view and see your own and other embodiments from directly above, much easier to orient embodiments in relation to others.
>
> Initially users struggled to co-ordinate their position in relation to other users, before it was pointed out that they were able to ’leave’ their embodiments and look down on them from above, making such co-ordination much more manageable. Here the ability to exploit the ’freedom’ of viewpoint in MASSIVE contributed to the successful completion of a task by making all participants aware of the work of others. The example also indicates that the form which embodiment in a CVE should take (whether one ’looks out of’ the blockie’s eye or prefers an ’out of the body experience’) is a contingent and activity specific affair

Of interest here for the purposes of this chapter and the development of e-scapes is not so much the fact that users exploited the ‘freedom of viewpoint’ in order to manage gathering in a circle within the CVE, but the *assumptions about the properties of the CVE* that provided for this strategy as one available to users for the orientation of their embodiments towards one another. The viability of adopting such a ‘bird’s eye’ view as a strategy of this kind is based upon a knowledge that all participants see the CVE as constituted in the same way: their *perspectives upon* the CVE may well be different (in this case they may not all share the same ‘bird’s eye’ view, for example) but they all share the *same understanding* of the CVE as a space within which they are all embodied. This intersubjective consistency therefore allows for a range of activity-specific strategies for interacting within the CVE – participants can be sure that they share a given common intersubjective space upon which they are taking individual perspectives. Were this intersubjective consistency *not* to hold, were participants all to have radically different understandings of the underlying constitution and properties of the CVE, it would make no sense to adopt the ‘bird’s eye’ view as a strategy for coordinating the position of one’s embodiment with others, since one could not be sure that others would be experiencing the space within the CVE in an intersubjectively consistent way. In short other objects, other actions, not accessible to the adopter of the bird's eye view might be constraining the action and interaction of those being viewed from above and hence the utility of this manoeuvre would disappear.

Approach

The notion of coupling in shared user interfaces is currently based on some form of derivation from a common interface. The common interface definition provides a basis for tailoring by users. Each user controls their own interfaces by altering their presentation of a common shared model. This model may be expressed as a user interface (as in SOL (Patterson et al. 1990)) or a set of shared state (as in Rendezvous (Smith & Rodden 1996)). Essentially this approach means that each subjective view results in a separate branch within a derivation hierarchy from the common shared interface, see figure 1.

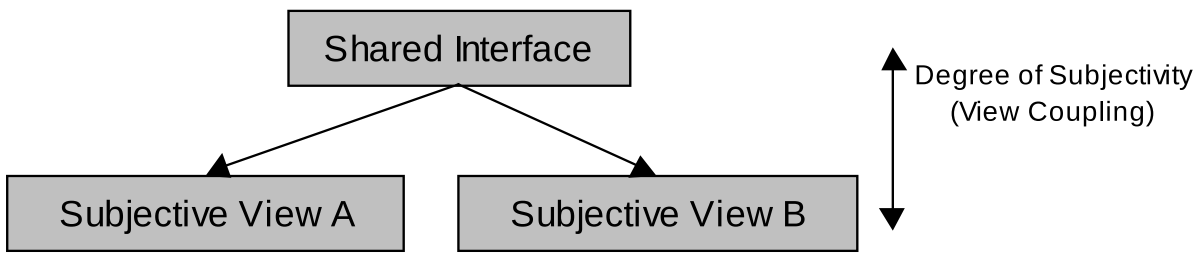

Figure 1. The subjective or personal view model

This traditional approach to developing subjective views means that each user has a separate perspective on the shared model. This separation of interfaces means that users can manage the local presentation of information and have personalised views of information intermingled with common elements. The problem we wish to address is the situation where we wish to bring these currently separate views together to allow users to share a greater common experience and to promote a greater degree of commonality by re-coupling elements of the interface.

Obviously one way to re-couple the separate interfaces is for users to revert to the common interface definition. However, this requires users to reconfigure their interface definition and is potentially rather costly. We wish to develop a low cost means of re-coupling these interfaces depending on the context of use. The approach we have adopted is based on the consideration that a user's subjective specification reflects their personal preferences, and should not have to be manually re-configured to suit transitory working conditions. Rather we wish to exploit an awareness of other users' activities and interfaces to automatically recouple the users interface.

The core of our approach is the explicit identification of the awareness users have of each other. The knowledge of this awareness is combined with the subjective information specified for each user, to generate an *awareness sensitive subjective view*. Figure 2, depicts the process where the knowledge of a user's subjective view is combined with *an awareness map* showing the awareness the user has of each object to create a new subjective view. This new subjective view is a second-generation subjective view derived from a combination of the user's initial subjective configuration and the awareness map derived from other users' subjective views.

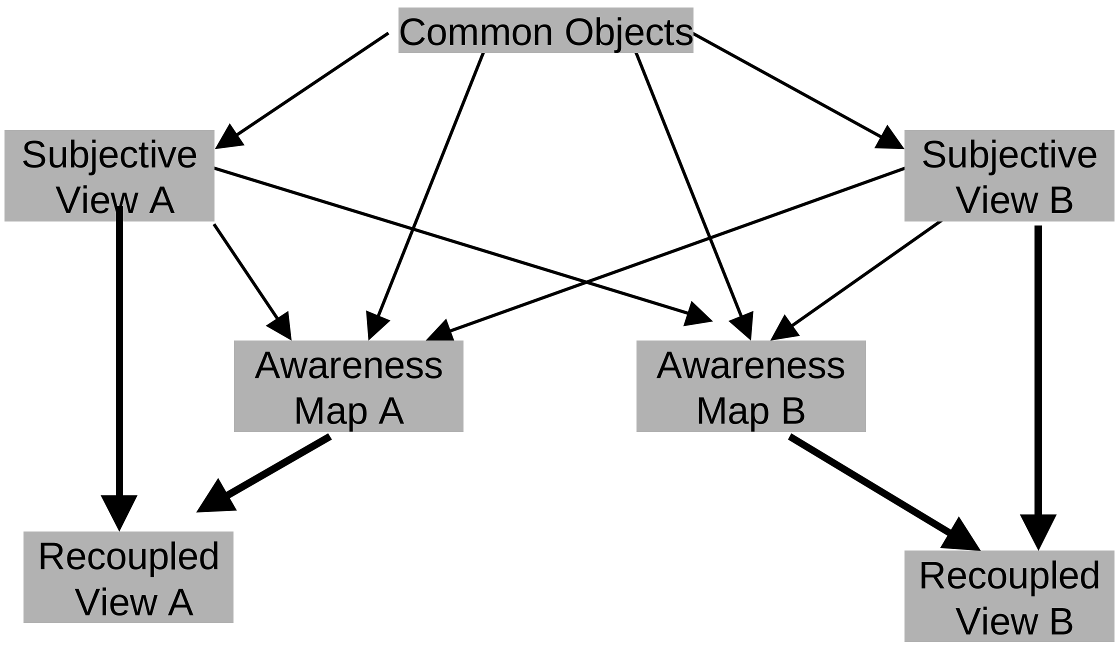

Figure 2. Generation of Recoupled-Subjective Views

An awareness Model to support subjective views.

To develop these re-coupling mechanisms we build upon an existing awareness model (Rodden 1996) that describes the interactivity between users. We extend this model to allow it to be used to further manipulate users' subjective views of the common model. The extension creates an additional level of user interface control based on the common interface model and users' subjective preferences.

The awareness model provides a generalised approach for understanding shared objects and multiple users of cooperative systems. The benefit of this model is that it enables us to focus on the shared manipulation of objects, without being restricted to a particular cooperative system.

The key elements of the model are motivated by a spatial consideration of users populating a shared pool of objects. The core of the awareness model is the view of these shared spaces as a set of interconnected objects. The key concepts used to describe awareness in the model are *location, focus* and *nimbus.*

-- *Location*: the object where the user is currently positioned

-- *Focus*: a set of adjacent objects the user is interested in

-- *Nimbus*: a set of adjacent objects that the user is effecting

It is possible for a user to have more than one position in this pool of objects and to have some form of aggregated focus and nimbus associated with them. Each of location, focus and nimbus can be considered as sets of objects. Figure 3 shows the general arrangement for a user ’tom’ in a pool of objects.

Using this model we can represent the objects where users are placed as sets of adjacent objects that can be used to determine how aware users are of each other. We wish to extend this model to allow us to describe the general mechanism of re-coupling independently of a particular user interface paradigm or implementation. In developing an extension of this model we seek to exploit the general nature of sets to develop an understanding of awareness across subjective world views.

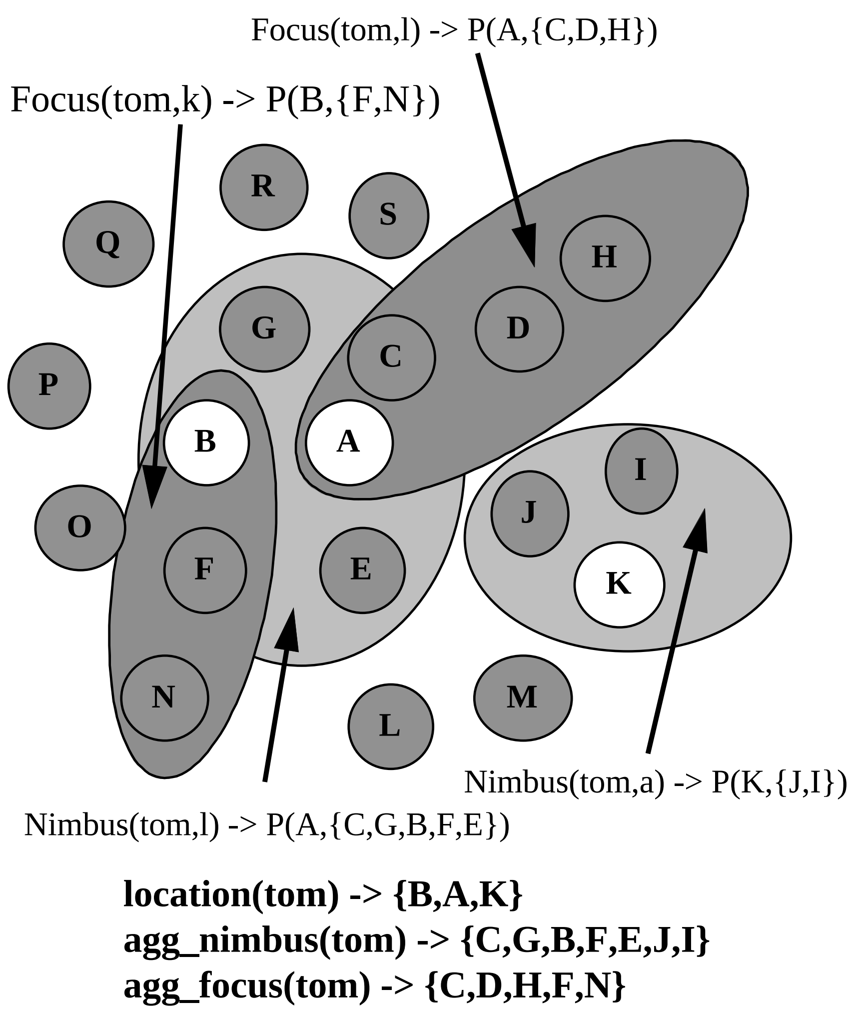

Figure 3 Location, aggregate focus and aggregate nimbus

The core of the awareness model is the use of simple sets to represent location, aggregate nimbus and aggregate focus and the comparison of these sets allows us to reason about awareness. For example, we can decide that we have an awareness strength measure by using a simple ratio of number of objects in the intersection of the focus and nimbus sets divided by the number of objects in the union of these sets.

$$\frac{|\,agg\_focus(u1) \cap agg\_nimbus(u2)\,|}{|\,agg\_focus(u1) \cup agg\_nimbus(u2)\,|}$$

Subjective Views of the Common Space

Our extension relies on understanding the particular derived views that an object has at any one moment in time and representing these within the model. This relies on the use of a *subjective pair* which show that one object has been derived from another. This subjective pair linking a common object (o_c) and the subjective object for a particular user (o_u) is represented as:

$$S(o_c, o_u)$$
$$\text{where } o_c \in O \text{ and } o_u \in O$$

This pairing allows us to construct a *subjective space* (SS(0)) to represent the derivation of alternative views from the common set of objects. The subjective space can be represented as.

$$SS(O) = \{ S(o_c, o_u) | \; o_c \in O, o_u \in O \}$$

We refer to the particular parts of the pair using a simple bracket notation. Where S is a subjective pair

Common(S) is the common object
Subjective (S) is the subjective object

We can also refer to the these components collectively where SS is a subjective space

Common_objects(SS) is the set of common object
Subjective_objects(SS) is the set of subjective object

A general function *SubjectiveModifier* is used to represent the decoupling of objects from the common space s

$$SubjectiveModifier: U \rightarrow SS(O)$$

This simply states that the user **U** has a subjective space consisting of the set of pairings of common objects linked to the subjective object replacing it. This allows us to reason about both the users common object space and the subjective object spaces presented to the user. To aid this we use a function *projected_space* that takes as a parameter a set of objects and a subjective space and replaces the objects in the common space with their subjective counter parts.

$$\text{ProjectedSpace: } SS(0) \rightarrow Op$$

Where Op is the set generated by replacing all the common_objects with the corresponding subjective_objects.

Let us illustrate this with a simple example. Consider two users tom and dick sharing a common space consisting of three different objects A,B,C. If tom modifies object A to create A' and dick modifies object C to create C' then both have subjective versions of the shared space. We can say that

$$\text{SubjectiveModifier(tom)} = \{S(A,A')\}$$

and

$$\text{SubjectiveModifier(dick)} = \{S(C,C')\}$$

We can also represent the two different versions of the shared space using the projected space function.

$$\text{ProjectedSpace (SubjectiveModifier(tom))} = \{A',B,C\}$$

and

$$\text{ProjectedSpace (SubjectiveModifier(dick))} = \{A,B,C'\}$$

Presenting the space

In developing our subjective view we also need to consider how a space is made available to users. Remember that we are considering users sharing a space constructed from sets of common objects and allowing an individualised projection of this common space by selectively replacing objects with alternatives. We do this by representing view characteristics of the objects in each user's view of the space. To do this we again exploit a simple pair linking an object (o) and a view modifier (m) represented as:

$$V(o, m)$$
$$\text{where } o \in O \text{ and } m \in M$$

The set M is a set of *modification tokens* that instruct the application how to display an object to the user. This token is application specific and particular interpretations of the model may use it differently. For example, in a simple GUI it could consist of the tokens {*invisible, normal, shaded*} and these would be interpreted in displaying these objects.

The view pairing allows us to construct a *view space* (VS(0)) to represent a user's view of a set of objects. The view space can be represented as.

$$VS(O) = \{ V(o, m) \mid o \in O, m \in M \}$$

A users view of a set of objects can then be represented by the function *Subjective_View* taking a set of objects and a user as a parameter.

$$\text{Subjective\_View: } U \times O \rightarrow VS(0)$$

In terms of the example used above we could imagine that object tom sees his modified object A as shaded and the others as normal. In contrast dick does not see A, sees object B as normal and his modified version of C as shaded. This is represented as.

Subjective_View(tom,{A',B,C}) = {V(A', shaded),V(B, normal), V(C,normal)}

Subjective_View(dick,{A,B,C'}) = {V(A, invisible),V(B, normal), V(C', shaded)}

Developing an Awareness Map

We can combine the common space and each user's subjective view to consider awareness. To do this we exploit an awareness pair that links objects(o) with an awareness modifier(a). This is represented as:

$$A(o, a)$$
$$\text{where } o \in O \text{ and } a \in A$$

The set A is a set of *awareness modifiers* that indicate how aware a user is of a particular object. This is a product of considering the other users' awareness of a user and their subjective view of an object. The generation of these awareness modifiers is determined by the detailed semantics of a particular application. Our general case of the model only considers the representation of these modifiers. We use these pairs to form an *awareness map (AM(0))* for each users. The awareness map is represented as a set of these pairs.

$$AM(O) = \{ A(o, a) | \; o \in O, a \in A \}$$

The generation of an awareness map is determined by considering the common space of objects(0) and each users' subjective space (using the different *subjective views* generated). This general arrangement can be represented using the simple function map_awareness that generates a map for a given user.

map_awareness: U -> AM(0)

Depending on the details of the particular application a given set of awareness values may result from a consideration of:

-- Awareness functions across the common space;

-- Awareness functions across users projected spaces;

-- Awareness functions across the objects unaltered by any subjective projection.

In addition to considering how some sense of awareness is calculated for a given space the awareness map shows how the particular defined subjective view should be altered for each object. Each of the awareness modifiers indicate how the presentation may be altered to reflect the way in which other users are seeing the object. These modifiers are core to the development of a recoupled view of these objects.

Recoupling the views

At any given moment we can recouple different user's views by exploiting each users awareness map and particular view. If we consider the display for a given user(tom) drawn from a set of users ({tom,dick,harry}) we can say the user(tom) has an subjective view space(VS) represented as

VS = subjective_view(tom, ProjectedSpace (SubjectiveModifier(tom)))

We can also say that the awareness map (AM) for this user can be represented as

AM = map_awareness(tom}

A final function *display* combines these two sets to give a modified view space that can be presented to the user. This is of the form

$$\text{display: } VS(0) \times AM(0) \rightarrow VS(0)$$

This arrangement allows each user's subjective display to be modified to reflect the effects of others subjective displays. The awareness based recoupling of these subjective displays can exploit a range of different particular algorithms. The framework presented in this chapter aims at being independent of the particular *policies* used to detail how two or more objects should be merged. These might range from the use of strength values to represent the view and awareness modifiers to the more complex re-coupling mechanisms offered by (Munson & Dewan 1994).

In the following sections we consider the implementation of the re-coupling model in different types of shared user interface systems. These different implementations of the model provide two different interpretations of the model The first is within a Collaborative Virtual Environments and the second is within a shared 2D user interface system.

Re-Coupling in a Virtual Environment

Cooperative Virtual Environments (CVEs) support the population of generic 3D spaces by a number of users. Such systems include DIVE (Carlsson & Hagsand 1994), dVS (Division Ltd. 1993), and MASSIVE (Greenhalgh & Benfird 1995), and allow users to navigate through the shared space and interact with objects and other users. A central facility of CVEs is the presentation of the shared environment, where the shared world is rendered onto the user's display. In a number of these systems, each users' presentation of the same shared world may differ through the use of subjective views where the underlying world object is replaced by a presented object. Subjectivity is used in these systems in a number of different ways, for example, a user may wish to find the location of the church, or be shown the quickest path from the church to the museum. Both these queries can be visualised subjectively, as depicted in figure 4.

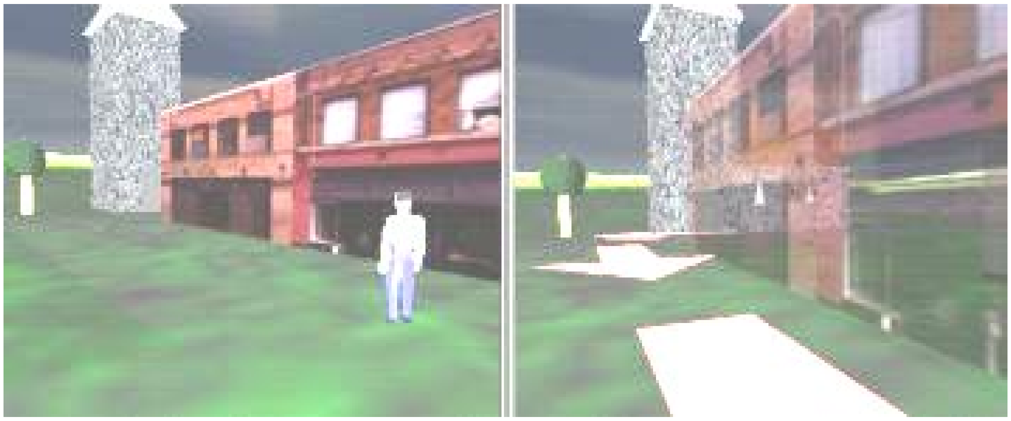

Figure 4. Highlighting locations of interest, by making other objects appear as transparent, and visualising paths between them.

Throughout, our discussion of the implementation an example scenario of engineering design (of aircraft) is used. Here two or more users inspect the internals of an aircraft, which normally contains a number of separate systems organised in close proximity. Figure 5 denotes such an area, where a section of the aircraft's hydraulics are located near some of the aircraft's electrical system.

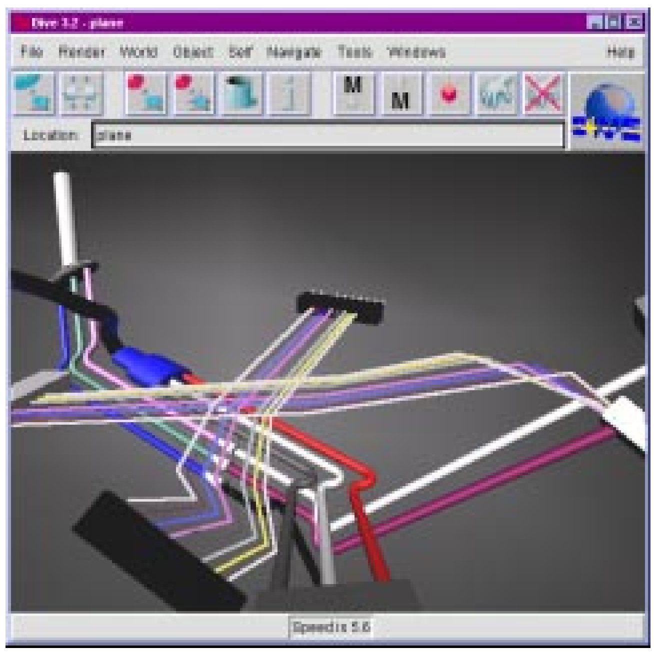

Figure 5. Inside the aircraft, a normal view

Using subjective views, each of the users' displays can be configured differently. This enables each engineer to concentrate on the objects most relevant to him or her. The CVE used in this example is Dive, which utilises the SOLVEN (Smith & Mariani 1997) model of developing subjective views. This allows objects to have a number of alternative representations, and supports a range of modifiers which manipulate the visual aspects of the geometry. For example, a wireframe modifier will cause the object to be drawn in wireframe mode. This forms a matrix of representations against modifers, as depicted in figure 6.

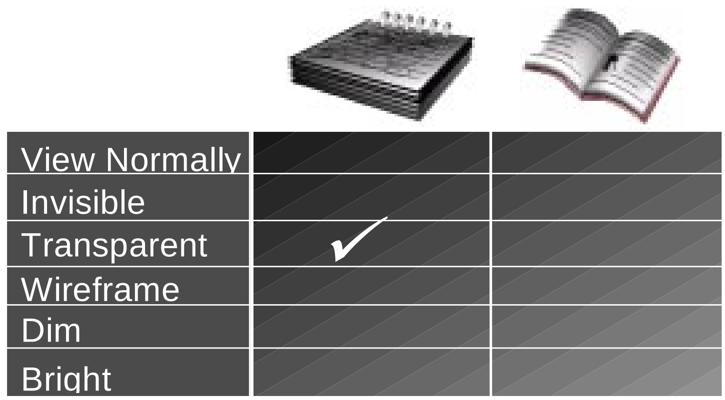

Figure 6. A view matrix of a book, showing two possible presentations and six modifiers.

In the above diagram, the user has chosen to represent a book in the style of a note-book (rather than a traditionally bound book), and that it should appear transparent. Using this model, users can selectively configure their interfaces.

In the engineering design example, each class of engineer may decide to remove the view of objects that are generally not relevant to them. Figure 7, depicts the isolated components that compose the subjective views of the hydraulics engineer and the electronics engineer.

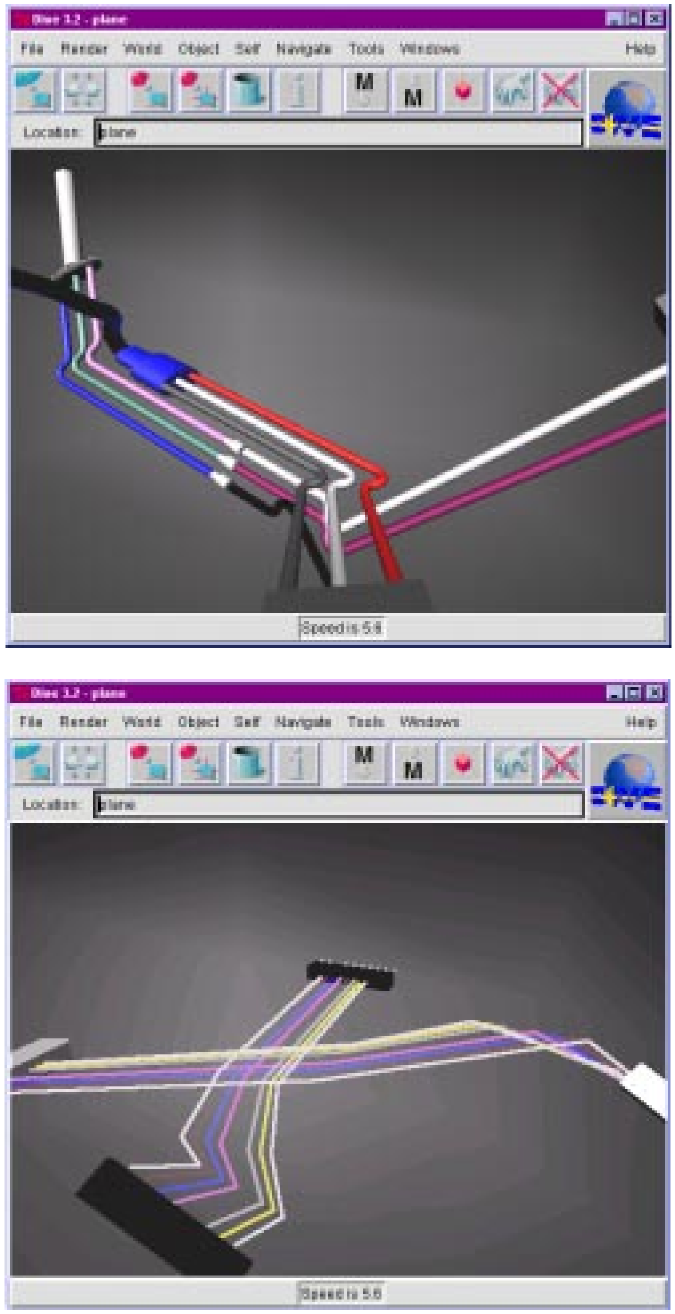

Figure 7. The hydraulic engineer’s view and the electrician’s view

When working on their individual tasks the subjectivity views enable them to concentrate on their specialised domain. However when two or more users are working across a similar set of objects their different views can lead to communication and referential problems due to the lack to common knowledge. To overcome this users involved would normally need to adjust their subjective specifications, requiring each to describe the state of their interface. In our case however we can exploit an awareness of the other users to allow a recoupling of these displays.

Mapping to the model

In this application of the general awareness driven recoupling model we exploit the following interpretations:

-- The set of common objects maps to objects in the CVE

-- Object replacements map to representations in the projected spaces.

-- The view modifier tokens map to the modifiers in the SOLVEN model

Describing the Subjective views

The specification of the subjective views relies on the different display characteristics the SOLVEN subjectivity mechanism. These display characteristics are mapped to view modifier tokens. We exploit a set of simple integer values to encode this information as shown in the table below.

Display Characteristic	View Modifier
Bright	*+1*
Normal	*0*
Dim	*-1*
Wireframe	*-2*
Transparent	*-2*
Invisible	*-3*

These values are used to represent the different presentations of objects in each user's display. In the case shown above we can consider each user's subjective display space as consisting of sets of subjective pairs linking objects to view modifiers. The subjective specification of the 'Wiring' and 'Hydraulics' objects for both the electronic engineer and the hydraulic engineer are detailed below:

Electronic Eng. Object	View Modifier
Wiring	+1 (bright)
Hydraulics	-3 (invisible)

Hydraulic Eng. Object	View Modifier
Wiring	-3 (invisible)
Hydraulics	0 (normal)

In the model this is represented as:

Subjective_view: (Elec_Eng,0) = { (Wiring,+ 1), (Hydraulics,-3) }

Subjective_view: (Hydr_Eng,0) = = { (Wiring, -3), (Hydraulics,0) }

These two different subjective views now need to be recoupled to take account of each user's view and the awareness the users have of each other. What we seek to do here is develop a view that allows for some form of intersubjectivity where users are aware of the presence of others and what they may also be seeing. The basis of developing this coupled view is through the development of an awareness map for each user.

Developing an Awareness Map

The core to deriving the re-coupled view is an encoding of other users' awareness of each other. This is calculated in terms of the common object model and the different subjective views of other users. We use this information to form an *awareness map* made up of a pairing of shared objects and an awareness modifier. In the case of the CVE the awareness modifier is calculated using a field awareness method that exploits each users nimbus on the common space.

Using spatial awareness

The spatial awareness approach considers awareness by examining a user's nimbus in terms of which other users' nimbi are in collision with it. Then, the view modifiers from those users are interpreted and used along with the awareness measure to construct the original user's awareness map.

In our implementation recoupling operates in a single direction (making objects *more* visible to a user). This reduces conflict when a number of users are in close proximity as the *most aware* (greatest number) value is chosen. Hence as a general rule, as more users, with different subjective settings, converge at a common location the more their displays loose their subjective values, and a more *common* interface is provided to each. To reason about awareness we use a visibility measure which is a percentage of how solid an object appears to users.

Consider the case depicted in figure 8. The subjective modifiers of a particular object have already been identified for three users. User C sees the object normally and it is invisible for Users B and C. The crossover of nimbi is used to calculate the awareness modifiers. User C affects User B's display making the object in that display 75% visible (or 25% transparent) and means that users B awareness map contains an awareness pair for the object encoding this.

In the case of our CVE development, the awareness calculation also considers the propagation of effects in building awareness maps. Before the arrival of User C, both User A and B did not see the object. However, now that User B can see this object it affects User A's display in the same way that User C affected User B. User A has their awareness map modified to indicate that the object can be made 15% visible as compared with User B's display. Hence the object in User A's display appears as 15% of 75%, which is 11% visible.

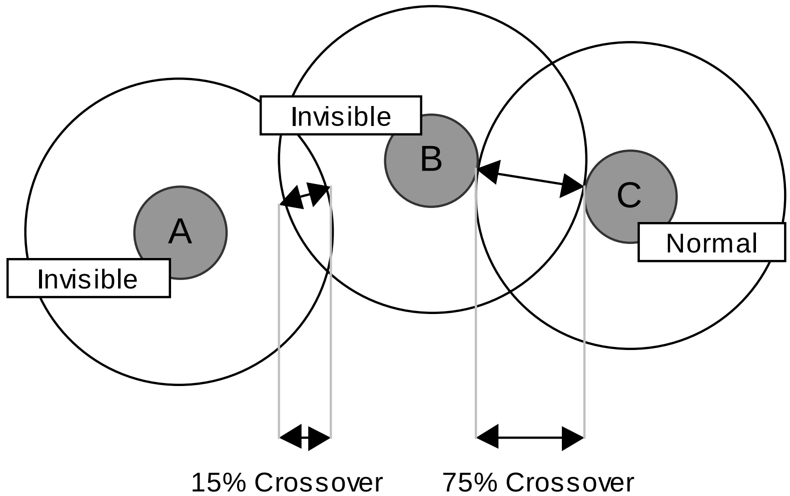

Figure 8. User C directly affects User B's display and indirectly affects User A's display

The awareness map for each user is modified to reflect changes in the environment. Consider the case where user C has moved closer to User A with a crossover of 60%, but the distances between Users A and B and Users B and C are unaltered. User A's awareness map is now directly affected by User C and User B. However, User B's impact on User A is to make the object 11% visible (through User C), and User C's impact on User A is to make the object 60% visible. Hence, the object in User A's awareness map is modified to give it 60% visibility (40% transparent), as our awareness mechanism chooses only the highest possible value.

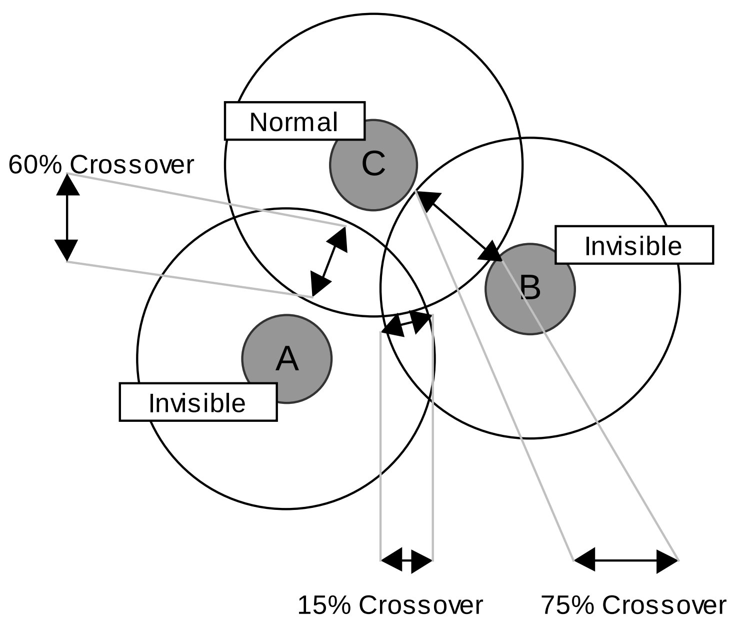

Figure 9. User C directly affects User A's and User B's displays

In some circumstances the re-coupling mechanism can cause unwanted disturbances. If this is the case it is possible for users to reduce the size of their nimbus (see figure 10), or switch it off completely. This provides a simple one-dimensional metric to control the amount of inter-subjectivity a user wishes to promote and receive. In certain conditions it may be advantageous to provide users with a range of differently shaped nimbi to choose from. For example, in a meeting room scenario, it may be desirable to choose a forward only facing nimbus, emphasising the activities in the centre of the room while reducing those behind the users.

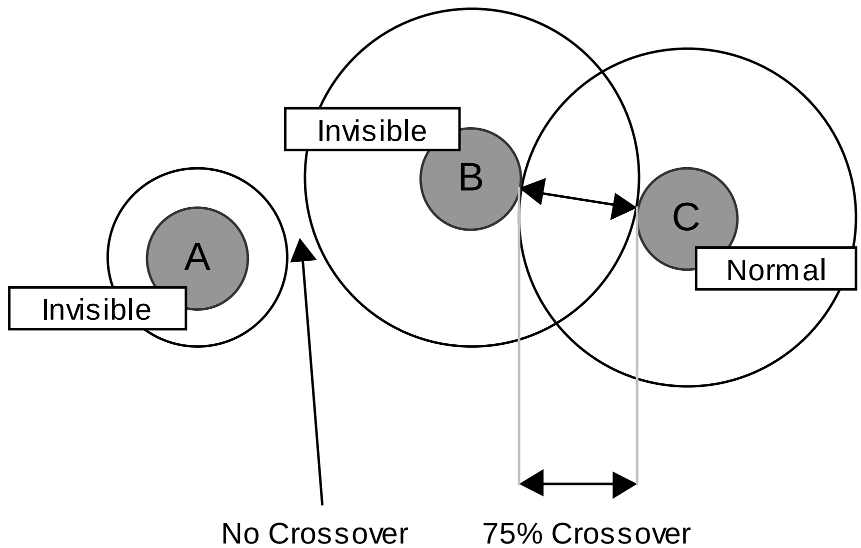

Figure 10. User A reduces the size of their nimbus to reduce intrusions from other users.

The recoupled displays

The object based awareness model detailed above is used to calculate an awareness map for each user that encodes the awareness effects of other's displays of objects. The re-coupling mechanism that generates the final display combines this information in the awareness map with the subjective view to give a final user display. This allows the displays to independently adjust themselves based on the awareness each user has of the objects in their display.

The resulting view modification is some product of an object's existing modification token and the awareness modifier. Our implementation operates in a increasing manner, in that objects should always become more visible through this mechanism. For example, in the electronic engineer's view the wiring is bright (+1), but invisible (-3) in the hydraulic engineer's view. In the case of the electronic engineer the existing value is the maximum and is unaltered by the awareness. In the case of the hydraulic engineer the view modifier (-3) can be less than that suggested by an awareness modifier. When it is greater the final display is altered within the range [-3 ➔ +1].

The degree to which the change is dependent on the awareness modifier value. In this implementation it is based on a percentage visibility of other users' view modifiers, see figure 11. The value of this visibility is linked to the overlap of nimbi and the view modification value for other users. This means that as the users' overlap increases the percentage visibility will increase and this will be applied to the highest view modifier.

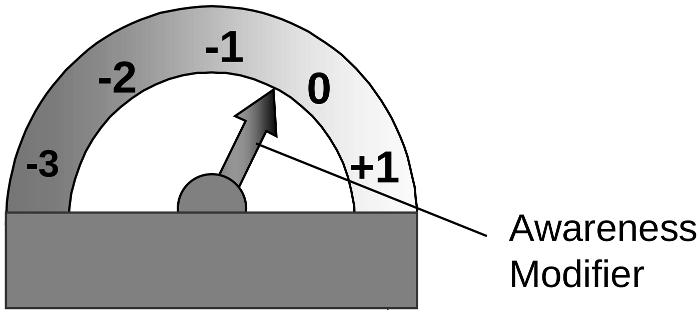

Figure 11. The selection effect of the awareness modifier

In the case of our example the electrical engineer has removed their view of the hydraulic system, but as the other engineer is aware of this system, this is reflected in the display of the electrician. Likewise the electrical system will appear in the display of the hydraulic engineer, see figure 12.

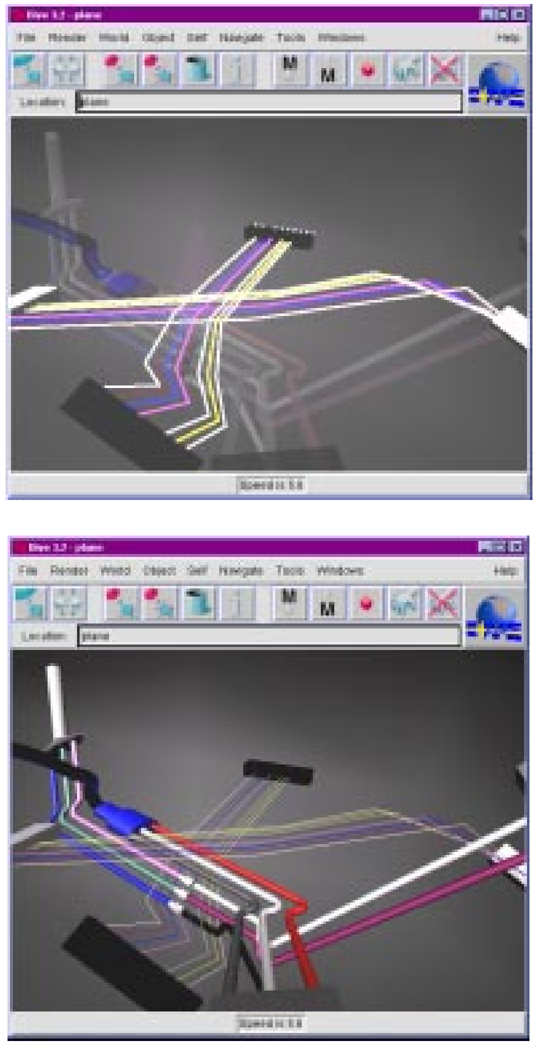

Figure 12. The common-subjective views of the electrician and hydraulics engineer.

Implementing Re-Coupling in non-3D environments

User interface re-coupling can be used to enhance shared user interfaces, other than those of 3D virtual environments. Many shared 2D user interface systems, such as SOL (Smith & Rodden 1995), Suite (Dewan 1990) and GroupKit (Roseman & Greenberg 1992) support individual user interface tailoring and a range of user interface coupling. The notion of awareness based re-coupling can be introduced into these systems to both promote a common understanding throughout the potentially different interfaces, and aid in the promotion of awareness between the two users. Coupling in many shared 2D systems involves user interfaces with relaxed spatial constraints. That is, the location of objects on one user's display may be at another location on another's. Such user interface systems are problematic when telepointers are used to portray users' activities within the workspace. A telepointer allows one user's mouse pointer to be replicated within the shared environment. As a user moves or points to an area within this environment, a representation of this user's pointer is replicated onto the other users' views. However, in a relaxed WYSIWIS scenario, when a user points at an object in their interface its coordinates may not be the same for all other users' interfaces, see figure 13. Hence for some users, the original user's telepointer may be gesturing at empty space. Thus the power of telepointing is lost, and produces meaningless gestures across the shared displays.

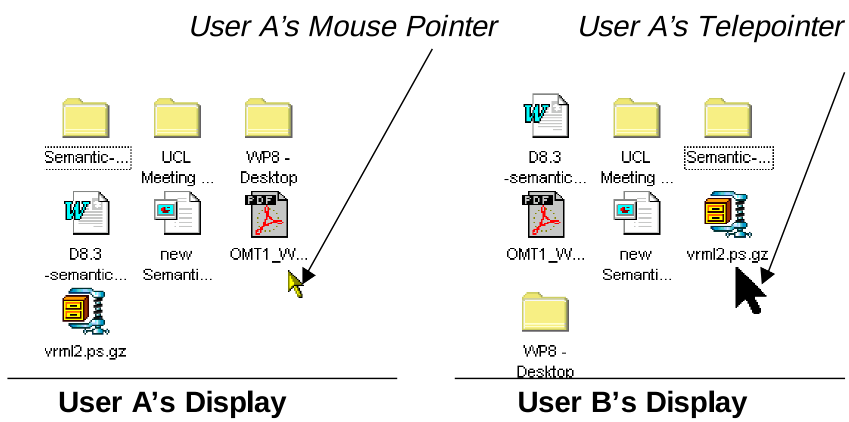

Figure 13 : Problems of telepointers in relaxed WYSIWIS views

The re-coupling mechanism can be used to implement a non-spatial telepointing system. Rather than appear as a 2D pointer, the telepointer may be realised as an increased brightness (colours tend to white) in the objects that are visually important to other users. These objects may be gained by choosing those semantically close to the originating user's pointer by some considering awareness based on the window hierarchy (Rodden 1996). For example, consider a simple control panel interface with three buttons. Each user's subjective view of the panel can be represented by the function:

Subjective_View (tom, O)={ (open, 0), (*close*', 0), (cancel, -1)}

Subjective_View (jon, O)={ (open, 0), (close, 0), (cancel, 0)}

This states that Tom has the *open* button as normal, a different representation of the *close* button and an invisible version of the *cancel* button (in this case –1 represents: invisible). The differences in Tom's display are encoded using the subjective modifier

SubjectiveModifier(tom) = {S(close,close'),S(cancel,cancel')}

Jon's display resembles the default presentation. Both subjective interfaces as depicted in figure 14.

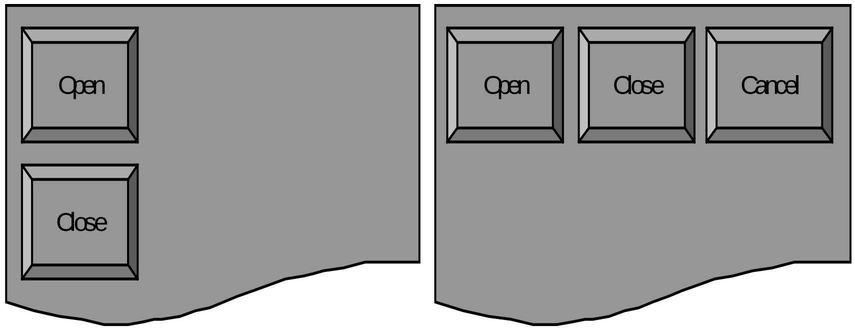

Figure 14. Two subjective interfaces.

The awareness map in this case considers these two users in close awareness given the common objects in both users' displays, and the proximity of their pointers to the closest object. For example, if two user's are pointing near the open button in each of their displays (which may not be spatially linked in the two displays), their awareness values will rise.

Hence, as the user Tom becomes more aware of the user Jon, their user interfaces become more re-coupled. The re-coupling function (display) defines the changes to each interface such that invisible objects appear, and objects that other user's are aware of become brighter. Hence, in Tom and Jon's interfaces the open and close buttons appear brighter. The cancel button also appears a little brighter in Jon's display, but in Tom's display the (once) invisible button appears in a faded (dark or near background colour) fashion.

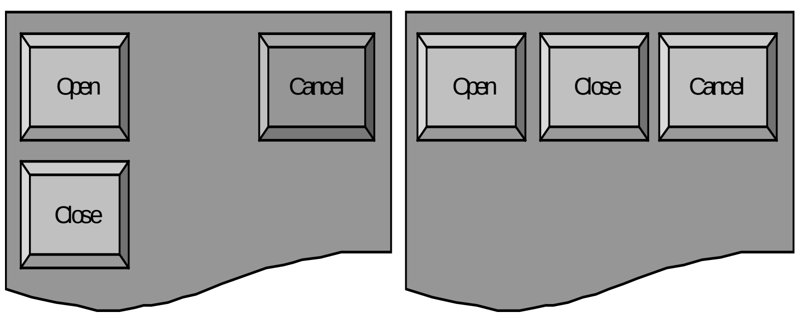

Figure 15. The re-coupled subjective interfaces.

Extending re-coupling – Personal Points of Presence

Previously in this chapter, we have described how a re-coupling model can be implemented to aid in the promotion of awareness among users, and consequently aid their collaborative activities. In addition to utilising the interface definition of other users to influence a user's interface, the re-coupling model can be extended to include the notion of 'Personal Points of Presence. A Personal Point of Presence is a representation of a user in the shared environment. The Personal Point of Presence is an intersubjectivity mechanism and need not be visualised, but it can be used to remember a user's awareness mapping.

For example, consider a 3D visualisation of the human body. The geometrical information available is large and much of the details will be of near irrelevance to most specialists. Hence, each medical specialist is provided with a subjectively tailored view of the body to reflect their expertise. For example, an osteopath visualisation will contain only the skeletal structure, whereas a surgeon's visualisation will highlight the liver and its closely related components.

Obviously the re-coupling model can be used to re-configure two doctor's views of body (if they are collaborating, or are merely working in the same area) to provide them with common subjective views, as described in the case of the electronic and hydraulic engineers earlier. However, consider the scenarios where the 3D torso is used for teaching or where it is used a reference frame for a group of specialists to discuss a complex medical problem. In both cases, each of the specialists may wish to annotate the 3D model with diagrammatic or textual information. Again, subjectivity can be used to reduce the visual complexity of the scene. However, each of the annotations are only visible to specialist who created them, and would normally be invisible to the other doctors (or students in the teaching scenario) unless they re-configured their subjective view specification.

A Personal Point of Presence can be used here to represent the doctor who is not present. Each of the annotations that a specialist makes may have a Point of Presence attached to it. Then, when that doctor is not present their annotations will still appear in the displays of the other doctors when they are investigating the same location of the body. This would allow the specialist to work asynchronously in the shared 3D environment, by annotating elements relevant to them, while also being presented with the annotations of the other doctors.

The above example describes the scene in a 3D e-scape. However, the above description can equally be applied to a shared 2D application, as both subjective views (relaxed-WYSIWIS) and the re-coupling mechanism can be used in this manner.

Conclusions

This chapter has considered the issues involved in users dynamically sharing interfaces that have been subject to some form of tailoring. We have argued the need to consider the provision of some form of intersubjectivity and suggested a means by which tailored user interfaces can be dynamically recoupled based on awareness. A general recoupling model was developed that build upon our previous work in subjectivity and a general purpose model of awareness. This model was also implemented and demonstrated.

We believe that the model and framework is sufficiently general that it can be extended to consider different views of shared information. This allows us to develop tailored presentations of CSCW systems that can be recoupled in terms of different users awareness of each other. The model leaves the particular views of the shared information and awareness mechanism used to drive the recoupling to the discretion of designers as these are very dependent on the particular semantics of the application.

References

Bowers J. O'Brien J. Pycock J. Practically Accomplishing Immersion: Co-operation in and for Virtual Environments. Proc CSCW 96. Boston MA USA. ACM Press. 380-389.

Dewan P. (1990). A tour of the Suite User Interface Software. *Proc. of the 3rd ACM SIGGRAPH Symposium on User Interface Software and Technology*. October 1990. 57-65.

Carlsson C, Hagsand. (1993) "DIVE - Multi-User Virtual Reality System", VRAIS '93, IEEE Virtual Reality Annual International Symposium. 394-400

The DIVE Homepage. http://www.sics.se/dce/dive/dive.html

Division Ltd. (1993) "dVS Technical Overview, Version 2.0.4. (www.division.com)

Ellis C, Gibbs S, Rein G. Groupware: Some Issues and Experiences. CACM 34, 1 (1991), 38-58.

Susan R. Fussell and Robert M. Krauss (1992) Journal of Personality and Social Psychology, 62, 378-391.

Greenberg S. Personisable Groupware: Accommodating individual roles and group differences. Proc ECSCW'91, Sept 1991. Kluwer Press. 17-31

Greenhalgh C., Benford S., MASSIVE: a Collaborative Virtual Environment for Tele-conferencing, ACM transactions on CHI, March 1995

Gutwin, C. Greenberg, S. Workspace Awareness for Groupware. CHI'96 Conference Companion (Vancouver 1996), 208-209.

Munson J, Dewan P. "A Flexible Object Merging Framework". Proc CSCW 94. Chapel Hill NC USA. ACM Press 1994. 231-242

Patterson, J. F., Hill, R. D., Rohall, S. L., Meeks, W. S. (1990). Rendezvous: An architecture for synchronous multi-user applications, *Proceedings of CSCW '90*, October 7-10, Los Angeles, Ca., ACM, 1990. 317-328.

Rodden T. "Populating the Application: A Model of Awareness for Cooperative Applications". Proc. CSCW 96. Boston MA, USA. November 1996. ACM Press. 87-96.

Roseman, M., Greenberg, S. GroupKit. A groupware toolkit for building Real-Time Conferencing Applications. Proc of ACM CSCW '92. ACM Press. 43-50

Smith, G., Rodden T. SOL: A shared object toolkit for cooperative interfaces. Int. J. Human-Computer Studies (1995) 42. 207-234.

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

Smith G., Mariani J. "Using Subjective Views to Enhace 3D Applications". ACM Symposium on Virtual Reality Software and Technology (VRST '97). ACM Press. Swiss federal Institute of Technology (EPFL), Lausanne, Switzerland. September 1997. Pages 139-146.

Stefik M, Foster G, Bobrow D, Kahn K, Lanning S, Suchman L. Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. CACM 30 (1), Jan 1987, 32.47

Appendices

Appendix One: Conceptualising and Explicating ‘Presence’

Craig Murray

Manchester University

Introduction: Conceptualising ‘Presence’

The concept of ‘presence’ has received scant attention in psychology and philosophy in any explicit sense. With the advent of new technologies, this previously supposed taken for granted state of consciousness has become an area for research and discussion. While there are over two decades of research on how the social presence of others are communicated to each other via various media, more recently there has been an interest in how someone may experience their own presence within a synthetic environment. Questions which arise from this interest include how someone might feel present within a space created by synthetic media, or how might someone be made present to others.

The concept of presence in real environments has only received implicit attention. With no full, empirically grounded understanding of the quality and qualities of presence, the use of the concept is rendered, seemingly, confused. It is a frequent comment that different writers mean different things when they employ the term. It will be argued later that ‘presence’ can be better understood by looking at the phenomenal properties of the experience it is said to demarcate. For now, we will examine how presence has been conceptualised by various disciplines and theorists, and attempt to elucidate the similarity in these different ways of thinking about presence.

The Oxford English dictionary (1935) includes some useful definitions, which are commonly employed by presence researchers. These definitions characterise presence as:

> ‘the fact or condition of being present; the state of being before, in front of, or in the same place with a person or thing; being there… denoting the actual person (or thing) that is present… [sometimes] used with a vague sense of the place in front of a person, or which immediately surrounds him.

However, other dictionary definitions of presence recognise at least two modes of presence: that of one’s own presence within an environment (as above), and that of perceiving the presence of others within an environment. Presence:

> ‘Being present in a place, *your presence is required*. …**3.** a person or thing that is or seems to be present in a place, *felt a presence in the room*.’ (Oxford Paperback Dictionary 1979)

While it is common to think of presence as a person’s subjective experience of being within a particular place or environment, in what follows I will make a broad distinction between a) the presence of oneself within an environment, and b) the presence of others to oneself within an environment.

Presence has been investigated almost entirely in relation to virtual environment and teleoperation technology. Much of this research, while of interest, has been largely speculative and theoretical. This work is exemplified by the studies of Held and Durlach (1992), Sheridan (1992a), Steuer (1992) and Zeltzer (1992). The few empirical studies that exist have investigated the influence of a number of factors, for example representation systems (Slater et al. 1994), a virtual body representation (Slater et al 1994) and pictorial realism (Welsh et. al. 1996) upon feeling present within a virtual environment.

Distinctions have been made between presence, telepresence and virtual presence (see Sheridan 1992b). Presence has been discussed above. Telepresence is reserved to refer to the feeling of being physically present at an ‘actual’ (i.e. ‘real’) remote site, conveyed by technological apparatus (e.g. real-time visuals via head-mounted displays (HMD’s)). Virtual presence refers to the experience of feeling present within a computer-generated environment. These teleoperation and virtual environment researchers are primarily interested in the subjective experience of one’s own presence within an environment. As such, they neglect the experience of perceiving the presence of other social actors. As will be discussed later, these researchers also tend to use experimental, quantitative methods to ‘measure’ presence within environments. It will be argued later that presence, as a phenomenal experience, cannot be adequately accessed via likert style questionnaires, psychophysical measures (such as heart rate), or purely quantified observable behaviour. Rather, it is a concept more readily available to qualitative methodologies, e.g. interviewing and disciplined self-reflection which lend themselves to a phenomenological approach.

Presence of Oneself Within An Environment

The Fluidity of Presence

The difficulty in conceiving and explicating ‘presence’ arises because of its lived nature. Like ‘being’ and ‘existence’, ‘presence’ is a metaphysical condition:

> ‘Perception is an oceanic feeling of a presence to which I adhere without distancing myself… The world is not present to me, I am present within the world, within presence itself. But for all that, presence is not being (*l’etre*), but rather it belongs to the being (*l’etant*) that is, and that is present …we remain captive to a metaphysics of presence.’(Dufrenne 1987:53)

The complexity of presence remains unappreciated. Researchers generally conceive of it as a stable experience, both within the real and virtual environments. This is apparent in their attempts to quantify the phenomena, ‘measuring' people’s perceived and observed presence. However, even in the real environment, presence is not an all or nothing phenomena. We can be within-the-world, or we can retreat, into sleep, dreams, or other worlds (daydreams, hallucinations). We can remain immersed within the experience, or we can make it an object of conscious thought. With reflection on the experience of presence, a distance between the person and the phenomena is introduced:

> '… the fullness of presence…never fails to be undermined as the subject, in order to constitute himself, breaks away from it and institutes presentation. Then it is that the subject separates himself from the object, alterity becomes objectivity, the real is won over.' (:58)

Therefore, in the following discussion of 'presence', it should be borne in mind that it is, even in the physical world, a fluid experience, with no stable, static, quantifiable moment.

The Body as Locus of Presence

Phenomenology reveals that the body in the world is both foreground and background. It constitutes our locus, our point of view, so that we are 'here' rather than 'there'. Yet, at the same time, the body recedes from conscious reflection. At once a holistic sense organ, and yet an assemblage of sense apparatus, the body recedes from awareness in its perceptual activity. It is the very disclosure of the world that the sense organs provide which leads to their mindful demise. The body recedes in activity but is always implicitly present and known. The body's 'presence is fleshed out' by its sensorium (Leder 1990), in particular proprioception frames the body. Similarly, Straus (1966) speaks of how the,

> 'morphological and functional characteristics belonging to this body not only ground our-thus-and-thus constituted live bodiness but also our very own presence, our unique existence' (:307).

In his philosophical essay 'Where am I?' Daniel Dennett (1978) addresses issues pertinent to the issue of personal presence. Dennett muses over a hypothetical case in which his brain is separated from his physical body, but is able to communicate via complex radio signals. Given that his physical body and brain are now able to exist in separate physical places, the question for Dennett is 'where am I?' He considers three possibilities: (1) that he is where his body is, (2) where his brain is, and (3) wherever he thinks he is.

The first of these is found dissatisfactory because a simple thought experience disproves it – i.e. if his brain was transplanted into another body he would still retain his own personhood within the locus of another body. The second possibility is found 'unappealing' because of an easily imagined situation whereby Dennett would be able to look upon his physically separated brain via the physical body. The third possibility suggests that 'At any given time a person has a *point of view*, and the location of the point of view (which is determined internally by the content of the point of view) is also the location of the person.' (:314).

While point of view is related to personal location, Dennett finds this 'an unclear notion'. He gives the example of the Cinerama viewer,

> 'who shrieks and twists in his seat as the roller-coaster footage overcomes his psychic distancing… Has he forgotten that he is safely seated in the theater? Here I was inclined to say that the person was experiencing an illusory shift in point of view.' (:314-315).

Dennett goes on to consider the experience of someone within a teleoperator system. These workers are able to interact with the remote environment, via feedback controlled devices they can 'feel the heft and slipperiness of the containers they manipulate with their metal fingers'. While they 'know' where they are, it is 'as if' they were inside the chamber into which they look. Via conscious thought, Dennett suggests, might they not be 'transporting *themselves* back and forth'? (:315).

Loomis (1992) identifies distal attribution as a closely related phenomenon of 'presence'. This is whereby perceptual experience is referred to external space, beyond the sensory organs through which they originate. He goes on to talk of the normal division of the phenomenal world into self and non-self. The body acts as a locus of the self, and its boundaries (the phenomenal body) and private subjective experience (such as thirst and hunger) are properties of the self. While the phenomenal and physical body do not always coincide (such as the amputees' phantom limb phenomena), there is a close match. Our locus of presence, then, is tied up in our phenomenal body.

Loomis considers the conditions under which distal attribution occurs. Presence and distal attribution beyond peripheral devices are similar phenomena. Presence occurs when sensory data supports only one interpretation of 'being somewhere'. Distal attribution to another site takes place when sensory data represent a remote location, as well as the peripheral devices that connect the observer to it.

Absence and Lack of Presence

It might be argued that the opposite of presence is absence. However, absence does not appear to have any perceptual component, although it may have other experiential detail, such as feeling 'left out', 'homesick', etc.

Leder (1990: 1) considers 'absence' with relation to the body,

> 'While in one sense the body is the most abiding and inescapable presence in our lives, it is also characterised by absence.'

So, while a phenomenological analysis can characterise experience in a particular manner, if we are to capture experience itself we might better look at 'lack of presence'. In any case, for our purposes, I would suggest we discuss 'lack of presence' rather than 'absence' in relation to the experiential aspects of personal presence.

We can ask how might a lack of presence be experienced in a real environment. A lack of presence can be experienced as ambiguous. While I do not feel present in any other environment than that in which I am located, this does not constitute a lack of presence in those environments from which I am absent. The experience of a lack of presence is an ambiguous amalgam of presence/non-presence within a particular environment of which we are in some manner sensorially aware.

As an instance of this ambiguity we can consider a concrete example the environmental experience of suddenly and profoundly deafened individuals. I take auditory experience as one example of our presence/presenceless experience. Ramsdell (1978) (previously discussed in relation to presence in Gilkey and Weisenberger 1995), although not explicitly a phenomenological reading, provides such insights into the world of these individuals. He asks us to imagine the world of the deafened person. This is a world in which rainfall makes no noise; in which crowds are silent; a world in which 'everything moves with the unreality of pantomime' (Ramsdell 1978: 500).

Ramsdell argues that it is the sounds of this most basic, routine, and pre-reflective level which maintains our experience of involvement in a living world, and of our being alive within it. It is precisely because this level of hearing is so familiar and pre-reflective that we remain unaware in our daily lives of its vital importance. So innocuous and hidden is it from our immediate awareness that even when deafened we remain unaware of the loss: the deaf person 'knows' only that they experience the world as dead. Ramsdell uses the loss of this hidden, pre-reflective nature of this coupling with the environment to explain the depression of many deafened adults who feel 'caught in a dead world'. Hence, this auditory dimension plays a pre-reflective role in our presence in the physical world.

While phenomenological literature on the experience of sound is rare (Ihde 1976), literature regarding the absence of sound is rarer still. A number of studies have induced hearing loss experimentally. Eriksson-Mangold and Erlandsson (1984) induced hearing loss by occluding the ear canal in a group of normal-hearing participants. In this study, participants experienced hearing deficiencies by two different occlusion methods: one group used plastic mass, another used expanderplugs. Participants spent a whole day with their ears occluded. Findings of the study included feelings of isolation, of not taking part in the life around them: *the surrounding world did not appear quite real*. Once again, the lack of presence that sudden hearing loss precipitates is evident.

As a more mundane example of how presence is undermined by a lack of sound, we can consider the development and evolution of architecture. The proliferation of glass in the twentieth century has created a rupture between person and environment. In Murray Schafer's (1989: 97) words, glass has 'shattered the human sensorium', dividing the perception of a visual world from its *own* aural, tactile and olfactory accompaniments.

Schafer focuses on the impact of glass, on the 'soundscape'. The window, argues Schafer, frames outside events in phantom-like silence. The increased thickness of glass has underscored a 'here' and a 'there', and results in 'a fission of the senses'. For Schafer, glass has resulted in insulated spaces which crave reorchestration. He compares the experience of looking out through a window to a ‘movie set’, with radio for a soundtrack. He gives his own experience of travelling through the Rocky Mountains by train, with ‘schmaltzy music’ coming through the loud speakers. He recalls thinking that what he was actually experiencing was a ‘travelogue movie’, ‘we are *not here* at all.’ (Schafer 1992)

Straus (1966: 19-20) has discussed hearing and the phenomenal ties it has between the perceiver and their world, using the example of a silent film:

> ‘When a film is shown without music, the pictures appear at a different remove - usually, remote; they are marionette-like and lifeless. We lack contact with what is being represented, which glides in front of our eyes in a spiritless, barren manner. We are spectators at, not participants in, what is occurring’.

Ihde (1976: 82-83) has further discussed this phenomenon which Straus describes. He notes that:

> ‘the fullness of *auditory presence* is one of an “animated” liveliness. ...when sound is added to abstract figures, they come “alive”.’ [italics mine]

But just as the addition of sound achieve this, the sudden absence of sound can ‘disembody a scene’:

> ‘At the instant of the disappearance of animating sound, the scene becomes eerie, a moving tableau which becomes more abstract and distant.’

Similarly, Lindsay (1988: 64) relates the account of one person with hearing loss which echoes Straus’ example, and is reminiscent of Schafer’s analysis too:

> ‘I was to come to terms with a world viewed largely through a plate glass window where other people live, laugh and suffer and barely know of my existence. ...I began to mistrust my own perception of the world and the people around me. How could I be sure of my impressions when I couldn’t hear?’ (:64)

Presence and Realism

A number of researchers (e.g. Welsh et. al. 1996) have focussed on ‘realism’ and its effect upon evoking a sense of presence within synthetic environments. As characterised by Lombard and Ditton (1997), this realism refers to

> ‘the degree to which a medium can produce seemingly accurate representations… that look, sound, and/or feel like the “real” thing.’

Lombard and Ditton make a useful distinction between social and perceptual realism. However, VR developers show more of a concern with the latter, as demonstrated in quests for photorealism.

Presence and Transportation

In this conceptualisation of presence, Lombard and Ditton (1997) identify three types of transportation involved: a) ‘you are there’; b) ‘it is here’; and c) ‘we are together’. Here we discuss (a), leaving (b) and (c) until later.

a) ‘You are there’ – research literature has pointed towards story telling, such as traditional oral narratives and later written ones, as having the potential to ‘transport’ the listener/reader to a different time and place. Lombard and Ditton point to the familiar 'welcome back' of TV presenters after a commercial break as an example of this type of transportation.

With regards to TV, Lombard and Ditton cite Kim (1996) as defining presence as a 'feeling of being part of the phenomenal environment created by television and not being part of the physical environment surrounding the viewer and the television set' (:27).

Presence as transportation is also invoked for virtual environments. So, there may be a 'suspension of dis-belief that they are in a world other than where their real bodies are located' (Slater and Usoh 1993: 222), they may feel that they are actually 'present in the environment generated by the computer' (Sheridan 1992: 120).

Presence and Immersion

Perceptual immersion may be characterised as 'the degree to which a virtual environment submerges the perceptual system of the user' (Biocca and Delaney 1995: 57), that is, the number of the user's senses which are inputted. This echoes presence in real environments, as discussed earlier, for instance when the sudden absence of hearing decreases presence in the real world.

One line of argument proposed is that for a sense of 'presence' in the VE of the computer, then the VB must closely resemble (both visually and sensorially) the body of the user (body as locus of presence). Some ask how the 'geometric mappings' of the body within the virtual and physical environments, relative to each other, contribute to a sense of presence (Sheridan 1992a). The general consensus is that 'identification, and therefore telepresence, would be increased by a similarity in the visual appearance of the operator' and the virtual body (see Loomis 1992), which can be contrasted with the speculative discourse on the polymorphous potentiality of VR (Murray 1996a).

Many writers (Held and Durlach (1992), Sheridan (1992), Steuer (1992)) argue that a sense of presence will generally increase with an increase in the number and fidelity of sensory inputs. As an example of one such input in VR, and it's relation to presence, we can return to our earlier discussion of presence as experienced by suddenly deafened adults. Consider the example of auditory displays. One of the closest phenomenological characterisations of presence that the VR community has made is in an article by Gilkey and Weisenberger (1995) (though see Zahorik and Jenison 1998 for an explicit attempt to characterise presence phenomenologically). These researchers look at the experience of sound in VR applications, and of it's importance for a sense of presence within an encompassing virtual environment. They illustrate this by drawing on the experiences of suddenly deafened adults, as reported by Ramsdell (1978).

Within this report, Ramsdell (1978) discusses the psychological experience of sudden and profound hearing loss. The world of these deaf observers seemed 'dead' and Ramsdell stresses the role of hearing in being 'part of' and 'coupling' with the environment. Gilkey and Weisenberger use this piece of research to argue that, even with the effective simulation of all other sensory detail, a VE without effective facilitation of auditory information will fail to provide an experience of ‘presence’. Their contention is that (faithful) auditory information is crucial for a sense of presence in a VE.

The experiences of these suddenly deafened adults, argue Gilkey and Weisenberger (1995), have parallels with the experiences of users of virtual environments. With sudden immersion into many VE’s, the person is plunged into deafness. Following on from Gilkey and Weisenberger’s (1995) theoretical exploration of the auditory dimension on the feeling of presence, Hendrix and Barfield (1996b) examined this empirically, finding that auditory information does have a facilitative affect.

I would argue that this is one reason why the ocean metaphor is so commonly employed in explaining the experience of virtual reality, and in the design of VE’s. Bricken (1992) pursues the ocean metaphor in their description of different VE’s. So, viewing 3-D graphics on a screen is compared to looking into the ocean from a glass-bottomed boat: ‘We see through a flat window into an animated environment; we experience being on a boat.’ Using a stereographic screen to look into a virtual world is compared to snorkelling: ‘We are at the boundary of a three-dimensional environment, seeing into the depth of the ocean from its edge; we experience being between at the surface of the see.’ Finally, using a steroscopic HMD is compared to ‘wearing scuba gear and diving into the ocean. ...We’re there’ (:364).

We can speak of the sense of immersion in current VE’s as incomplete. There is an inability to become fully subject, or present, in a VE. While the phenomenal experience and properties of presence are yet to be explicated in a phenomenological analysis, it is possible to identify several likely contributing factors: visual resolution; failure to accommodate all sense modalities; discrepancies between the physical and the virtual; time and aural lag, etc. The state of semi- or quasi-presence can be characterised by a phenomenological analysis. In a sense a phenomenal ‘bubble’ or ‘pocket’ surrounds the immersed participant: they are part of, yet separate from, the VE.

Meredith Bricken’s analogies of immersion are particularly telling: using a stereoscopic HMD is seen as like wearing scuba gear and diving into the ocean - ‘were there’. If we consider the experience of water immersion, it to has the phenomenological ‘bubble’ characteristic: sounds are dampened or absent, as in most VR applications (see Gilkey and Weisenberger 1995); movement is laboured and slow; and both involve peripherals which mediate the environment.

Current use of immersive VE’s is characterised by an experience that is at once inclusive and excluding. We are able to act and perceive in the VE, but these experiences are characterised by a feeling of distance, which is not spatial or temporal, but rather existential. We are ‘removed’, not in space and time, but rather, as a rupture between person and (simulated) environment.

The ocean metaphor in VR design has been taken to an extreme, that is in a literal sense, in the system ‘Osmose’, created by Char Davies and co-workers. In this system, Davies creates the sensation of ‘gently floating’ via navigation techniques of balance and breathing. These techniques were inspired by, among other things, Davies’ experiences of scuba diving:

> ‘Sensors were used to measure the tilt of the immersant’s spines, and the expansion and contraction of their chests. This allowed us to move immersants horizontally through the virtual space in whatever direction they tilt, and move them upwards when they fill their lungs, and downwards when they empty their lungs (similar to scuba diving). ...Immersants are, as in scuba diving, discouraged from reaching out and touching things in Osmose (we intentionally did not track their hand positions).’ (Davies and Johnson n.d.)

Presence as Transparency and Incorporation

The transparency of visual, kinaesthetic, aural and other displays are considered a key determinant of feeling present in a VE (Held and Durlach 1992). The continuity of the-body-as-I in VR calls for an assimilation of both technological peripherals and the virtual (re)presentation (Murray 1996b). In the following, I discuss the phenomenological body image and the embodiment of physical artefacts, namely tools. This discussion is pertinent to the discourse on the desirability of transparency in VE displays.

The ‘Body Image’ and Tools

In psychology, the body image concept has often been defined as the way in which one’s own body is perceived and/or conceived. Perception, cognition, attitudes, feelings and ego have all been incorporated into psychological definitions of ‘body image’. Psychological interpretations have used the term to denote ‘knowledge’ of one’s own body, present at the borders of consciousness, or functioning unconsciously’ (Tiemersma 1989: p.1).

Phenomenological approaches to the body image have concentrated on the perceptual aspect of body image. Corporeal experience provides us with a 'body image', a phenomenological understanding of our bodies extended in space and time.

In Heidegger’s (1962) analysis of tool use, he uses the example of a hammer. For a tool to be ready-to-hand it must, in Heidegger’s terms, ‘withdraw’. In so doing, the tool becomes the means rather than the object of the experience (Ihde 1990). This has a parallel with the experience of the body itself: Sartre (1970) sees the body as the perpetually ‘surpassed’. The tool itself is also surpassed as it withdraws into the architecture of the body, forming what Ihde (1990) terms ‘an embodiment relation’.

Various phenomenological authors have invoked the example of the blind person’s cane (Heidegger 1962; Merleau-Ponty 1970; Schilder 1978; Ihde 1990). For Merleau-Ponty, the cane for the blind man is no longer an object, but an extension of the realm of the senses. With the cane as a ‘familiar instrument’, touch is experienced at its end point (‘its point has become an area of sensitivity’ (:143)) - rather than at the hand.

This body familiarity (or ‘habit’) of the tool is the result of being ‘transplanted into them, or conversely to incorporate them into the bulk of our own body’ (:143). Merleau-Ponty (1970) speaks of the blind man’s cane as ‘an instrument with which he perceives. It is a bodily auxiliary, an extension of the bodily synthesis’ (:152). Similarly, Leder (1990) contends that ‘In its use of tools and machines the body supplements itself through annexing artificial organs’ (:30).This incorporation of the tool into the body gestalt is what Leder (1990: 34; see also Grosz 1994) refers to as a ‘phenomenological osmosis’, ‘the body allows instruments to melt into it’ (Kujundzic and Buschert 1994: 207-208).

We can further illuminate our understanding of the above with Ihde’s (1990) phenomenology of technics (defined as ‘the symbiosis of artefact and user within a human action’ (:73)). Ihde employs the example of eyeglasses, and expresses the optically mediated world when wearing them with his trivariate separation of ‘I-glasses-world’. However, in wearing eyeglasses, the glass is ‘transparent’. The weight of the glasses on the ears and the bridge of the nose also ‘withdraws’. In so far as we take technologies into our experiencing by perceiving through them, the technology becomes embodied. ‘I-glasses-world’ becomes ‘(I-glasses)-world’. Ihde maintains that the same process occurs for the hearing aid (for hearing) and the blind man’s cane (for tactile motility).

The Presence of Others to Oneself Within an Environment

Presence and the Look of the Other

When we are looked upon, this reveals both the presence of the other and oneself in an environment. Biocca (1997) suggests that ‘The perception of the other is the empathetic simulation of internal states of another “if we were in the space” over there.’

According to Sartre (1957), being-seen-by-the Other is the truth of seeing-the-other. Experience and being-looked-at reveals to oneself the Other-as-subject, a condition of our own being-as-object.

One interactive video installation (currently at ZKM, Karlsruhe) is Luc Courchesne’s ‘Portrait no1’. The user is presented with the image of a woman, ‘Marie’. The image is at face-to-face level, and the user is able to engage Marie in ‘conversation’ via choosing questions from a text menu. Apart from the artistic intentions of the artist, we can look at this piece for the phenomenal aspects of the user’s interaction. Because Marie is at the same height as the user, as well as appearing at a reasonable distance, the interaction is evocative of two people in conversation. ‘We can face her, flirt a little, and make eye contact.’ (Scharz 1997: 94)

> ‘You may try to get her attention: when selecting "Excuse me..." on the display, *Marie suddenly stares at you*; then, selecting "Do you have the time?", "Are you staring at me?" or simply "May I ask you something?" starts a conversation that will develop according to visitor's curiosity or Marie's moods.’ (italics mine)

Although the answers to any questions are, ultimately, pre-arranged, the user experiences an engagement of two social actors. The piece in no manner passes a ‘Turing test’, however the user does experience ‘the look of the other’. They feel looked upon. Marie has a presence which arises through practical activities – asking and answering questions. The practical activity of a conversation, coupled with the gaze of Marie, provides us with the presence of an other as well as our own presence. The emotional experience of ‘my-being-looked-at’ reveals the other as a subject (Sartre 1957; Eeke 1975).

Ghostly Presence

Another form of presence which is often referred to is a ‘ghostly presence’. This means more than the presence of ghosts, but calls attention to a particular form of presence. In contrast to the vividness, fleshliness of living bodies, spiritual, ghostly bodies are often characterised as translucent, transparent, invisible, or ephemeral. A ghostly presence may be conveyed by rushes of cold air, sudden changes in temperature, feelings of being watched, hearing voices or other sounds, (corroborating) accounts by other people, or seeing a ‘loose apparition’.

Varnado (1987) discusses the ‘numinous’ as the feeling of the supernatural. Although the Oxford dictionary defines numinous as ‘indicating presence or influence of a god’, Varnado, following Otto (1958) refers to a ‘sense of haunted presence’, which may be associated with landscapes, mountains, streams, and other impressive natural objects.

Feelings of eeriness, or a place as haunted may on occasion become explicit recognition of a transcendent something ‘a real operative entity of a numinous kind’ (Otto 1958). This entity can be daemonic as well as holy. As argued by Varnado,

> ‘The very nature of the numinous experience is such that it cannot be perceived unless the subject feels that there is indeed something objective before which he is a creature.’

Presence and Social Richness

Lombard and Ditton refer to social presence theory (Short et. al. 1976) and media richness theory (Rice 1992) as the basis of this conceptualisation of presence. These theories derive from the study of different communication media (from the telephone to the computer). Communication media are said to differ, amongst other things, with regards to whether they can ‘transmit the social, symbolic, and nonverbal cues of human communication’ (Rice 1992: 452).

The social presence theory suggests that the absence of visual, aural, and other information characteristic of FTF contact makes people less aware of their audience. Conversely, the presence of a communicator would be lower for ‘narrow’ bandwith media such as e-mail. For our purposes, these theories suggest that the presence of others can be made manifest to a lesser or greater degree due to the particular medium and the activities involved. This conceptualisation of presence has much in common with current VR discourse on the presence of self within an environment. They both focus on the technology, in particular with the number of sensory channels catered for, along with their fidelity.

Presence and Transportation

This conceptualisation of presence there involves three types of transportation: a) 'you are there'; b) 'it is here'; and c) 'we are together'. (a) was discussed earlier. Here we will discuss (b) and (c).

b) 'It is here'. Instead of feeling as though their presence has been transported somewhere else (other than which they started off in), the media user may have the impression that objects and people are brought from another place to their environment. Television is one example of this, where pictures, people, and events are brought into the home of the viewer. Lombard and Ditton (1997) cite examples of early cinema goers who reportedly reacted with panic at the site of an oncoming train.

c) 'We are together' (shared space). This type of transportation characterises videoconferencing as well as some multi-user virtual environments. In videoconferencing this might involve the feeling of sharing space with other people. Whereas virtual environments may also involve this, on desk-top applications in particular it may be a general feeling of being together. A telephone conversation between two or more people may also involve this kind of feeling.

Having presented the different ways in which presence has been theorised and conceptualised, we now move on to the proposed determinants of presence and its quantification.

Theorised Determinants of Presence

Steuer (1992) discussed technological variables that they believed influence telepresence. These include vividness and interactivity. Vividness refers to the representational richness of a mediated environment, which can vary in terms of sensory breadth, i.e. the number of sensory dimensions which can be simultaneously presented, and sensory depth, the resolution within each of these perceptual channels. For example, Zeltzer (1992) asks what level of photorealism is required for a feeling of presence.

Interactivity refers to the extent to which people can participate in modifying the form and content of a mediated environment in real time. This can vary in terms of speed, range and mapping. Range refers to the rate at which input is assimilated into the mediated environment. Range refers to the number of possibilities for action at any given time). Finally, mapping refers to the ability of a system to map its controls to changes in the mediated environment in a natural and predictable manner

Similarly, Held and Durlach (1992) discuss sensory factors in relation to presence. High resolution and a large field of view are seen as important for the transparency of the display system, and, therefore, for facilitating a sense of presence within the presented environment. Additionally, information across the senses should be consistent, a wide range of sensorimotor interaction should be provided, there should be high correlations between the body movements of the user and those of a slave robot (or virtual body), there should be similarity of appearance between user and robot/virtual body. Lastly, Held and Durlach suggest telepresence may increase with familiarity/use of the system.

Sheridan (1992) provides 3 determinants of presence: (1) extent of sensory information; (2) control of sensors in environment (e.g. head turning); and (3) ability to modify the environment (e.g. move objects). With the optimalisation of all three of these, ‘perfect presence’ is said to occur.

Techniques for ‘Measuring’ Presence

The following is a brief overview on how researchers have suggested and attempted to ‘measure’ a person’s sense of presence within a (real, teleoperation or virtual) environment.

Held and Durlach (1992) expressed a concern with quantitative measures of telepresence, involving standardised scale-construction techniques, psychophysical and physical techniques (e.g. startle response). Schloerb (1995) proposed a quantitative measure of telepresence involving ‘both objective and subjective measures’. It was proposed that a measure of subjective telepresence should involve a psychophysical test, and then analysed using signal detection theory. Sheridan (1996) identified three methods which have been proposed to measure presence

(1) reflexive responses (Held and Durlach 1987; Loomis 1992) such as ducking before looming objects, etc.

(2) subjective category rating scales, characterised by verbal description, and

(3) discrimination: the (in)ability to discriminate between a real and a tele or virtual world. (see Schloerb 1995).

The closest these researchers have come to allowing participants to express their own experience of presence is with Likert-style questionnaires (see the work of Slater and colleagues). Here, on a scale of 1 to 7, participants rate their sense of presence in a VE. A similar measure was employed by Hendrix and Barfield (1996a,b), who asked participants:

> ‘If your level of presence in the real world is “100”, and your level of presence is “1” if you have no presence, rate your level of presence in this virtual world.’

And:

> ‘How strong was your sense of presence, “being there,” in the virtual environment?’

In the above study the wording of the survey questions were deliberately made simplistic to avoid individual differences in the interpretation of the questions. In addition, in keeping the questions brief and nonexpansive it was hoped that added sources of biases in the questionnaire construction would be avoided.

The above methods, I would argue, are inadequate as the sole means to research presence. ‘Presence’, as a phenomenal experience, can be better addressed by returning to how it is lived by people, rather than abstracted via quantitative approaches. In what follows, I present a phenomenological approach which may allow access to participants’ experience of presence, and of the qualities that constitute this experience. Consequently, this method is advocated for investigating presence within other, namely virtual, environments.

Qualitative Techniques to Explicate Presence

Zahorik and Jenison (1998) discuss presence as ‘being-in-the-world’, drawing on the work of phenomenology of Heidegger and the ecological psychology of Gibson. In a view of presence steeped in these traditions, the authors argue, presence can be seen as tied to one’s successfully supported actions in the environment:

> ‘Presence is tantamount to successfully supported action in the environment… the coupling between perception and action is crucial in determining the extent to which actions are successfully supported.’

Such an approach does not exploit the full potential of phenomenology. While it proposes a method for supporting presence, it does not explicate the properties of presence – how presence is lived. The following proposals set out how a phenomenological analysis can achieve this aim.

Employing the Phenomenological Approach

Phenomenological methods are employed which will uncover the essential meanings of lived experience by describing ‘psychological realities’ (Baker e al. 1992: 1357). Some researchers have argued that the only way for psychology to understand human behaviour and experience, as it is actually lived, is to gather together descriptive accounts of such experiences in their given contexts (e.g. von Eckartsberg 1986). Phenomenology offers a philosophical and methodological approach which is consonant with this aim. Through a process of deep, detailed description, we are able to bring to our awareness a pre-reflective life-world, revealed to us as psychological meaning (Vale and King 1978).

We have already seen that phenomenology concerns itself at least implicitly with profound psychological realities or meanings. When applied to psychological research, phenomenology is employed to tease out these meanings, to make explicit these taken for granted psychological realities.

The Qualitative Interview

The qualitative research interview elicits phenomenological of lived. The phenomenological interview proceeds in a ‘presuppositionless’ manner. It is not prescriptive, with an exhaustive list of pre-conceived questions. Rather, such an interview is responsive to the interviewee. While an interview protocol, a list of suggested areas of discussion derived partly from a theoretical sensitivity and preceding interviews, is necessary, they shouldn’t dictate the direction of the interview. The participant leads the way in the interview, while the interviewer carefully crafts questions grounded in the participant’s discourse.

Kvale (1984) identifies the features that characterise the qualitative research interview. Most importantly, it seeks to understand the lived experience of the interviewee, and therefore is centred on their life-world. This approach renders the interview qualitative, descriptive, and specific. The interviewer attempts to be presuppositionless and focuses the interview on certain themes. The interview is open for ambiguities and changes. The success of an interview depends, in large part, upon the sensitivity of the interviewer.

Applying Phenomenological Analysis

A number of researchers have attempted to explain what the phenomenological method involves, which although advocated by Husserl, was left largely unexplained as to how it should proceed.

Vague descriptions of what a phenomenological method involves still persist. For instance, Spinelli identifies three steps of the phenomenological method. The first is that of ‘epoche’, Husserl’s term for the bracketing of presuppositions, of having an openness to experience which allows it to be made manifest in its natural state. The second stage Spinelli identifies is to ‘describe, don’t explain’. This can be seen as a continuation of bracketing presuppositions, of allowing the experience to be made manifest. (However, Spinelli rightly notes that no description can ever be totally free of explanatory components). The last step identified by Spinelli is the ‘horizontalization’ or ‘equalization’ rule, whereby all items of description have equal value or significance. However, such a description of a method is still devoid of the needed procedural components.

Van Manen (1984) identifies four procedural activities in conducting phenomenological research. These are:

(a) turning to the phenomenon which is of interest and which commits us to the world;

(b) investigating the experience as it is lived rather than as it is conceptualised;

(c) reflecting on the essential themes which characterise the phenomena;

(d) describing the phenomena through the act of writing and re-writing.

Conducting the Analysis

Giorgi (1985) identifies four activities part of phenomenological analysis:

(1) First, reading through an entire description provided by a participant, in order to get a general sense of the whole statement.

(2) After grasping the whole, the text is read once more in order to discriminate "meaning units" from a psychological perspective and with a focus on the phenomena being researched.

(3) The third step involves an exploration of all 'meaning units' in which the psychological insight that each provides is extracted and expressed 'more directly'. One first discovers the relevant meaning unit (or theme), then, based upon a subsequent analysis, explicates its full import.

(4) The final step in Giorgi's phenomenological analysis involves the construction of a general description of the phenomena under consideration which encompasses all of the explicated meaning units.

Justifying Themes

Miller and Hodge (1997) argue that our interpretations of themes should include justifications of each theme with regards to certain criteria, including:

(1) What rationale was used to construct the theme? What types of reason can be given?

(2) What specific supporting evidence is used to justify the choice of this theme as being appropriate?

Straightforward and Reflective Experience

A distinction should be made between straightforward and reflective experience. 'Straightforward experience' is as lived. In asking participants about their experience of 'presence' they are distanced from it. No reflection can catch the moment. This is what Zahorik and Jenison (1998) refer to in their phenomenological reading of presence.

They write,

> 'In order to report mentally represented things, one must step back from the primary mode of existence. But stepping back removes one from exactly the existence mode of interest, the primary mode of existence, that of concernful action.

These researchers do not present the above as an argument against the possibility on reflecting on presence, in fact they appear unaware of the full import of their statement. However, in the same article they present Heidegger's characterisation of the hammer in embodied activity, and Merleau-Ponty (and others) have done the same for the experience of the body.

Such a task may be difficult for participants, however a skilled interviewer will ease the process. In addition to eliciting participants accounts of presence experience, it is also possible to take the approach of Heidegger and Merleau-Ponty, that is, to reflect on our own experience of presence.

Role of Researcher In Phenomenological Research Outcomes

A positivist stance identifies the qualitative researcher as a potential source of bias and contamination of data, as well as selective in both the collection and analysis of data. As Bogdan and Taylor (1975) note, such critics fail to recognise that the researcher acts as a sieve in all forms of research. A phenomenological approach is one in which the researcher is recognised as an intimate part of the research process, rather than an element to be factored out. As part of the naturalistic paradigm, phenomenology rejects the 'dualistic distinction between knower and known, leading to a realization that the personal is always present in research.' (Henwood and Pidgeon 1992: 100).

'Ability' of Participants and the Researcher To Convey Experience

While phenomenological inquiry seeks to illicit rich descriptive accounts of experience, such an aim can be rendered problematic by the (in)ability of both the participant and the researcher to accurately and vividly convey particular experienced phenomena.

Language in the phenomenological interview is both the method of recalling and interpreting experiences. The difficulty arises from both our ability to use language to convey our experiences (e.g. Dreher 1994), and with regards to the adequacy of language for this task in the first place. The problem arises in the fact that we always live more than we can ever express in language. However, with this recognition in mind, von Eckartsberg (1986) argues that we can express much more than we usually do if we put in the effort for a careful description.

> 'The interrelationship of language is a difficult psychological conundrum. How is it that we can say what we experience and yet always live more than we can say, so that we could always say more than we in fact do? How can we evaluate the adequacy or inadequacy of our expression in terms of it doing justice to the full lived quality of the experience described?' (von Eckartsberg, 1986: 17).

In terms of performing an analysis and interpretation of 'life-texts', the implications of the above is that we should take care that our interpretations are grounded in the differing experiences of our participants and not a reflection of the differing abilities of our participants to use language (Dreher 1994).

(In)Consistencies in Participants' Data

During an interview, or over a series on interviews, participants may appear inconsistent in their statements or responses to questions. It is important that these 'inconsistencies' be fully explored with the participants, in order to see whether a change of opinion has occurred, or whether we do not have a full grasp of the complexity of our participants experience.

> 'What appears to be false or inconsistent according to your perspective and your logic may not be according to your subjects'' (Bogdan and Taylor 1975: 11).

Applying the Phenomenological Approach to the Study of Presence

Having discussed the phenomenological approach in detail, along with a methodology (and its problems), we can turn towards the application of this approach to the study of presence in the real environment, as well as its subsequent use in virtual (and teleoperator) environments.

Before studying presence in virtual environments, a conceptualisation of the phenomena in the real world needs to be obtained. Interviewing participants about their experience of presence involves explicating their embodied experiences within the world. Rather than asking questions such as ‘to what extent do you feel present within the environment (this room, etc.)’, an interview should address such issues as ‘what is it like to be here?’, ‘how would you describe where you are?’, ‘how do you know you are here rather than somewhere else?’ In fact, questions such as these should be reflected upon by the researcher first of all, and their own reflective descriptions should be transcribed.

From a researcher’s own reflections, themes can be generated. These are added to, elaborated upon, and expanded by participants own discourse. From this data, a phenomenological characterisation of presence can be distilled. Rather than a set of determinants of presence, a phenomenal description of presence, and its qualities (rather than determinants) within the real world is obtained.

In order to investigate presence within virtual environments, similar procedures are adopted. First of all, a self-reflection can take place (as above). Questions can be asked such as ‘what is it like to be ‘in’ this environment?’, ‘What am I aware of at this moment’? Again, a set of themes can be gathered from this researcher’s own reflections. Next, descriptions from others can be sought. Instead of a standard interview, participants could be encouraged to narrate their activity. The spontaneous language that participants use can be employed to generate further themes, as well as deepen others. The researcher can ask questions about their discourse and activity, providing further responses which can be analysed qualitatively.

Next, a comparison of the themes generated in the real environment, with those generated in the virtual environment can be made. Those themes particular to the real environment, those particular to the virtual environment, and common themes can all be identified. In this way, a characterisation of virtual and physical embodied experience can be made. An evaluation, and characterisation of presence can take place with any system in this manner.

This suggested framework for investigating presence within real and virtual environments is currently being undertaken. Findings of this work will be presented in future escape work.

Conclusion

The aim of this appendix has been to present the various ways that the concept of ‘presence’ has been conceived. We have also overviewed how ‘presence’ has been theorised and investigated empirically in virtual environments.

This paper has adopted a phenomenological approach to widen the conception of presence, and to suggest other ways in which it can be conceived. In drawing attention to the perceived presence of other social actors, a previously neglected area has also been highlighted for future theoretical and empirical work.

Finally, a phenomenological method of inquiry has been outlined and advocated for the explication of the phenomenon of presence. Future work will focus on implementing this.

References

Biocca, F., & Delaney, B. (1995). Immersive virtual reality technology. In Frank Biocca & Mark R. Levy (Eds.), *Communication in the Age of Virtual Reality* Hillsdale, NJ: Lawrence Erlbaum Associates. (pp. 57-124).

Biocca, F. (1997) the cyborg's dilemma: Progressive embodiment in virtual environments, *Journal of Computer-Mediated Communication*, 3(2).

Bricken, M. (1992) Virtual worlds: No interface to design. In M. Benedikt (Ed.) *Cyberspace: First Steps*. MIT Press: London. (pp.363-382).

Davies, C and Harrison, J. "Osmose: Towards Broadening the Aesthetics of Virtual Reality", *ACM Computer Graphics: Virtual Reality* (Volume 30, Number 4)

Dennett, D.C. (1978) *Brainstorms: Philosophical Essays on Mind and Psychology*. Penguin Books: London.

Dufrenne, M. (1987) *In the Presence of the Sensuous: Essays in Aesthetics*. Humanities Press International: New Jersey. (Edited by M.S. Roberts and D. Gallagher.

Eriksson-Mangold, M.M. and Erlandsson, S.I. (1984) The psychological importance of nonverbal sounds: An experiment with induced hearing deficiency, *Scandinavian Audiology*, 243-249.

Gilkey, R.H. and Weisenberger, J.M. (1995) The sense of presence for the suddenly deafened adult, *Presence: Teleoperators and Virtual Environments* 4(4), 357-363.

Grosz, E. (1994) *Volatile Bodies: Toward a Corporeal Feminism*. Indiana University Press: Bloomington and Indianapolis.

Heeter, C. (1992) Being there: The subjective experience of presence,

Telepresence, *Presence: Teleoperators and Virtual Environments*, 1(2), 262-271.

Heidegger, M. (1962) *Being and Time*. Harper and Row: New York. [Translated by J. Macquarrie and E. Robinson].

Held, R. and Durlach, N. (1992) Telepresence. *Presence: Teleoperators and Virtual Environments*, 1(1), 109-112.

Hendrix, C. and Barfield, W. (1996a) Presence within virtual environments as a Function of visual display parameters, *Presence: Teleoperators and Virtual Environments*, 5(3), 274-289.

Hendrix, C. and Barfield, W. (1996b) The sense of presence within auditory

Environments *Presence: Teleoperators and Virtual Environments*, 5(3), 290-301.

Ihde, D. (1996) *Listening and Voice : A Phenomenology of Sound*.: Ohio University Press, Athens.

Ihde, D. (1990) *Technology and the Lifeworld: From Garden to Earth*. Indiana University: Bloomington and Indianapolis.

Kim, T. (1996). Effects of presence on memory and persuasion. Unpublished doctoral dissertation, University of North Carolina, Chapel Hill. Cited in Lombard and Ditton (1997).

Kujundzic, N. and Buschert, W. (1994) Instruments and the body: Sartre and Merleau-Ponty, *Research in Phenomenology*, 24, 206-215.

Kvale, S. (1984) The qualitative research interview: A phenomenological and a hermeneutical mode of understanding, *Journal of Phenomenological Psychology*, 14(2), 171-196.

Leder, D. (1990) *The Absent Body*. The University of Chicago Press: Chicago and London.

Lombard and Ditton (1997) At the heart of it all: The concept of telepresence, *Journal of Computer-Mediated Communication*, 3(2

Merleau-Ponty, M. (1970) *The Phenomenology of Perception*. Routledge and Kegan Paul: London. [Translated by C. Smith].

Miller, S.I. and Hodge, J. (1997) Phenomenology, hermeneutics and narrative analysis: Some unfinished methodological buisness

Murray, C.D. (1996a) Experiencing the body: Informing speculative discourse of polymorphous bodies in virtual reality. In the short paper proceedings of *the Third UK VR-SIG Conference*, De Montfort University, Leicester 3rd July. (pp.77-87)

Murray, C.D. (1996b) Technologies of embodiment: Wielding the scalpel in virtual reality. Paper presented at: *Ciber@rt: First International Conference on Virtual Reality*. Valencia, Spain. November 4th-8th.

Ramsdell, R. S. (1978) The Psychology of the Hard-of-Hearing and the Deafened Adult. In H. Davis and S.R. Silverman (Eds.) Holt, Rinehart and Winston: New York. (pp. 499-510).

Sartre, J-P. (1957) *Being and Nothingness : An Essay On Phenomenological Ontology*. Methuen: London.

Sartre, J-P. (1970) The body. In S.F. Spicker (Ed.) *The Philosophy of the Body*. Quadrangle: New York.

Schafer, R.M. (1992) The glazed soundscape, *The Soundscape Newsletter*, 4

Sheridan, T.B. (1992a) Musings on telepresence and virtual presence. *Presence: Teleoperators and Virtual Environments*, 1(1), pp.120-125.

Sheridan, T. (1992b) Defining our terms, *Presence: Teleoperators and Virtual Environments*, 1(2), 272-274.

Short, J.; Williams, E. and Christie, B. (1976) *The Social Psychology of Telecommunicatons*. John Wiley: London.

Slater, M., & Usoh, M. (1993). Representations systems, perceptual position, and presence in immersive virtual environments, *Presence: Teleoperators and Virtual Environments*, 2(3), 221-233.

Slater, M. and Usoh, M. (1994) Body centered interaction in immersive virtual environments. In N.M. Thalmann and D. Thalmann (Eds.) *Artificial Life and Virtual Reality*. John Wiley and Sons: Chichester. (pp.125-147).

Slater, M., Usoh, M., and Steed, A. (1994) Depth of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 3(2), pp.130-144.

Spinelli, E. (1989) *The Interpreted World*. Sage: London.

Straus, E.W. (1966) *Phenomenological Psychology*. Tavistock Publications: London. [Translated, in part, by E. Eng].

Varnado, S. L. (1987) *Haunted presence : The numinous in Gothic fiction*. University of Alabama Press: London.

Vale, R.S. and King, M. (1978) An introduction to existential-phenomenological thought in psychology. In R.S.Vale and M. King *Existential-Phenomenological alternatives for Psychology*. Oxford University Press: New York. (pp.3-17).

Van Manen, M. (1984) Practicing phenomenological writing, *Phenomenology and Pedagogy*, 2(1), 36-69.

Von Eckartsberg, R. (1986) *Life-World Experience: Existential-Phenomenological Approaches in Psychology*. Center for Advanced Research in Phenomenology and University Press of America: London.

Welsh, R.B, Blackman, T.T., Liu, A., Mellers, B.A., Stark, L.W. (1996) The Effects of pictorial realism, delay of visual feedback, and observer interactivity on the subjective sense of presence, *Presence: Teleoperatorsand Virtual Environments*, 5(3), 263-273.

Zahorik, P. and Jenison, R.L. (1998) Presence as being-in-the-world, *Presence: Teleoperators and Virtual Environments*, 7(1), 78-89.

Zeltzer, D. (1992) Autonomy, interaction, and presence, *Presence: Teleoperators and Virtual Environments*, 1(1), 127-132

Appendix Two: eSCAPE Art Works : the Commissioning Context

Sally Jane Norman

ZKM

Introduction

To develop new approaches to the design and integration of electronic landscapes, notably by prompting original reflection on the issues of presence and representation in virtual spaces, eSCAPE is building a number of prototype works with artists. This involves setting up collaborations between commissioned artists and computer hard- and software developers. The development process is being enriched by feedback from the social sciences, which have already acted as a driving force for many pioneering CSCW developments.

As an eSCAPE partner, and in its capacity as a centre devoted to art and new media, the ZKM Institute for Visual Arts launched a first call for artists' proposals in August 1997. Issued as a document (see appendix 1), the call summarised eSCAPE objectives, described the ZKM's role within the project, outlined the commission goal and requirements, and laid down specifications for the sought proposal, including the deadline and schedule. The draft document was sent to eSCAPE partners for feedback, duly amended, then the final call was sent to the pre-selected group of seven artists on August 11th.

What initially looked like a simple assignment turned out to be quite complex. Our first task as commissioner of works was to perform an extensive review of current artistic practices using leading edge technologies, so as to determine those most apt to lend themselves to development across a platform such as eSCAPE. To this end, we were not only focussing on technical features (i.e. exploitation of hard- and software tools pertinent for our project resources and goals), but were moreover concerned to identify artists likely to actively participate in dialogue with counterparts from a wide array of disciplines, and provide energetic feedback to other project members. This in itself turned out to be an exclusion criterion for certain potential candidates, for various reasons which we shall attempt to elucidate below.

Regarding formulation of the actual call to proposals, it was not easy to draw up this document with enough precision and, at the same time, enough margin to appeal to a number of very different artists; nor was it easy to target artists likely to be interested in and relevant for our predefined eSCAPE themes. We felt that it was important to stress "shareability" of the targeted experimentation, and likewise felt that it was necessary to reproduce the original eSCAPE document caveat on copyright, to preclude any misconceptions that might later jeopardise project development. Clearly stipulating these aspects, at the risk of sounding tedious, was considered preferable to omitting them and thereby leaving the path open for awkward misunderstandings. Finally, imposing deadlines and commission dates to ensure that the artistic contribution tied in sufficiently with the overall project and workpackage schedule was another potentially dissuasive factor : we were requesting high availability, by contacting artists in August for a commission start-up barely four months later. Responses were received from four of the seven addressees, and subsequent questions put to non-responders to determine the reasons for this, and to amend this and future calls for proposals accordingly, have gone unanswered.

Ultimately, then, our first eSCAPE commissioning exercise proved to be problematic in terms of both procedure and results. This is not entirely surprising : despite the ZKM's substantial experience inviting and hosting artists in residence, and its broad range of contacts within the international arts community, it was the Centre's first experience commissioning works in a context involving so many specific constraints (generally, artists are invited to ZKM after submitting an "open" project proposal - i.e. not unduly thematically or platform-bound - specifications essentially bearing on the time frame required to achieve the work, and equipment needs). Since the way artists can be associated with developing new electronic community models is a crucial aspect of the eSCAPE undertaking, we felt that it was equally crucial to analyse and document this experience, in order to elaborate appropriate strategies for the future. In addition to discussing the ins and outs of the commissioning process with the artists who actually responded to the call, we have raised this issue with many others, and with our technical counterparts. The account that follows is thus based on exchange which took place over an approximately six-month period. The positions adopted are shared by many, and can be considered generic; for this reason, and in answer to several explicit requests, these positions have been left anonymous.

Artist - developer teams and raisons d'être

Artists working with digital technologies in ways likely to be of interest to a project such as eSCAPE have necessarily acquired substantial experience in developing creative concepts with new tools : they are familiar with the team work this requires, and have frequently been involved in close collaborations with programmers and implementers. In the course of their experience, some artists have built up close and loyal relations with technical counterparts whom they consider and acknowledge as full-fledged participants in the creative process. For such persons, embarking on a project which requires close collaboration with a totally new set of partners may be difficult or even inconceivable. Apart from the fact that this implies recommencing the often laborious exercise of creating the common terms which subtend real dialogue, the relations of confidence steadily built up between an artist and programmer may extend to the tools they have forged in common. These tools often comprise an essential starting point for future creations, practically and conceptually.

The formidable arsenal of secrecy and confidentiality agreements that have to be signed by anybody who works for leading edge technology companies is just one rather eloquent indication of how industry tends to view pioneering research. Artists aware of the importance of their contribution to "hot" R&D fields may justifiably be reticent about becoming involved in consortiums where they are pushed to the hilt for their creative input, only to find that after giving impetus to teams generating new software and interfaces, they cannot even access, let alone keep and re-use, the tools wrought by this process. When the creative work injected into a collaborative project itself constitutes the fruit of years of effort, this kind of stalemate is clearly unacceptable. Consequently, a symbiotic relation with a trusted programmer, in the knowledge that tools developed will continue to stand as available resources and assets for future work, may be preferred to starting up collaborative work within an unknown, largely unforeseeable structure.

Existing tandems or teams may be felt to constitute an essential force in elaboration of a work, in that previous shared experience will have "tuned" technically expert partners to the types of demands artists are likely to place on them. A common language will have been established and honed - often after much trial and error experimentation - allowing the artist to freely communicate concepts, and the engineer to readily recognise and bring about their optimal translations in technical terms. Indeed, understanding in some cases may be fluent to the point where an engineer is able to anticipate on requirements, and spontaneously suggest tools most apt to convey lines of creative thinking that characterise an artist partner. This building up of a dialogue, and the mutual confidence it implies, by no means occurs automatically, and those who have battled at some length to create this situation may not be prepared to readily forego it. Artists who have a long history of working with the same programmers thereby share a rich heuristic background with the latter : the overall working process, which consists of seeking and finding solutions to conceptual and technical issues, constitutes an inestimably valuable common culture, and a precious resource for making new work.

Certain creative research units developed as offshoots of major industrial firms may trigger suspicion amongst artists they invite to momentarily join their R&D teams. However generous they might seem, residencies and grants issued via "brand-name" organisations (e.g. Canon, Rank Xerox) are sometimes viewed as an ill-disguised bid for cultural bonification on the part of companies simply seeking to redeem their public image through altruistic "cultural" activity, thereby distinguishing themselves from their competitors (1). Reining in artists to catalyse certain kinds of corporate research would logically seem to imply preferential, if not exclusive use of a given set of tools, but creators looking for an openly experimental situation may find this demand inappropriate and too restrictive. Moreover, even if they consider the proposed experimental conditions to be challenging, artists may be hesitant about undertaking research that, in the long run, is likely to engender a perverse form of "technological habituation", i.e. strong dependence on the tools and human competence to which they have enjoyed privileged - albeit fleeting - access. Consequently, they may shy away from whirlwind "flings", from short-lived, highly concentrated collaborations with brilliantly specialised technicians, and prefer the constancy of less dramatic, but more lasting associations with polyvalent technical partners.

On several occasions in the course of our discussions, the necessity to engage "artist + programmer", i.e. an existing team, rather than a solitary individual, arose instantly, chiefly for the above-mentioned reasons. Some artists insisted heavily on the major creative role ensured by their technical collaborator(s), making it quite clear that this function had become indissociable from their own - often to the extent that their works were systematically co-signed. Certain structures offering residencies focussed on art and technology undertakings, like Ars Electronica in Linz (Austria), seek to answer this particular need by issuing calls for proposals from preconstituted interdisciplinary teams, made up of persons with artistic, theoretical, and technical competence which is synergised across an original creative vision (2). Gerfried Stocker, director of the Ars Electronica Center, has a dual artist-engineer background which has made him keenly aware of the specific contexts that must be nurtured to generate new kinds of technology-grounded works. For the time being, though, this type of strategy, developed by a pioneering organisation well versed in the realm of new cultural practices, remains exceptional.

Another complex case is that of artists' groups, where several individuals have consolidated an integral creative identity across joint work. In their most dynamic form, such groups are comprised of complementary persons with different fields of competence and, at the same time, high affinities for certain kinds of conceptual exploration. When these teams are in turn bent on expressly working with certain technical partners, commissioning an art work from them can become a fairly weighty undertaking. Yet it is precisely this kind of structure that risks being a propitious platform for interdisciplinary encounter, since its multiple facets likewise constitute multiple potential hold points for setting up a dialogue.

Apart from loyalty to trusted technical partners, artists may also feel obligations and loyalties towards partners from other spheres of activity who assume an important public relations function; these persons can play a decisive enabling role in their creative research. Certain cross-disciplinary undertakings are extremely fastidious to set up, and require careful bridge-building between different people and fields; artists may be closely dependent on "ambassadors" to instigate initial contacts. These are likely to be highly committed institutions and/or individuals, prepared to energetically fight for creative work they believe in. Once more, such loyalties must be recognised and respected, the more so in that such "ambassadors", whatever the ostensible duration of their involvement with the artists, will often - and rightly - be viewed as having (had) a determinant influence on the scope and evolution of a creative itinerary. Building up a network of reliable allies - a human resources platform, as opposed to a platform comprised of sheer technical resources - is often a long-term undertaking. Artists who have managed to do this realise that knowing the right people in the right places, directly or through their "ambassadors", is worth as much as or more than any financial grant, since it is this type of network that ensures their access to tools and to competent technical counterparts.

Spotlighting R&D

While artworks grounded in sophisticated digital tools clearly necessitate an intimate blend of creativity and technical expertise, enabling a certain ease of crossover between the poetic and technological spheres, this does not necessarily imply that the works themselves are developed in a climate of “open” or “public” interdisciplinarity and debate. On the contrary, many artist-engineer tandems or groups operate in somewhat private worlds, which they consider to be the prerequisite for uninhibited experimentation. The idea of spotlighting fragile early phases of R&D, and building up vigorous open discussion on work in progress, is anathema to such teams. Their integration into larger development structures thus requires careful groundwork to build up the climate of confidence that nurtures open debate.

Apart from artists keen to develop technologically-anchored works with a maximum of freedom and privacy for “disinterested” creative reasons, others may be reticent to rally an eSCAPE-type consortium because they are conscious of the potential ramifications of their artistic concepts for future industrial and economic partners. Artistic “jealousy” is not necessarily kindled by purely poetic vindications, but may be fired by the artist’s own long-term strategies to transpose innovative concepts to the arena of mass production. It is important neither to underestimate the importance of such tactics, nor to brashly stigmatise them as mercantile, since they are indicative of major new trends in contemporary art practice, its insertion and its viability within a production-oriented society. The fact that growing numbers of industrial developers today scout electronic art events as a source of new ideas has awoken artists to the fact that there are potential market outlets for their work, and whetted their desire to valorise these outlets.

Indeed, for many, only hard-headed strategies of this sort can provide a guarantee of financial autonomy, thereby upholding their possibility to continue freely as creators. Consequently, involvement in projects announcing free, down-loadable shareware, may be just as suspiciously viewed as involvement in projects announcing lucrative spin-offs. Whether they are dealing with research-oriented institutions or with commercial enterprises, artists have as much reason as anyone else to think in terms of medium- and long-term survival. Indeed, they may have more reason than others : artists whose creative energies and integrity have led them to abandon certain mainstream cultural circuits (galleries, art fairs and markets), which are anachronistic but can still guarantee a livelihood, are often stuck with having to eke out a very precarious existence. This situation may contrast starkly and rather paradoxically with that of the institutional or industrial technology developers with whom these same artists feel and seek intellectual kinship.

Artists as (expendable?) beta testers

The fact that they may be confronted with flagrant discrepancies between their own often difficult professional conditions and those of their programmer-engineers, does not make it any easier for artists to elaborate coherent strategies as actors within the technological world. This is something that crops up regularly : programmers and engineers who have closely collaborated on artistic projects involving radical ideas and technologies are often considered by the industrial work market as a valuable resource. Such people tend to be ingenious, audacious solution-finders, capable of imagining and developing novel tools to instantiate sometimes very abstract, unwieldy concepts; this profile is in high demand amongst developers.

In contrast, although artists are increasingly sought after as beta testers, in the knowledge that their unpredictable conceptual approaches and demands are likely to challenge previously obvious and accepted technological limits, it is ironic that their role as catalysts in the development process tends, retrospectively, to be underestimated. Indeed, they are too often seen as expendable and replaceable. Yet productively challenging existing tools is not a matter of peremptorily "throwing down the gauntlet" and requesting a few technological frills (so-called artists who do this are indeed, fortunately, expendable). The only creators who can place really stimulating demands on tool design are those who, in their hermeneutic activity manifest as much rigour and intelligence as their technological counterparts. Such persons are few and far between, yet commissioning bodies purportedly seeking new, original creations often seem to idly assume that the painter-in-the-garret style romantic will be able to come up with a relevant contemporary work. Indeed, the fact that it is financially much easier to invest in such an artist than to back more onerous collaborations is one sadly mercenary reason for the astonishing durability of outdated visions of art. Another is the profusion of "Sunday painters" able to convince gullible sponsors that their works belong to the contemporary avant-garde.

Our insistence on the need to recognise and respect specificity of the artist's working environment and needs, particularly as manifest through past and ongoing collaborations, should not be read as a plea to respect the high-strung poet : a comparable situation arises whenever an individual researcher developing truly innovative material is taken on board a large group project. Software designers, for example, may express similar reluctance to be caught up in anonymous groups, and to share volatile ongoing research with a large number of people. The essential difference, though, resides in the fact that whatever their claims to solitary genius, software designers as a species date back a couple of decades only, whereas centuries of western social practice have tended to reinforce the image - and in many cases, the harsh social reality - of the isolated creator.

Informed versus Decorative Aesthetics

One area of flagrant misunderstanding that frequently arises between persons from the technical and artistic communities concerns the nature of the artist’s potential contribution to technological development. Many engineers developing innovative computer tools (notably software and interfaces) tend to see the artist as a source of "added value", i.e. as someone likely to provide a "cultural bonus" capable of differentiating their technology from that being developed by their competitors. Such persons consider the role of art as being essentially ornamental or decorative : dauntingly austere technologies can be made enticing and exciting if correctly window-dressed (similar issues arise when artists working alongside scientists are employed solely as illustrators, instructors, and popularisers of scientific theory, their creative talents being used to make obscure concepts accessible and palatable). Indeed, the widespread, successful integration of new tools, their user-friendliness and learnability, may be largely hinged on displays made attractive through original use of shape and colour, enhanced by the "intuitive" appeal and legibility of skilled layout work, i.e. on features that characterise any good design. This begs the question of whether and to what extent design and art are synonymous. Still closer to the bone is the underlying question : what is art ?

Definitions of art are as unverifiable as they are innumerable. They range from conceptions which espouse notions of craftsmanship and free individual expression, through to ideals of revelatory masterpieces, bearers of truth and singular visions. In many societies, art is not separated from the magma of cultural practices, but comprises an integral part of everyday activities and experience, and cannot be differentiated from them. In European society, artistic authorship and incumbent social status can only rarely be dated back any earlier than medieval times, despite the fact that archeology has revealed thousands of years of production of what could be and often are called art works (needless to say, by modern Westerners first and foremost). Moreover, even in the brief period corresponding to western art history, definitions and ideals of artistic practice vary considerably, in ways that sometimes convey strongly reactive switches in thinking. Hence, the English Romantic loners who haunted the early nineteenth century later ceded to Ruskin’s and Morris’s ideals of anonymous arts and crafts workers, striving for the betterment of society. During the same period, the French academic stars of the atelier painters’ system were supplanted by a group of anti-academic outdoors enthusiasts, less concerned with the limelight than with natural light, who collectively launched the derogatively dubbed Impressionist movement. These are just a couple of recent examples which testify to the lability of the artist’s social status and role, and to the variable importance ascribed to solitary genius. Given the transformations in cultural practice now being wrought by digital technologies, and notably the ways they are impacting notions of individual creativity and authorship, it could well be useful to bear such examples in mind.

Trying to formulate a simple working definition of art in this current context is particularly arduous : certain artists have themselves become technology developers, thereby assuming a determinant role in creating tools they feel are apt to optimally convey creative expression. They are indeed designers, but in a very fundamental, rather than ornamental way. People like Russian-American artist and software developer Lev Manovich (for whom "the aesthetics is the interface") have a strong influence - at both the pragmatic and conceptual levels - on the artists’ community at large (3). Even those wary of the artist-developer profile, those who view its technical component as undermining and weighing down the "freedom" of artistic expression, cannot deny or counter the emergence of this new creative channel, and indeed are often amongst the first to benefit from the tools thus forged.

Other artists consider that their timeworn hermeneutic role as poetic interpreters (as opposed to scientific elucidators) of life’s mysteries, necessitates awareness of and focus on the aspects of human existence which today are prone to the most radical transformations. It is these extremely unstable, and at the same time portentous areas of existence that sharpen creativity, that most pressingly solicit the power of freely interpretative energies. As a corollary, poetic insight alone allows us to embrace certain highly complex notions and events in a readily communicable, sensible fashion. Sheer decorative aesthetics are clearly far removed from this conception of the artist’s task, seen as the sensing, monitoring, and interpretation of epistemological breakthroughs. One of our eSCAPE computer science partners has summed up this opposition as the difference between "aesthetically decorative" and "aesthetically informed" work (4).

When we refer to attractive, engaging visualisations, it is important to define what we mean by these terms : displays which at first glance look relatively austere can convey powerful aesthetic intention and meaning through instrinsic structuring and ordering mechanisms. This may make them conceptually engaging and attractive in very rich, subtle ways. Conversely, visualisations loaded with seductive techniques learned from marketing specialists may not even bear up to a second glance, if these display tricks have been superficially tacked on like tinsel to a structure with which they share no conceptual ground. In the latter case, there cannot be any further aesthetic engagement once the initial decorative thrill has palled.

Unfortunately, however, it is precisely the more ground-breaking visions of the artist’s role that tend to meet with the most incomprehension, if not downright resistance, on the part of non-specialists - the general public and potential technological collaborators alike. Decorative design work is obviously easily dealt with at a number of levels, be they conceptual or economic : as mentioned earlier, commissioning a painting tends to be simpler - and cheaper - than commissioning a work which deeply exploits and questions new tools and concepts. Similarly, commissioning an artist to "dress up" technology seductively is less problematic than commissioning an artist likely to subvert it for creative purposes. That said, it would be misleading to oversimplify the design issue : the many understandable reasons for hiring a designer to make new products more engaging range from unabashed marketing strategies to the most altruistic pedagogical goals.

Conclusions

The various contradictions, conditions and constraints revealed through this commissioning process emphasise the need for a much deeper and wider grasp of what constitutes a viable working framework for artists involved with new technologies. In the eSCAPE context, the "artist/ designer" dichotomy raises questions as to where and how creative energies are and should be invested in multidisciplinary technological development work.

At the ZKM, as commissioners of art, we have tended to focus on creation which can be characterised as "aesthetically informed", rather than "aesthetically decorative". This is because we feel that the engagement arising from committed conceptual exchange between artists and developers is more likely to engender profoundly new approaches to e-scapes than engagement at a superficial design level. As employed here, the notion of "superficial" is not meant to be pejorative : our goal is neither to hierarchise nor to make simplistic value judgements on the designer’s and artist’s respective roles (indeed, in some cases these two roles may be fulfilled by one and the same person). Nevertheless, we believe that for artists to be effective partners in technological development, they should be involved as far upstream in this process as possible. Only under these circumstances can they be truly influential in helping to shape technological concepts and tools. Moreover, only under these circumstances are artists likely to truly benefit from exchange with their technological partners, as opposed to being called in on an ad hoc basis as gifted "hired hands", to dress up new products for public showcasing.

Artists have been romanticised, mythified, and ostracised as unusual individuals (for many, this is a euphemism) over a period of several hundred years, and precedents for artists’ associations with industrial and/or technological developers remain sparse, even if this is changing fast. The extent to which software designers and artists are in the process of becoming one and the same person is a moot point at this stage, although this point is bound to recur in a more focal and very interesting way during eSCAPE development. Current difficulties involving artists in technological research projects may well turn out to be transient, as the two communities gradually merge and new professional profiles emerge. In this rapidly evolving setting, eSCAPE with its three-year lifespan offers a particularly valuable vantage point for studying the elaboration and "in vivo" testing of new kinds of interdisciplinary encounter between artists and technology developers.

Notes

1. "Manichean oppositions of humanists and technologists are caricatures at best, but today's situation is all the more complex in that the digital revolution has engendered a new set of cultural premises and actors, with disconcertingly mixed values. The most altruistic trusts and foundations may launch on commercial production ventures to cover costs and thus uphold their autonomy. Conversely, the most mercantile businesses may allocate substantial resources to cultural institutions through a genuine interest in the arts, a desire to redeem their public image, a strategy to open up new lines of research and development, a scheme to alleviate company taxes, or for all of these reasons." S.J.Norman, "Culture and the New Media Technologies", paper for the Stockholm Unesco Cultural Conference [http://www.unesco-sweden.org/Conference/Papers].
2. See the "Ars Electronica" chapter of our study on "Transdisciplinarité de genèse de nouvelles formes artistiques"; report commissioned by the Délégation aux Arts Plastiques, Ministère de la Culture, France, November 1997 [http://www.culture.fr/culture/mrt/bibliotheque/norman/norman.rtf].
3. [http://jupiter.ucsd.edu/~manovich]
4. This distinction was made by Adrian West, of the Advanced Interfaces Group at the University of Manchester, during an eSCAPE plenary debate in February 1998.

Call for Proposals

DESCRIPTION OF eSCAPE OBJECTIVES

ZKM, together with two British partners (the Computing Department at Lancaster University, and the Psychology Department at Victoria University of Manchester), and one Swedish partner (the Swedish Institute of Computer Science), is launching a long-term European research programme called eSCAPE, focussed on the theme of Electronic landSCAPEs, or shared virtual environments.

Shared virtual environments, where VR technology is used to visualise a space shared by multiple individuals, are of increasing interest from a variety of perspectives. It is not uncommon for the proponents of shared virtual environments to claim that they may support social interaction in ways which go beyond what is possible using more familiar technologies such as videoconferences, 'mediaspaces' or shared desktops. Crucially, shared virtual environments permit users to become *embodied* within a shared space by means of an embodiment or ’avatar’. It is often argued that avatars permit a degree of self-expression and experimentation with self-identity for users, and many systems support the end-user configuration or design of embodiments. It has also been claimed that appropriately designed shared virtual environments enable users to sustain mutual awareness about each other’s activities.

As such environments proliferate and become more commonly known, a critical research challenge emerges. How should these different environments be interconnected so that users can truly access a virtual *world* and not merely a set of isolated local virtual spaces? In current systems, connections between virtual environments are typically by means of ’portals’ or ’gateways’ and travel between environments is a form of ’teleporting’. There is no integrating space, no environment which provides the interconnections between environments. This is precisely the function of an eSCAPE : a shared virtual place where other shared virtual places meet and where all those places are inhabited information spaces which can contain representations of persons (avatars), objects/information and artificial agents.

Description of ZKM’s ROLE WITHIN eSCAPE

Within the eSCAPE project, our British and Swedish partners will be providing expertise in the social and computer sciences, while ZKM’s role essentially consists of ensuring artistic expertise, namely through commissioning works from artists concerned with issues of presence and representation in electronic and/or hybrid (real/virtual) environments. Although these works are to be realised over a relatively short period (6 months) , they will be commissioned bearing in mind long-term eSCAPE objectives, which include promoting cooperative interaction in distributed multimedia systems via development of a number of novel software and platform demonstrators.

In keeping with these objectives, and leading on from the presence/ representation issues to be approached in the early stages of this research, we shall be working on two specific thematic spaces called the "Virtual Cityscape" and the "Virtual Planetarium". The former instantiation of the eSCAPE concept presupposes use of a fairly concrete, familiar metaphor as the means of providing interconnection between shared virtual environments. That said, existing on-line cities have been criticised for reproducing traditional institutions (e.g. virtual city halls and the like), thus underplaying the potential of inhabited virtual environments to develop *new* social and institutional forms. Accordingly, in eSCAPE we intend to explore more artistic and poetic interpretations of the metaphor of a Virtual Cityscape, where, for example, buildings and other features are not confined to a realistic reproduction of real forms but can be composed of textural, pictorial or abstract elements. The Planetarium metaphor is intended to allow for more abstract visualisation of shared virtual environments and their relations than the Cityscape metaphor may provide even in its most poetic forms.

COMMISSION GOALS AND REQUIREMENTS

In light of the overall structure of our research programme, we are concerned to commission works that will lend themselves to further development in keeping with these themes, and to involve artists keen on being involved at the seminal stages of long-term research of this kind. Indeed, these works will be the object of analysis by computer scientists and social scientists on the eSCAPE team, to determine their implications and potential ramifications at the technological and sociological levels. They may serve as test-beds for software and interface experimentation, giving rise to "eSCAPE demonstrators", i.e. works conveying novel approaches to shared virtual environments. Consequently, in addition to being able to capitalise on substantial resources and experience offered within the eSCAPE group, commissioned artists will be actively engaged in an ongoing dialogue with researchers - social and computer scientists - employed by our partner institutions. A certain amount of on-site availability at Karlsruhe thus constitutes an essential part of the commission. The selected artists will be required to provide regular feedback on their work, notably with eSCAPE collaborators who are designing tools which may be in line with the artists' technical requirements.

Ownership protocol

The eSCAPE project will focus on the production of material for public use and dissemination. To this end all eSCAPE deliverables will be considered public with electronic versions of the deliverables freely available. Rights to ownership of background information brought to the eSCAPE project by project partners will remain with the partner, and ownership of all intellectual property (including copyright, software and know-how) will reside with the originating party. The ownership of intellectual property for all jointly produced systems will be shared between those involved in its production.

FORMULATION OF THE INITIAL PROPOSAL

For the initial submission, we are seeking a brief description of the proposed work (text and summary visuals) , with an outline of the basic concept and its realisation (technical infrastructure). Indications concerning foreseeable implementations across different platforms would also be relevant (display interfaces, network parallel to localised installation, etc). Where pertinent, reference to previous works should be accompanied by documentation on these works (printed material, visuals and audiovisuals, web sites, etc.).

SUBMISSION DEADLINE, NOTIFICATION, & COMMISSION SCHEDULE

- Deadline for submission of initial proposals : October 1st 1997

- Notification of selected proposals : December 1st 1997
- Commission starting date : January 1st 1998
- Commission end date : July 1st 1998

For all further information, please contact ZKM, Institut für Bildmedien, specifying that queries concern the "eSCAPE" submission.