

Staging a Public Poetry Performance in a Collaborative Virtual Environment

Steve Benford, Chris Greenhalgh, Dave Snowdon and Adrian Bullock
Department of Computer Science, The University of Nottingham, UK

We discuss the design of a CVE poetry performance and experiences arising from staging it to two hundred members of the public. The design, a collaborative effort between computer scientists, artists, poets and producers, addresses issues of virtual world structure; embodiment of performers and audience; navigation interfaces; temporal structure of the event; and mixed reality presentation. Experiences include virtual audience members often ignoring the poets and conflicting attitudes towards embodiment. New CVE design possibilities are proposed, including object centred interaction, context sensitive interaction and tools to define and manage the spatial and temporal structure of events.

1. Introduction

In this paper, we present the design of a poetry performance in a collaborative virtual environment (CVE) and discuss the experience of staging this performance to an audience of two hundred members of the general public. The performance took place in Nottingham on 11th November 1996 as part of the NOWninety6 arts festival and involved an extensive collaboration between computer scientists, poets, a graphic artist and an event production team. This paper explores a range of CSCW issues arising from this experience, including:

- developing rich and engaging forms of content for CVEs;
- exploring the tension between autonomy of action for audience members and the management of a performance so as to encourage appropriate behaviours at relevant times and places;
- considering the role of 'virtual architecture' in enabling different modes of awareness and communication between participants as an event unfolds;
- developing appropriate embodiments for performers and audience;

- developing appropriate navigation interfaces for different participants (e.g. for members of the public with no previous experience of VR as compared to the performers and 'camera-people');
- exploring different ways of situating interfaces to the virtual environment within different physical spaces based on notions of mixed reality as proposed in (Benford, 1996);
- understanding the issues involved in designing and managing such an event and how these might influence world construction tools for CVEs.

The exploration of these issues is particularly relevant at the present time for several reasons. First, we believe that so-called 'citizen' applications of CVEs, such as performance, art and entertainment, represent a significant future application area for this technology. Recent months have seen a rapid growth of interest in publicly accessible virtual worlds on the Internet as discussed at the closing panel at CSCW'96 (Damer, 1996). Furthermore, in addition to the considerable interest of artists in VR technology as a medium for generating interactive artistic works, several people have begun experimenting with VR as a medium for performance and social interaction. Recent examples include the Membrane series of performances staged by Fahlén and Bowers (Fahlén, 1995) and 'The Mirror', a UK based public trial of a graphical and textual CVE running over the Internet and alongside the BBC Television series *The Net*. As the underlying technology matures, so the issue of providing appropriate content will become increasingly important. Second, this work involves a combination of artistic and technical perspectives on the design of a CSCW application. Just as the multi-disciplinary combination of social science and computer science has been at the heart of CSCW, so a synergy between art and technology might stimulate new innovations in systems design, especially in terms of user interfaces and new forms of content targeted at citizen users. Third, we believe that it is important to take CSCW technologies out of the laboratory and to demonstrate them and experiment with them publicly.

Section 2 offers a brief summary of the event. Section 3 then discusses key design issues in detail. Following this, section 4 reflects upon the experience of staging the event and key lessons learned. Finally, section 5 proposes a number of new design possibilities for general CVEs based on these lessons.

2. A brief overview of the event

Our poetry performance involved four hip-hop poets, Remi Abbas, Roger Robinson, Akure Wall and Dave "Stickman" Higgins, performing in turn in a collaborative virtual environment designed by the artist Sean Varney and supported by our own MASSIVE-2 CVE software. In addition to the performers, the environment was populated by ten virtual audience members at a time, where each audience member was an autonomous, mobile and embodied participant who could explore the environment and could interact with the other participants over live audio links

(active participation and social interaction by audience members was one of our principal goals). The virtual environment was structured as a central stage area in which the virtual poets appeared, surrounded by four outlying worlds containing different graphical designs and embedded text from the poems. The audience members were encouraged to explore these outer worlds during breaks in the performance, thereby adding an exhibition aspect to the event. The virtual poet embodiments featured a moving head and hands driven by Polhemus motion trackers attached to the head and hands of the physical poet. The audience members were embodied as "angels" and used a conventional workstation.

The performance was located in two quite different physical spaces. The ten workstations used by the audience members were located in a crowded café-bar. In contrast, the poets performed in an adjacent cinema which was structured as a more traditional performance space with a stage and seated audience. The use of projected graphics and audio links provided a limited form of mixed reality between the two physical spaces and the virtual one. In particular, the seated audience members in the cinema were given a view into the virtual environment so that they could see the virtual poet embodiment moving alongside the real poet and could also see and sometimes hear the virtual audience members.

The event lasted for four hours and was attended by a total of 200 paying members of the public. Given a cyclic performance schedule, 60 of these managed to experience some poetry as an audience member in the virtual world. The others were able to wander between the two physical spaces, watching a mixture of live performance and the virtual audience members using the system.

The software used was our own MASSIVE-2 system (Benford, 1997), a successor to MASSIVE as reported at ECSCW'95 (Greenhalgh, 1995). Like its predecessor, MASSIVE-2 supports communication in shared virtual worlds using graphical, audio and textual media. In addition, it supports an extended version of the spatial model of interaction which allows for greater scalability and also the implementation of different structural effects on awareness such as nested bounded regions and dynamic crowds. This is mapped onto an underlying network architecture based on a dynamic hierarchy of multicast groups. This software was run over a purpose built Ethernet connecting a mixture of Silicon Graphics computers (Onyx, Impact, O2 and Indy) connected to various video projectors and Polhemus trackers. The default machine for audience members was an O2.

3. The design of the event in more detail

In this section we focus on several key aspects of the design in the event in greater detail. These include: the structure of the virtual world; the embodiment of the different participants; the temporal structure of the event; the presentation of the virtual world within the physical world; and navigation interfaces. However, before progressing further, we should briefly mention the design process itself. The con-

cept of the event was first formulated by a combination of Computer Scientists and the NOWninetysix festival producers in March 1996. The poets and VR artist were then invited into the team and the concept evolved considerably over the following nine months. The VR artist was given the task of designing the virtual world and embodiment graphics, subject to some fairly severe constraints on complexity. These designs were provided in VRML format and a series of meetings over a period of several months were needed to refine them and integrate them into the MASSIVE-2 platform. A few of these meetings involved the poets who consequently had a limited opportunity to participate in the world and embodiment design (the major constraint being travel). Finally, there were several in-laboratory rehearsals followed by a full two day period just before the performance where the whole system was built from scratch in an isolated environment for final rehearsals.

3.1. The structure of the virtual world

The virtual world was designed to house a combination of performance and exhibition activities and consisted of a central stage area surrounded by four outlying worlds (one per poet). The screen-shot in figure 1 gives an overview of this structure, with the four outer worlds shown in collapsed form as seen from outside, i.e. as simple coloured cones (see below). The content of these outer worlds was designed by the VR artist so as to offer a visually engaging experience to an observer and to naturally lead them to some textual fragments of poetry located at their centres. Figure 2 shows the contents of each of the outer worlds in detail. The stage area consisted of a central raised platform on which a virtual poet embodiment would appear, surrounded by ten marked locations which acted as virtual seats for the audience members and which defined their individual home positions (see the discussion of navigation below and figure 4).

This structure was also designed to have a number of effects on the mutual awareness of its inhabitants. In particular, it used MASSIVE-2's notion of third party objects (an extension to the spatial model of interaction described in (Benford, 1997) to define the outer worlds as separate bounded spatial regions. Each of these outer worlds was defined as a region with the following properties:

- When viewed from the outside, its internal structure would be hidden away and replaced with a simple external view (a coloured cone). This extended to the region's inhabitants who would be invisible to those on the outside.
- On entering a region, it would be unfolded to reveal its contents (graphical structure and other inhabitants) in detail. In addition, those people located inside the region would still be able to see the structure of the outside world (i.e. stage area) and its inhabitants in normal detail.
- The audio medium was treated in the same way as the graphical one; all conversation occurring inside the region would not be heard by those outside, but any audio from the central space (e.g. the poetry) would still be heard by those on the inside.




Figure 1: Overview of the virtual world as seen from the stage area




Figure 2a: Stickman's world




Figure 2b: Roger's world




Figure 2c: Akure's world




Figure 2d: Remi's world

In other words, the boundaries of the four outer regions were defined to have an asymmetric effect on awareness. From the outside, they were opaque (both visually and aurally), but from the inside they were transparent. There were two motivations for this design. First, the regions were designed to offer private zones for social interaction, such that conversations within the regions would not interfere with the poetry in the stage area or the conversations in other regions and yet the poetry (and interactions of those near the stage) would still be visible and audible within the regions. Second, the regional structure aided network and computational scalability; the detailed structure of each region would only be rendered by its occupants and there would be no transmission of audio or movement data from those inside a region to those outside. In this way, the artist could be given freedom to design more complex graphical structures than would otherwise have been possible.

3.2. Embodiment of the different participants

The two different classes of participant in the virtual world, poets and audience members, were embodied in different ways.

The poets' embodiments needed to be visually interesting. In particular, we felt that they should be capable of gesturing. Thus, they were provided with relatively expressive embodiments, uniquely designed for the individual poets (see figure 3). Each poet embodiment consisted of a graphical object for a head, a separate body and two separate hands (in one case, one of the "hands" was actually a staff). The body section showed the poet's position within the virtual world. The movements of the graphical head and hands were driven by the movements of the physical poet through the use of Polhemus magnetic position trackers which were attached to some light-weight elastic loops that could easily be put on by the different poets. Consequently, the physical poet's head orientation and hand movements were mimicked by its virtual embodiment, allowing a range of gestures to be performed (although it took several rehearsals for the poets to adapt themselves to the exaggerated gestures that were required to produce something visually interesting on their graphical embodiment). A simple graphical mouth was also provided which would change shape whenever the poet was speaking.

Following early rehearsals, it was decided to locate two virtual poet embodiments back to back in the virtual world. This allowed for a circular stage design with each audience member getting a front-on view of the poet. It also supported the development of a 'camera control' interface for the projected display in the cinema (see below) which could keep both a frontal poet view and several audience members in shot from any position around the stage.

The audience members were provided with somewhat simpler embodiments (given that ten of them would be in the world at a time), being represented as graphical angels. Each angel was given a different colour and number, a name label above its head and a graphical mouth that would appear whenever its occupant was speaking. Unlike the poets' embodiments, the angels weren't articulated, and only served




Figure 3: The Poet's embodiments




Figure 4: Angel embodiments within the virtual world

to indicate 3-D position and orientation within the virtual world. Figure 4 shows several angels within the virtual world alongside one of the poets.

The design of the embodiments included the definition of foci and nimbi for the management of awareness and particularly for the spatial control of audio communication. All participants (poets and angels) were given a constant valued nimbus across the whole world, effectively providing an unconstrained opportunity for overhearing them. The audio focus of each participant was a multi-valued field with a peak of maximum focus within a range of five (virtual) metres, gradually dropping off to a zero at a range of three hundred metres (from the centre of the stage to the centre of one of the outer worlds was two hundred metres). Thus, the volume at which people were heard would decrease with distance, although it would still be possible to hear audio from the stage while in one of the outer worlds.

3.3. The presentation of the virtual world within the real world

An important aspect of our design was the ways in which interfaces to the virtual world were situated within different physical spaces and also the ways in which occupants of these various spaces were provided with some degree of mutual awareness to create a form of mixed reality as initially proposed in (Benford, 1996).

The event was distributed across two distinct physical spaces. The first was a fairly traditional performance space created in a cinema. This involved a seated audience facing a stage on which the physical poets performed. The second was a “cyber café” created in an adjacent café bar. Here, people could sit down at workstations in order to control an angel in the virtual world.

Our aim was to create two quite different environments, with people wandering between the two. The cinema was intended to provide a well managed and disciplined environment within which the poets could perform. The café bar, on the other hand, was intended to be a noisy and relatively unstructured environment. Several further linkages were made between these two physical spaces and the virtual world:

- A view of the virtual world was projected in the cinema alongside the stage so that the cinema audience could see the virtual poet embodiment moving alongside the physical poet and could also witness the behaviour of the angels. This view was controlled by a dedicated “camera person” using a specially created interface (see below). In addition, audio from the virtual world (i.e. conversations between the angels) could be mixed into the house PA in the Cinema.
- A view as seen by one of the angels was projected onto the wall of the café-bar so that all of its occupants (up to a hundred people) could keep track of events in the virtual world. The audio from the poets was also broadcast into this space.

To complete the mixed reality, it was initially intended to provide video views into the two physical spaces to the angels within the virtual world similar to those

used in the Internet Foyer application described in (Benford, 1996). However, this idea had to be dropped due to lack of time, concern about the impact on the network and the lack of suitable available hardware.

3.4. The temporal structure of the event

Given budgetary constraints and hence the limited availability of workstations, one of our biggest problems was how to give an exciting experience to as many people as possible. We also needed to balance the needs of the audience in the cinema, who would be experiencing a fairly conventional style of performance, against those of the virtual audience members who needed to mix attending a performance with exploring the outer worlds. These issues all related to the temporal structure and rhythm of the event.

Our solution to this problem involved a cycle of performances. In the cinema, each cycle was an hour long: fifteen minutes change over; fifteen minutes performance from a poet; a short change followed by fifteen minutes from a second poet; and finally fifteen minutes "touring" where the cinema audience would be taken on a tour of the virtual world and would observe the angels exploring and interacting with one another. The cycle was quicker in the café bar involving a half-hourly change over. Each batch of ten angels would either explore for fifteen minutes and then experience fifteen minutes poetry (corresponding to the first half of a cinema cycle) or vice versa (corresponding to the second half). The whole event involved three cinema cycles, allowing sixty people to experience the virtual world. These were preceded by an introductory talk and followed by a discussion session.

In order to manage smooth change overs within the Cinema and Café Bar, additional workstations dedicated solely to management tasks were located in each space. These were used to start up and shut down different processes (e.g. user clients, poets embodiments and the world itself) and to control the display of change over notifications on different workstations. Extensive use was made of UNIX talk between these two machines in order to coordinate the overall event. These machines were continually manned by technicians. Other technicians were provided to help with the tracking equipment at the stage and each angel workstation was also assigned its own assistant.

3.5. Navigation interfaces

Given that our system had to be used by members of the public with no previous experience and given some of our previous observations concerning the difficulties of navigating CVEs (Greenhalgh, 1995), the design of the user interface was assigned great importance. In fact, we designed and implemented two quite different navigation interfaces for the event, intended for use by different kinds of participants: an angel interface for members of the public and a camera interface to be used to drive the projected view in the cinema.

We begin with the angel interface. We decided to allow the angels to be able to fly up and down off of the ground plane and to rotate forwards and backwards, even though the whole event could quite easily have been experienced from ground level in the virtual world. However, we took several steps to minimise the risk of getting lost and disorientated. First, we separated all controls for movement off of the ground plane into separate areas of the interface. We provided buttons to reset position to ground level and orientation to straight ahead and we also retained the idea from MASSIVE of providing different camera angles, giving both in and out of body views (straight ahead, over the shoulder, mirror and birds-eye were provided).

A particular concern was how to encourage people to move about the virtual world at appropriate times (e.g. how to easily get them to the outer worlds for exploring and back to the stage in time for a performance). We introduced two mechanisms to support this. First, we defined a number of travelators from the stage to the outer worlds (visible as grey pathways in figure 1). People moving onto one of these would gently drift towards an outer world unless they deliberately moved otherwise. Second, we provided a "back to seat" button which, when continuously held down, would turn its angel to face the stage, move them back to their seat and then turn them to face the centre of the stage in a smooth animated movement. In this way, we hoped to make it easy for people to move to appropriate places without taking away their autonomy of control.

We also implemented two further general navigational aids: a 2-D plan view of the world with a zoom function and with labelled positions of the other angels; and stereo panning of audio so that people to one's left would be heard on the left hand side of the stereo mix and vice versa. We did consider the use of "solid" objects as another way of avoiding disorientation, similar to the approach of games such as Doom where one is constrained by solid boundaries to move through corridors and other enclosing spaces. However, given the relatively open structure of our world combined with the irregular design of the outer worlds and the relatively high computational cost of collision detection, we decided that this would not be an appropriate or cost effective solution for this particular event.

The camera-control interface was quite different. First, it was defined so that all movements were about a single point of focus - the stage area. Thus, the camera always pointed towards the performers and the only available movements were various rotations and zooms about a central point. The camera person could navigate by specifying a destination for the camera and then requesting a smooth animated transition to this new location. They could also create a list of stored positions and could animate transitions to these. In this way, we ensured that the projection interface in the Cinema always kept the performer in view and also that it offered smooth transitions between attractive viewpoints (e.g. sweeps around the performer, zoom into a poets face, move to a birds-eye view and so on).

This concludes our discussion of the design of the poetry event and the various features that were introduced in order to help manage participation and navigation. In the following section we reflect upon what actually happened at the event.

4. Experience

The following observations have been derived from a variety of sources including video recordings of the event; informal interviews with participants during the event; a final public discussion session involving the performers, artists, technicians, producers and members of the audience; and our own reflections. In many ways we consider the event to have been a great success, especially given its innovative and experimental nature. The virtual audience members clearly enjoyed themselves, there was plenty of social interaction between them (perhaps too much, as we shall see below!), they seemed to learn to navigate reasonably well, the software and hardware performed to our expectations and the management of the event and change overs progressed fairly smoothly. However, we did experience a number of unforeseen problems which suggest new design possibilities. It is these problems that we focus upon in this section.

4.1. The angels ignoring the poets

The most obvious and major problem with the event was that the angels tended to ignore the virtual poets. They often moved away from their seats within seconds of a performance beginning; they constantly spoke to one another over the poetry; and they sometimes even wandered into the centre of the stage during a performance. We suspect that there were a number of causes for this behaviour, spanning social, production and technical issues.

A key technical problem was that the poets were simply not loud enough in the virtual world. First, there were some problems with getting a clear audio feed from the house PA system into the virtual environment. Second, giving the angels the same foci and nimbi as the poets may have over-empowered them. Giving them smaller nimbi would have reduced their ability to interrupt the poetry. However, this step might also have reduced the potential for encounters and social interaction between angels during the exploration phases of the event (our initial design was based on the assumption that relatively large foci and nimbi were needed to ensure adequate levels of awareness for encounters to occur between only ten participants distributed across four outer worlds). Ideally, we would have needed to dynamically adjust the size of the angels' foci and nimbi at different stages of the event or at different locations, contracting them during the poetry and expanding them during exploration.

A second technical problem concerned the navigation interface. It was very easy for the angels to leave their seats - much easier than it would be in a conventional physical performance space. A physical theatre seat constrains one's view and movements to a much greater degree than our virtual seats which basically just marked a home position. On reflection, it may have been better to have given the angels a more performer centred navigation control during the performance, similar to that used to control the projected display, effectively forcing them to keep the

poet in view as they moved about. We could also have made it more difficult for them to leave their seats or perhaps could have introduced a gravitational pull back to their seats during the performance based on the travelator mechanism used to transport them to the outer worlds. However, like the reduction in nimbi discussed above, these changes would probably have been inappropriate for the exploration phases of the event. From a social perspective, it may be that, given the novelty of the multi-user VR experience for our audience members, the temptation to fly around and talk to other people was just too great. In many ways, the problem might even be seen as indicative of the attractiveness of the idea of social interaction within CVEs. This may have been exacerbated by the temporal structure of the event - half of our virtual audiences were expected to sit still for the first fifteen minutes of their experience.

Finally, there was clearly a lack of feedback from the poets to the audience members as we shall discuss in the section below.

It is interesting to note that we did not experience this problem when rehearsing with the system in the laboratory. Of course, our audience members for rehearsals tended to be more experienced users, were more personally involved in staging the event, were possibly more familiar with their expected behaviour, were in a different social setting and had consumed less alcohol. There were also fewer angels present in rehearsals. However, whatever the causes, the behaviour of the angels on the night would have been difficult to predict in advance and may also have operated in a vicious circle; once some people started to talk over the poets so it became harder for others to pay attention. Again, this points towards a requirement to be able to dynamically influence social interaction, perhaps through dynamic manipulation of spatial model mechanisms such as focus and nimbus, as an event unfolds.

4.2. The poets ignoring the angels

A contributing factor to the above problem was undoubtedly that the virtual poets were clearly not addressing the angels. Although the virtual poets appeared in the performance space, gestured and spoke, they were driven by the physical poets who were performing to the physical audience in the cinema. Consequently the poets did not, and indeed could not, directly engage individual angels using gesture or gaze. Given their limited awareness of what was going on in the virtual world (the projected view could be seen to the side of them, but only with difficulty), the poets could not even refer to events in the virtual world. Thus, there was no reaction from the virtual poets when the virtual audience talked over them or otherwise behaved "badly". In short, it was too easy to ignore them without appearing to cause offence.

In some ways, this problem stemmed from a deliberate design decision. We had decided that, given their lack of experience with the system, it would be safest for the poets if they could perform to a live audience. As an alternative, we could have structured the event so that they performed to the virtual audience and were only observed by the physical one. The inclusion of video views of the physical poets

and audience in the virtual world as originally aimed for might also have helped with this problem. Either way, the result of our strategy was that both the physical performance space and the virtual world worked well as spaces in their own right, but not really in a properly integrated way. To put it more formally, an integrated mixed reality was not successfully created. One of the causes behind this problem may have been the lack of a stable and consistent spatial frame of reference that spanned both the physical and virtual spaces. In particular, the use of the camera interface in the Cinema with its changing viewpoint, although aesthetically pleasing, made it difficult to establish any consistent reference between the two spaces.

4.3. Embodiment and symbolism

One of the key issues raised in the post event discussion was the symbolism of the embodiments and of the virtual world design. As black poets, much of whose work dealt with issues of race, the performers were obviously highly interested in the symbolism of their embodiments. So was the artist (who incidentally was white), but with a different perspective; he was clearly interested in the aesthetics of the environment and also in reflecting a specific design concept. In this case, he designed the world around a religious theme of angels (audience members) and devils (the poets). This clearly raised some concern with the performers! Much of this had completely passed by the computer scientists (also all white) who judged the poets embodiments on their aesthetic value, with little consideration of their symbolism. Some broader observations arise from this. First, as we have seen before with CSCW, inter-disciplinary working raises new problems. In this case, we were seeing quite different perspectives on a design issue from performers, artists and technologists, all of whom were essential to the event. Although there had been some discussion between all of these parties on the design of the embodiments, resulting in several changes, the issues of their symbolism had not been raised! Second, this problem raises yet another issue to be considered by embodiment designers in addition to those discussed in previous papers (Benford, 1995; Bowers, 1996).

4.4. Problems with starting-up

Our final problem involved starting up the system simultaneously for our ten audience members. Due to concerns about the stability of the software if run uninterrupted for several hours, we decided to restart the system at every audience change over. Although we had run with this many users before, we had never started them all simultaneously. The result of so doing was severe delays in initialising processes, time-outs and system crashes. The short term solution to this was to stagger the start-up for the audience members. This would have been much more difficult over a wider area network due to the large amount of coordination involved. Although some of the problem was down to specific software problems with MASSIVE-2, we suspect that initialisation for large numbers of users is going to be a problem area for many CVE systems.

5. Towards new design possibilities

Although interesting in themselves, the problems described above are of particular relevance because they offer general insights into the development of CVEs. In this section, we explore some of the resulting design possibilities.

5.1. Object centred interaction

One way of focusing an angel's attention on the poets might have been to provide them with a similar navigation interface to that used with the projected display, where all movements are relative to a common point of focus. We might generalise this to a notion of "object centred navigation" - a new way of navigating in CVEs which involves designating objects as a focal point and then moving about them. Such objects might be elements of the world structure (e.g. a stage), data objects in an information visualisation, elements of a 3-D design or even other participants. We have recently extended MASSIVE-2 to include such a navigation interface. Our interface allows a user to dynamically designate an object in the virtual world to be their current navigation focus. They can then move towards this object, away from it or around it. At any time they can select another object and make this their navigation focus. Other aspects of interaction might also be dealt with in an object centred way. For example, "object centred awareness" would involve several people who were locked onto a common focus experiencing heightened mutual awareness (e.g. two people focusing on the same element in a 3-D information visualisation might automatically get an audio channel allocated to them). In fact, the mechanism of third party objects from the extended spatial model already provides lower level support for this idea; a third party object is able to adapt (amplify or attenuate) existing awareness relationships according to how aware it is of the other objects or vice versa (Benford, 1997). Thus, a third party could amplify the mutual awareness of objects that were highly aware of it (i.e. were focused upon it).

Consistency might also be dealt with in this way; objects dealing with a common focus might experience a more consistent world view than those with different foci. In general, what we are proposing is a general consideration of how common objects of interest might be used as the basis for mediating different aspects of interaction in CVEs, including navigation, awareness, consistency and perhaps others. In its most general form, we might refer to this idea as "object centred interaction".

5.2. Context driven interaction

The poetry event involved the angels in two distinct activities: attending the performance and exploration. Each of these has different interaction requirements. For example, the performance requires smaller foci and nimbi for the angels and performer centred navigation whereas exploration may require larger foci and nimbi (to enable communication among a more widely dispersed population) and a more

traditional 3-D style of navigation. To generalise, aspects of interaction such as awareness, navigation, consistency and so forth might best be dealt with in a context or activity sensitive way. Two especially interesting aspects of context are: spatial context (i.e. where people are) and temporal context (i.e. what stage is an event at). In other words, it should be possible to trigger changes in the interface and system functionality according to both spatial location and temporal events. As noted above, aspects of adapting focus and nimbus according to spatial context have already been addressed through the third party extension to the spatial model (Benford, 1997). Our experiences with the poetry event suggest that third parties might be extended to influence other aspects of interaction such as navigation. We also need to develop models of the temporal structure and flow of events which can be used to manage interaction. Given such models, we can consider extending virtual world modelling tools such as (Colebourne, 1996) to include systems support for both temporal and spatial structuring of events in CVEs and the management of transitions between them. Finally, dynamic management of such facilities is required so that event producers can react to evolving circumstances and behaviours.

6. Summary

This paper has discussed the design of a public poetry performance in a CVE. At the heart of this design was the objective of enabling the virtual audience members to be autonomous and socially active participants in the event. Our design addressed a broad range of issues including developing a virtual world structure that would support a combination of performance and exploration; developing appropriate embodiments for both the performers (a semi-immersive interface to support gesturing) and the audience members; adapting mixed reality concepts for the display of the virtual world within two quite different physical spaces; determining an appropriate temporal structure for the event; and developing different navigation interfaces for the audience members and for projection interfaces.

A number of interesting issues were raised by the actual staging of the event. First, for a combination of technical, social and production reasons, it transpired that the virtual audience and performers practically ignored one another. Second, the post-event discussion revealed that the design of the various embodiments, especially their symbolism, was an important and controversial issue. Finally, a number of specific technical problems were revealed such as those caused by simultaneously starting the virtual environment for ten audience members.

Several more general conclusions can be drawn from this work. The design of this kind of event requires that many different issues are addressed against a backdrop of very specific technical constraints (networking, performance of computers etc.). *The staging of this kind of event in a public forum provides a powerful focus for defining and dealing with such issues. It is also possible to generalise the lessons learned from such experiences. In this case, section five raised to possibility of new*

approaches to managing interaction in CVEs, "object centred interaction" and "context centred interaction".

Finally, the staging of this event has required extensive collaboration between computer scientists, artists, performers and producers. Although not without its problems, this has been an extremely positive experience. Furthermore, as the technology of CVEs matures and is directly taken up by the public (as we believe it will be within the next few years), so bringing artistic and production expertise into the development of CVE applications will become increasingly important. Thus, "content will be king", and events such as the one described in this paper will become an increasingly important aspect of research into CVEs.

Acknowledgements

We would like to thank the EPSRC and the European Commission (ACTS programme) for funding the development of the underlying technologies and The University of Nottingham and Nottingham City Council for funding for the event itself. We would like to thank artist Sean Varney; poets Dave "Stickman" Higgins, Akure Wall, Roger Goodwin and Remi Abbas; producer Andrew Chetty and all of the NOW'96 team for working with us on creating and staging this performance. We would like to thank Jolanda Tromp and Mike Fraser and other colleagues at Nottingham for helping out during the event and for their input into its design.

References

- Benford, S. D., Bowers, J. M., Fahlén, L. E., Greenhalgh, C. M. and Snowdon, D. N., User Embodiment in Collaborative Virtual Environments, Proc. CHI'95, May 7-11, Denver, Colorado, USA, ACM Press
- Benford, S. D., Brown, C. C., Reynard, G. T. and Greenhalgh, C. M., Shared Spaces: Transportation, Artificiality and Spatiality, Proc. CSCW'96, Boston, pp. 77-86, ACM Press.
- Bowers, J., Pycock, J. and O'Brien, J., Talk and Embodiment in Collaborative Virtual Environments, Proc. CHI'96, ACM Press, 1996.
- Benford, S. D., Greenhalgh, C. M. and Lloyd, D., Crowded Collaborative Virtual Environments, Proc. CHI'97, Atlanta, Georgia, US, March 22-27, 1997 (in press).
- Colebourne, A., Rodden, T. and Palfreyman, K., VR-MOG: A toolkit for building shared virtual worlds, Conf. FIVE Working Group, London, 1995, 109-122, M. Slater (ed).
- Damer, B and Bruckman, A (moderators), Peopled Online Virtual Worlds: A New Home for Cooperating Communities, A New Frontier for Interaction Design (panel), Proc. CSCW'96, Boston, USA, Nov 16-20, 1996, pp 441-442, ACM Press.
- Fahlén, L. E., The Synthetic Hyper Theatre: The Conceptual Symbiosis of VR, Art and Aesthetics, Invited paper Conf. FIVE Working Group, London, 1995, 109-122, M. Slater (ed).
- Greenhalgh, C. M. and Benford, S. D., Virtual Reality Tele-conferencing: Implementation and Experience, Proc. ECSCW'95, Stockholm, September, 1995, North-Holland.