Virtual Reality International Conferences 2000
Laval 18-21 Mai 2000

Designing Virtual Environments for Social Interaction

Andy Crabtree & Mark Rouncefield
Centre for CSCW Research
Computing Department
SECAMS Building
Lancaster University
Lancaster LA1 4YR
United Kingdom
Tel. +44 1524 593097
Fax. +44 1524 593608
Email. a.crabtree@lancaster.ac.uk

Steve Pettifer
Advanced Interface Group
Department of Computer Science
University of Manchester
Oxford Road
Manchester M13 9PL
United Kingdom
Tel. +44 161 275 6141
Fax. +44 161 275 6204
Email. srp@cs.man.ac.uk

Keywords: shared virtual environments, requirements, social interaction, ethnography, situated evaluation.

Abstract

The Distributed Legible City (DLC) is an interactive multi-media art installation enabling a number of 'cyclists' to participate in a shared virtual environment. In this paper we describe the evolution of the DLC guided by ethnographic study. We consider the implications and problems associated with designing a virtual environment where the requirements to 'support social interaction' are ambiguous and open ended. Although our own work is but exploratory, we identify trans-situational features of social interaction which may be oriented to and explicated in other settings in designing virtual environments to support social interaction.

1. Introduction

Despite a great deal of hype surrounding the notion of 'virtual reality', virtual environments (VEs) are of burgeoning technical, social and commercial interest, promising substantial benefits to the accomplishment of distributed activities in overcoming real-world constraints of space and time. At the current point in time, however, constructing anything but the most trivial of virtual environments is a technically challenging task, particularly where computer-mediated human interaction is concerned.

Until recently much of the effort in the field has been directed towards demonstrating the basic feasibility of VEs as a mechanism for human-computer interaction. It is only of late that processors, graphics engines, networks, and software infrastructures have matured to the point where high-level issues such as computer-mediated interaction may be realistically examined. Though there are many important issues to be addressed in the development of VR technologies, there is growing consensus that the basic challenges are understood. What may be accomplished 'within' a VE, however, and what paradigms are appropriate to supporting the *achievement* of distributed activities, is a much more open-ended question?

In delivering what is essentially a report on our own experiences in the development of shared virtual environments, we do not claim to have found the 'silver bullet' to the 'wicked' problems encountered in the attempt to support distributed activities. Our work is largely exploratory in character and treats

general issues of computer-mediated interaction through the evolutionary development of what is basically an electronic artwork. As such, the particular 'installation' is of little relevance to understanding the many and varied organisational contexts within which VEs will ultimately be developed. We place an emphasis, however, on design methodology. That is, on treating the design of VEs supporting interaction as a *methodological* enterprise (in sharp contrast to a theoretical enterprise). It is in this respect that we feel that the lessons learnt in developing the artwork are of particular purchase (eSCAPE D4.0).

Of particular concern here are issues pertaining to the gathering of requirements supporting computer-mediated interaction across distributed sites. What kinds of interaction are we to design for? How are we to provide for the collaborative performance of activities? How are we to support the coordination of activities? How do we engender a sense of co-presence among participants immersed 'within' some virtual space? In short, how are we to provide for the accomplishment of intersubjectivity through the use virtual technologies?

Below we describe the design, implementation, evaluation and evolution of the Distributed Legible City (DLC), a shared virtual environment based on a multi-media art installation. We consider the technical challenges encountered during implementation, and examine issues associated with the uncertain character of requirements. We explicate the expectations and intentions of the designers in formulating the 'plan in the machine' (Suchman, 1987) and changes made to the installation in light of the results from the evaluation studies of the installation-in-use by members of the general public. In conclusion, we reflect upon methodological implications for the design of shared virtual environments supporting social interaction emerging from our current work.

2. Background to exploratory work

The DLC was constructed in the context of the ESPRIT long term research project eSCAPE, which aims to investigate the paradigms and technologies appropriate for the construction of large-scale virtual environments (or 'inhabited information spaces'). The project places an emphasis on public use and the research programme brings together engineers, social scientists and artists. The aim of the project is to inform the future construction of shared virtual environments that, whilst being visualisations of electronic information, are inhabitable and engaging places in their own right. One strand of research focuses on the development of 'facsimile' environments which closely resemble real-world spaces and places and which are, in many respects, based upon a 'cityscape' metaphor. On the premise that city-like structures are commonly understood and contain many features supporting interaction (such as maps, signposts, information kiosks, etc.), a number of techniques for generating 'virtual cityscapes' have been investigated within the project (eSCAPE D4.2).

The DLC was the first of the eSCAPE prototypes specifically aimed at exploring social interaction in a city-like environment. Whilst the low level technical needs (interactive frame rates, low device lag, etc.) are ostensibly the same as for any other VE, the 'purpose' of the environment and 'just what' activities it should support, is very unclear. Nonetheless, social interaction is, first and foremost, an intersubjective accomplishment in which persons *collaborate* to 'fit' their individual lines of action together in order to get their activities done (Blumer, 1966). The central problem we must address in the effort to support social interaction in VEs, then, becomes one of providing for the accomplishment of intersubjectivity. Or, alternatively, for the co-construction of mutually intelligible courses of action 'within' virtual time and space.

3. The DLC: an exploratory vehicle

The Distributed Legible City (DLC) was developed as a vehicle for exploring issues implicated in the co-construction of mutually intelligible course of action. The DLC is based on The Legible City, a 1990 multi-media art installation conceived by Jeffrey Shaw, Director of the Institute for Visual Media at ZKM in Karlsruhe, Germany (Shaw, 1998).

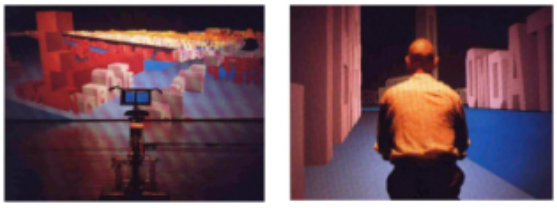

Figure 1. The Original Legible City

In its original form a single cyclist, seated upon a modified touring bike, is able to pedal around three virtual cityscapes. This original work is situated inside a darkened square room, with the bicycle facing a large back-projected screen. The screen displays the cyclist's view. The view is not of buildings and street furniture but instead, of solid coloured 3D letters forming sentences appropriate to the urban location. Thus the tourist is able to cycle around Manhattan, Karlsruhe or Amsterdam

taking in routes with commentaries about the respective city or fragments of texts associated with the area. Mounted on the handlebars of the bike is a liquid crystal display that shows a 'you are here' map and a large button that, when pressed, transports the cyclist between the three locations.

The first incarnation of the Distributed version of the Legible City was constructed with the MAVERIK (Cook *et al.*, 1998) and DEVA (Pettifer, 1999) VR systems, and consists of a number of 'stations' connected via networking technology. Each station was fitted with a 21-inch monitor mounted in front of a modified exercise bike. The cyclist was provided with an audio headset, with headphones and a boom microphone, via which to communicate with other users. A button mounted on the handlebar of the exercise cycle activated an overview map that was superimposed upon the current city's view. Dots on the map displayed the location of others in the environment and an arrow marked the rider's position and direction within the virtual space.

Figure 2. The DLC: the station, the overview map, and the cyclist's avatar

The user's 'body' was represented to the other inhabitants of the environment by an animated cycling avatar, the pedalling and steering motion of which was coupled with the actual pedalling and steering of the user's exercise bike.

4. The plan in the machine

Seen from a design perspective, the problem of providing for the accomplishment of intersubjectivity was in the first instance construed as one concerned with 'giving purpose' to the environment. In the effort to develop an appreciation of the manifold and subtle ways in which people accomplish intersubjectivity, we sought to avoid confronting users with overly complex interactive tasks by constructing an uncomplicated environment that provided enough of 'interest' to stimulate and support engagement and interaction. The use of the Legible City was motivated by informal observations of use at ZKM where users would happily explore the environment without any expectations of social interaction. These observations revealed that on a subtle level there is actually much 'to do' in the original Legible City. For example, a tour of much of Manhattan is possible following the musings of a taxi driver.

The intention in extending the Legible City to be a multi-user environment was that the subtle and somewhat esoteric content of the environment would be a 'talking point' promoting social interaction. In other words, the design plan had it that the novel structure of the Legible City – the text-form cityscape – would provide a basis for social interaction. Thus, the sequence of actions providing for the accomplishment of intersubjectivity consisted of: traversing the cityscape by cycling around; using the 2D overview map to navigate the space and locate others; using the headset to establish communicative relations with other users; and having made contact, orienting to the 3D text-form as a resource for interaction.

5. Encountering practical troubles

The DLC was exhibited at the 1998 Information Society Technologies conference (IST 98) in Vienna. Use of the DLC was subjected to ethnographic study (Crabtree *et al.*, to appear), the results of which are outlined here. Although meeting low-level requirements (high graphical update, low device lag, etc.), the studies revealed a number of 'usability' problems having direct impact upon the accomplishment of intersubjectivity.

5.1 Engaging with the environment

In the absence of an instructor, users of the exercise bike were unsure of just what to do. Although users often had an intuitive sense of aspects of the plan in machine engagement proceeded serendipitously and was often short-lived. A design decision meant that only cyclists close together could hear one another. The intention was that this would encourage users to find each another in order to hold a conversation. In practice, users would don the headset, pedal around a little saying "hello, hello .. is there anyone there?" and in the absence of a response, would cease interaction.

5.2 Navigating virtual space

The overview map was intended to be an intuitive navigational device supporting the 'finding' of other users 'within' the VE. Achieving competent use of the map was, however, a practically troubled affair. The action of 'pulling the map up' by pressing the cycle's handlebar button was not at all intuitive. Making sense of the map's features required instruction - that some features of the map represented "cities", that "this dot is you", and that "you are heading in that direction", for example. In the absence of instruction users would frequently abandon attempts at engagement.

5.3 Achieving conversational orientation

Users who did manage to locate one another in the virtual environment found it important to orient to, and achieve, a 'face-to-face' (avatar-to-avatar) position before engaging in conversation. This proved an extremely difficult and onerous task. As a consequence of the effort involved in achieving face-to-face positioning the conversation that ensued was often little more than perfunctory (e.g. "who are you?", "where are you?", "nice to meet you", "I'm going now .. bye bye") with some riders becoming visibly frustrated with the machine and ceasing engagement abruptly.

6. Rethinking the plan in the machine

The practical troubles encountered by users in the course of engagement highlighted a number of surprising and unanticipated discrepancies between the design plan and the situated actions of users.

6.1 Re-orientation of purpose

The first, and primary, result from the studies was a radical re-orientation of the perceived 'purpose' of the environment. As noted above (Section 4.), the intention during design was to provide an environment that enabled social interaction within an interesting context (the text-form cityscape). However, accomplishing interaction was, quite witnessably, the abiding practical concern of users. From the user's point of view the esoteric environment transpired to be anything but 'interesting' let alone taken as a resource for interaction. Although the cityscape was composed of novel textual forms which could be cycled straight through, it was treated instead in the ways in which real-world urban arrangements of space are *naturally* treated by cycling 'down' the DLC's 'streets' and 'around' its 'buildings' (Crabtree, 2000). Users cycled around the city streets, trying to avoid colliding with, what for them were ostensibly, text-form buildings. The plan in the machine was, as it were, in some disarray.

6.2 Use of the overview map

Compounding matters, although integral to engagement and the occasional accomplishment of intersubjectivity, the intelligibility and use of the overview map left a great deal to be desired. The map was designed as a familiar artefact which would intuitively furnish users with information about the structure and layout of the environment; position users 'within' the environment, and in being employed as a way-finding device, enable users to locate one another. It did not do so. The map occasioned practical troubles that were resolved either through instruction or serendipitously through experimentation and interpretation.

6.3. Achieving conversational orientation

The original design plan did not take into account the fact that users would orient their avatars to face one another in order to initiate communication. Though the cycle as an input device had initial benefits in terms of being easily comprehended, problems arose where the installation did not mimic a real bike. It was not possible to pick the exercise bike up and to turn round on the spot; it was not possible to "look over one's shoulder". To achieve a conversational orientation users would cycle round in ever decreasing circles until the correct orientation was achieved occasioning considerable frustration and frequently abandonment. The issue here is subtle, yet important to appreciate, drawing our attention to embodied practices for coordination and the accomplishment of intersubjectivity that may need support in VEs.

6.4 The new design plan

Attention to the situated actions in and through which use was engendered and intersubjectivity 'within' the virtual space occasionally accomplished, provided for the re-formulation of the design plan. Although the textual structure of the VE was maintained, an 'interest' in the cityscape itself was no longer presumed on the part of users. Instead, an abiding concern with the accomplishment of social interaction was assumed on the part of users and the new design plan sought to support that assumption in addressing the practical troubles users encountered.

In the light of the observable difficulties experienced by users the map was dropped in favour of a 'tour guide'. This took the form of an

animated, flying bird that positioned itself so as to appear to be flying in the direction of the closest other user whilst maintaining 'line-of-sight' contact with its 'owner'. Although it is clearly not natural to 'follow a flying bird around' in undertaking navigation, the particular affordances of the environment in terms of hardware and software suggested that this might be a feasible approach and solution. The idea was further motivated by the observation that users would keep the 2D map 'pulled up' for extended periods, thus ignoring the 3D environment. Implementing the tour guide offered the opportunity of encouraging thorough immersion in 3D space.

To support coordination in the achievement of conversational orientation, the decision was made to abandon a fixed point of view and implement a Head Mounted Display (HMD) on each station. This move occasioned significant changes in the architecture and provided for the exploration of fully immersive VE supporting the accomplishment of intersubjectivity. Each station was fitted with a HMD and position sensor, enabling the system to determine the position and orientation of the cyclist's head and to render an appropriate viewpoint. This gave the user the freedom to look over their shoulder or to one side without the need to re-orient the cycle. The animation of the virtual cyclist was modified such that the head orientation reflected that of the user, informing other users that 'eye contact' had been made.

7. The Immersive DLC

The second iteration of the DLC was exhibited at the Launch of the European Union's Fifth Framework Programme, in February, 1999. Users were 'introduced' to the installation, prompted to 'experiment' with the installation and instructed to 'follow the birds'.

Figure 3. The Immersive version of the DLC

7.1 The Immersive DLC in public use

In terms of supporting navigation and locating other users in the environment the tour guide proved a significant improvement, requiring less interaction between the demonstrator and users in the achievement of competent use. While users encountered the birds as unfamiliar objects, they clearly found it easier to learn and achieve competent use of them in situ as a result of the simplicity of the tour guide notion and an economy of instruction. In practice, once instructed on the significance of the birds, users of the Immersive DLC experienced little difficulty in using them to find others 'within' the VE, requiring at best a simple instruction to "look up for the birds" in the course of navigating and locating the other.

With the overview map removed from use, users were much more attuned to the 3D environment and undertook exploration of the space with a vigour not previously observed. Although exploration of the space was itself promoted by the HMD to some degree, it was for intended purposes – i.e. the coordination of action in the achievement of conversational orientation - that the Head Mounted Display came into its own. The natural and reflexive orientation to the speaker was provided through the VR headsets, thus supporting user interaction 'within' the virtual space. Just as the problem of 'overshooting' another user's avatar could be remedied by the turning of the head, so too the co-ordination of orchestrated movements was easily remedied by users. The VR headsets supported the natural inclination to orient to the conversational partner and consequently interaction increased.

8. Plans and situated actions

Whilst acknowledging the exploratory character of our work in supporting distributed activities through virtual technologies we feel we have identified some important and generalisable design lessons. In particular our research contributes to an appreciation of design issues in the face of what can only be described as the uncertain circumstances of VR development in general, and shared VE development in particular.

Insofar as the given end is the performance of distributed activities through the computer-mediated accomplishment of intersubjective,

mutually intelligible courses of action 'within' a virtual space, then the *formulation of the design plan* becomes a matter of specifying a sequential order of procedures for engagement. Thus the sequence of procedures providing for the accomplishment of intersubjectivity 'within' the DLC consisted of traversing the cityscape by cycling; using the 2D map to navigate and locate others; using the headset to communicate with others; and employing the 3D text-form as a resource for interaction. This approach towards design assumes that users will be able to recognise the plan in the machine and act accordingly. Yet, as the first study of the DLC in served to demonstrate, users often failed to recognise the underlying plan. Why was the plan unrecognisable?

One common theme in the literature on the character of plans is that plans are generative mechanisms of action. However, this is an idealised notion of the part plans play in the accomplishment of situated action (Suchman, 1987). Although plans are constituent of practical activities, it is not as generative mechanisms but as resources or artefacts for reasoning about situated actions. Plans are realised through embodied skills and practices. As a great many of the embodied skills and practices implicated in realising the plan for the DLC were not anticipated in formulating the plan, support for them was not implemented and there was then, from a users' perspective, something observably missing 'within' the VE. If we are to develop shared VEs supporting distributed activities, there would appear to be some necessity to develop an appreciation of embodied skills and practices implicated in the accomplishment of programme executable plans.

8.2 Some trans-situational features

'Just what' embodied skills and practices will require support in the development of particular shared VEs is a matter for occasioned inquiry. Although purely exploratory in character, user interaction with the DLC nonetheless serves to sensitise us to some significant trans-situational features of situated action that may be fruitfully oriented to and 'fleshed out' elsewhere.

Natural attitude and interactional competence
Understanding actions as the uncomplicated product of instructional sequences fails to recognise that people construct their own sense of situations in the course of accomplishing them. In a similar vein, interests in space and place that simply construe space and place as a container 'within' which persons construct their courses of action is similarly inattentive to situated action. Being irremediably situated in space, people adopt a 'natural attitude' that furnishes grounds for interaction and which is integral to the accomplishment of intersubjectivity. For example, in approaching the DLC, users do not see the installation as an 'artwork' but as a 'cityscape'. This enables users to establish a distinct sense of to how to 'go about' interacting 'within' the space. Spaces and their arrangements do not simply contain action but are thoroughly implicated in the accomplishment of action. Designs for VEs resembling real-world spaces may be fruitfully informed by attending to the interactional competences in and through which space and place are implicated in the accomplishment of everyday activities.

Schemes of interpretation
Activities take place against a 'seen but unnoticed' background of expectancies. People employ expectancies in everyday life as schemes for making sense of the settings, scenes and events they encounter. Thus, and for example, when encountering the overview map when undertaking engagement with the DLC, it soon becomes apparent that persons expect the map to explicate itself. The features by which it should do this (such as an adequate index) are not there however. The map fails to comply with people's common-sense expectations of maps, and in the absence of instruction, users often become frustrated. The significance of this for designing VEs is an open issue, for nothing says that VEs must correspond to real-world environments, only that they support intersubjectivity. Nonetheless, the accomplishment of intersubjectivity clearly relies on publicly available schemes of interpretation. Much may therefore be learnt in attending to the schemes of interpretation at work in particular settings through explicating the embodied ways in which persons make sense of artefacts and activities in the course of accomplishing intersubjectivity.

Instructed actions
One means of sensitising users to the 'plan in the machine' is to furnish users with a sequential course of instruction. Again, however, this is an idealised notion of how plans are made to work that produces a range of practical troubles in using novel technologies. There is an 'irredeemable incompleteness' to text-based instructions consisting of an absence of a description of the embodied actions whereby the instructions are to be realised. These details are missing. Yet, somehow, instructions must be turned into action. How? What does 'following an instruction' turn upon?

Suchman (1987) suggests that there is a collaborative aspect to instructional activities and that making instructions work is a co-production of participants. The studies of the DLC in use show that use consists of a course of practical instruction.

That course of work consists of an unfolding, sequentially ordered 'walkthrough' furnishing 'just enough' practical instruction for the user to accomplish the ordered steps of the sequence. It consists not simply in verbal (or textual) instruction but also and simultaneously, in instructed actions that give the situated talk (or text) its concrete 'just this' sense. When users are being instructed to 'follow the birds' for example, they are told to 'look up . this one is you . the other one is there'. Instructed action is underpinned by a practical method of demonstration-by-showing-and-doing. The method makes it observable 'just what' embodied actions are necessary. There is a generality to this method and insofar as users are interacting directly with the machine then the method may be incorporated into the plan in machine in devising support for the methodical communication of instructed actions.

Coordinating individual lines of action

We have seen how the coordination of 'face-to-face' (avatar-to-avatar) meetings in virtual space relies on ordinary conversational *methods* of orienting to the speaker, and of making inter-personal exchanges. In attending to the methodical ways in which people coordinate their individual lines of action through VEs we are primarily concerned with the ways people accomplish and concert their activities. Our interest is in 'just how', in observable details, people come to coordinate their individual and joint lines of action and thus come to get their activities done. The DLC demonstrates that failure to support the methodical ways in which activities come to be naturally coordinated can only result in rejection of the VE by users. Design must attend, then, to the inter-actional methods whereby distributed activities are 'put together' in the actions and interactions of parties to them (Bowers *et al.*, 1996).

8.3 Planning the work of development

In the absence of any established paradigm for VE design, we have chosen to adopt an evolutionary model in which the work of design is accomplished through a process of exploratory prototyping. The construction of early, only partially complete prototypes supports the elaboration and exploration of potential requirements. Their utility lies in making potential systems available to assessment by end-users, for changing ineffective designs, and for elaborating and refining good ideas. The strategy underpinning and guiding our prototyping activities might be characterised as an ethnographically-informed approach towards cooperative design (Crabtree, 1998). The strategy is ethnographically-informed in two distinct respects. Firstly, design is predicated on ethnographic studies of specific target domains (eSCAPE D4.1); secondly, as reported in this paper, the strategy is ethnographically-informed through the situated evaluation of prototyping sessions.

Situated evaluation of prototyping sessions draws attention to the cooperative character of our approach to design (Greenbaum & Kyng, 1991; Grønbæk *et al.*, 1997). Recognising that end-users are the real experts in the production of activities and accomplishment of intersubjectivity, the approach places users at the centre of the design process. Prototyping with end-users facilitates analysis of the utility of the potential system, and identifies breakdowns, glitches, and problems, by confronting the prototypes with practical circumstances of use. Situated evaluation of prototyping sessions facilitates the assessment and subsequent elaboration of the means whereby intersubjectivity is achieved.

Adopting the ethnographically-informed approach towards cooperative design as means of accomplishing the work of the eSCAPE project has seen the emergence of a 'reactive-proactive future strategy'. This strategy affords the exploration of more radical interaction technologies by grounding them within practical circumstances of intersubjective use. Consequently, in considering the future, two distinct but interrelated possibilities are evident. One the one hand, a reactive future strategy takes as its starting point in the issues raised by ethnographic studies of cooperative work situations and the situated evaluation of virtual environments. This strategy allows us to respond directly to issues implicated in the accomplishment of intersubjectivity and elaborate the emerging application. On the other hand, a proactive future strategy allows us to exploit the familiarity gained from the on-going development of thematic landscapes to expose users to more radical interaction approaches and techniques. This enables users to exploit the familiar nature of existing spaces to develop a working understanding of the potential utility of future spaces. Thus, in an iterative, evolutionary fashion, a reactive-proactive strategy affords the possibility of 'working up' the plan in the machine through the exploratory elaboration of situated action in the real and virtual world through the active participation of end-users and in-depth knowledge of practical circumstances of intersubjective use. Such a strategy allows for the potential emergence of a paradigm deeply attuned to the real-world character of social interaction as an intersubjective accomplishment.

Acknowledgements

This research was funded by the European Network for Intelligent Interfaces (i^{3}net), Esprit Long Term Research Project 25377 (eSCAPE). The authors

wish to acknowledge the invaluable contributions of all the partners and project members.

References

Blumer, H. (1966) Sociological Implications of the Thought of George Herbert Mead, *American Journal of Sociology*, vol. 71, p.535-548.

Bowers, J., O'Brien, J., Pycock, J. (1996) Practically Accomplishing Immersion: Cooperation in and for Virtual Environments, *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*, p.380-389, Cambridge, Massachusetts: ACM Press.

Cook. J., Hubbold, R., Keates, M. (1998) Virtual reality for large-scale industrial applications, *Future Generation Computer Systems*, p.157-166, New York: Elsevier Science.

Crabtree, A. (1998) Ethnography in Participatory Design, *Proceedings of the 1998 Participatory Design Conference*, p.95-103, Seattle: Computer Professionals for Social Responsibility.

Crabtree, A. (2000) Remarks on the social organisation of space and place, *Journal of Mundane Behaviour*, 1 (1), p.25-44.

Crabtree, A., O'Brien, J., Nichols, D.M., Rouncefield, M., Twidale, M. (to appear) Ethnomethodologically Informed Ethnography and Information Systems Design, *Journal of the American Society for Information Science*, New York: John Wiley & Sons.

eSCAPE Deliverable 4.0 (1999) *Towards a Common Methodology* (eds. Crabtree, A., Hughes, J.A., Rodden, T.), Esprit Long Term Research Project 25377, Lancaster University: Computing Department.

eSCAPE Deliverable 4.1 (1999) *The Library Abstract eSCAPE Demonstrator* (eds. Mariani, J. & Rodden, T.), Esprit Long Term Research Project 25377, Lancaster University: Computing Department.

eSCAPE Deliverable 4.2 (1999) The Tourist Physical Landscape Demonstrator (eds. West, A., Pettifer, S., Hughes, J.A., Rodden, T.), Esprit Long Term Research Project 25377, Lancaster University: Computing Department.

Greenbaum, J. & Kyng, M. (eds.) (1991) *Design at Work: Cooperative Design of Computer Systems*, Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Grønbæk, K., Kyng, M., Mogensen, P. (1997) Towards a Cooperative Experimental System Development Approach, *Computers and Design in Context* (eds. Kyng, M. & Mathiassen, L.), 201-238, Cambridge, Massachusetts: MIT Press.

Pettifer, S.R. (1999) *An Operating Environment for Large-Scale Virtual Reality*, Ph.D. Thesis, University of Manchester: Department of Computer Science.

Shaw, J. (1998) The Legible City, Presence and Representation in Multimedia Art and Electronic Landscapes, eSCAPE Deliverable 1.1, *Presence and Representation in Multi-Media Art and Electronic Landscapes*, Esprit Long Term Research Project 25377, p.28-35, Lancaster University: Computing Department.

Suchman, L. (1987) *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge: Cambridge University Press.