

Orchestrating Real-Time Participatory Experiences

**Steve Benford, Adam Drozd, Chris Greenhalgh, Mike Fraser, Holger Schnädelbach,
Boriana Koleva, Rob Anastasi, Martin Flintham, Terry Hemmings, Andy Crabtree**

The Mixed Reality Laboratory
The University of Nottingham,
Jubilee Campus, Nottingham, NG8 1B, UK
{sdb, asd, cmg, mcf, hms, bnk, rma, mdf, tah, axc} @cs.nott.ac.uk

Ju Row-Farr, Matt Adams, Nick Tandavanitj

Blast Theory
Unit 43a, 4th Floor, Regent Studios
8 Andrews Rd, London E8 4QN
{ju, matt, nick}@blasttheory.co.uk

John Bowers

Centre for User-Oriented IT Design (CID)
Royal Institute of Technology,
Stockholm, Sweden
bowers@nada.kth.se

Dirk vom Lehn, Christian Heath

The Management Centre
Kings College London
Franklin-Wilkins Building London SE1 8WA, UK
{dirk.vom_lehn, christian.heath} @kcl.ac.uk

Abstract

Real-time participatory experiences such as multiplayer games, on-line role-play and performances require orchestrating; an on-going process of managing participants' activities from behind the scenes. We describe how orchestration was achieved in three contrasting experiences: an on-line role-play called Avatar Farm, a touring mixed reality performance called Desert Rain, and a mobile mixed reality game called Can You See Me Now?. We show how each adopted a different approach to orchestration according to its goals and operating constraints, including the use of dedicated interfaces for following and manipulating characters and objects in virtual worlds. A comparison of the three draws out common issues for orchestration including: admission to an experience, training and familiarisation, establishing engagement, avoiding fractures in engagement, monitoring, intervening, coordinating behind the scenes activity, managing pace and timing, and closing an experience. We summarise techniques and tools that are available to designers to address each of these issues. Finally, we reflect on how orchestration needs to respond to the challenges of increasing scale and mobility of future participatory experiences.

Note to the reviewer. This paper builds on and extends two conference papers: (Koleva, 2001) and (Drozd, 2001).

1. Introduction

In *Computers as Theatre*, Brenda Laurel proposed an approach to interaction where computers are considered as a form of theatre rather than as tools, and where the focus of design is on engaging users with content rather than with technology (Laurel, 1992). She suggested that various behind-the-scenes activities are required to maintain engagement and to orchestrate users' experiences.

Nearly ten years later, collaborative virtual environments (CVEs) are being used for on-line games, performances, role-play and other leisure and entertainment applications (Dodsworth, 1997). These applications are quite literally computers as theatre and so have to deal head-on with the challenge of orchestration. What activities are required to ensure the smooth running of an experience, and how can participation be guided and shaped, especially in real-time?

One approach, directly derived from traditional theatre and television, is to employ a production crew to monitor events and to intervene where necessary. We first experimented with this approach in *Out of This World*, an inhabited television show in which members of the public and professional actors staged a gameshow in a CVE that was also broadcast to a viewing audience (Benford, 2000). Like most conventional live television shows, *Out of This World* was driven by the desire to create fast-pace action that could be captured live by (virtual) cameras. Our producer and director expected the show to unfold to a tightly defined schedule and participants to be "on their marks" at prescribed locations in the virtual world with precise, to-the-second timings. Previous experiences had shown how difficult it can be to achieve fast-pace coordinated action in a CVE (Benford, 2000). Consequently we opted for a tightly managed approach in which the entire show was structured as a series of phases. Each phase applied its own movement constraints to the participants, allowing us to move them around the world, limit the extent of their exploration, or root them to the spot. A member of the production crew monitored the show and under instruction from the producer and director, manually triggered its different phases. This approach was generally successful in terms of creating a coherent and relatively fast-pace experience and demonstrated the importance of developing new orchestration techniques and supporting tools.

This paper takes an in depth look at the process and technologies of orchestration. It begins by describing three subsequent explorations of orchestrating real-time interactive experiences. The first was an inhabited television show called *Avatar Farm*. In contrast to *Out of this World*, *Avatar Farm* had a more improvised and non-linear form, raising new challenges for orchestration. The second was a touring mixed reality performance called *Desert Rain*. This provided an intimate experience for a small group of players who journeyed through a series of interconnected physical and virtual spaces, and so raised the issue of orchestrating activity that spans physical and virtual environments. The third was a mobile mixed reality game called *Can You See Me Now?* in which up to twenty on-line players were chased across a map of a city by three performers who were running through its streets. This experience shed light on the process of orchestrating an activity where participants were physically dispersed and where some of them were mobile. We describe the content and implementation of each and provide an ethnographic account of how orchestration was achieved in practice. Following this, we compare the three experiences in order to draw out more general issues, design guidelines and future challenges for orchestration.

2. Avatar Farm – a non-linear, improvised drama in a virtual world

Out of This World been criticised for a lack of empathy with its characters and for directly adopting a clichéd format from conventional television. Our goals for *Avatar Farm* were therefore to engage members of the public in a more richly dramatic experience and to explore the potential of CVEs to support new narrative forms.

2.1 An overview of Avatar Farm

Avatar Farm drew its inspiration from a previous virtual community called The Ages of Avatar (Craven, 2000). This provided us with a resource of existing virtual worlds and players who shared a common history. We recreated the virtual worlds in the MASSIVE-3 system, running on a dedicated local area network in our laboratory at Nottingham. We invited four key members of the community to join us for two days in June 2000, where they collaborated with seven professional actors, a storywriter and a production crew, to improvise a drama loosely based around their familiar characters and worlds. We chose actors who were experienced in engaging members of the public in more conventional role-play situations and improvised theatre. The result of these activities was Avatar Farm, a two-hour long improvised drama structured as four 25 to 40 minute long chapters, involving 15 virtual characters, played by 11 people that was both web-cast live and also recorded.

Avatar Farm was a fable involving gods, tricksters and innocents abroad. The four players from the Ages of Avatar were reawakened in the familiar virtual worlds to find that they had been repossessed by some feuding gods and their sidekicks. The players were initially enslaved as pawns in the gods' struggles. In chapter one, the players were reawakened and immediately separated and taken to different worlds to meet the gods for the first time. Chapter two involved the players learning how to gain special powers such as flying, changing appearance and becoming invisible. They also learned how to trigger time-rifts – ghostlike playbacks of scenes from the past (part of a backstory that had been recorded by the actors on previous days using MASSIVE-3's 3D record and replay mechanism (Greenhalgh, Purbrick, Benford *et al.*, 2000)). In chapter three, the players' loyalties to one another were tested to the point where they rebelled. Finally, in chapter four the players, rose up to overthrow the villains of the piece.

From the point of view of this paper, a key characteristic of this drama was its non-linear narrative structure. The core of the story was based upon the four players' experiences. For much of the time they were separated and involved in parallel scenes that were taking place in different worlds. As each followed their own thread through the story, their paths would cross at various points and occasionally they would all meet for a pivotal scene before splitting up again. Even when not directly involved with the four players, the actor-controlled characters remained active, carrying out their normal background activities.

A second key feature was the degree of improvisation involved. Before the event, our storywriter had planned an overall structure for Avatar Farm in considerable detail and had fleshed out the backstory, the actors' characters and the effects of key virtual objects. However, individual scenes were to be improvised rather than scripted (especially in terms of audio dialogue) and it was an open question whether the public players would follow the intended script, would chose to go their own ways, or would even cooperate at all.

Finally, the plot required the players to collect and use various virtual objects in order to achieve special effects that were central to the story. These interactions involved relatively complex sequences of utterances and gestures and could have quite profound effects on the world and especially on other players' experiences.

2.2 Improvising magic – our approach to orchestrating Avatar Farm

Avatar Farm's non-linear and improvised nature posed problems for orchestration as the crew had to monitor and manage concurrent scenes in different locations, cope with a wide variety of unforeseen circumstances, and react to developing storylines on the fly. Realising some of the more extreme effects of virtual objects required marshalling considerable human and technical resources and was subject to resource constraints. For example, creating a time-rift required concerted action by several people and our technical set-up meant that only one time rift was possible at any moment, even if several had been requested by the players. Given these challenges, we chose an approach to orchestration that combined three key elements.

1. We adopted the metaphor of ‘improvising magic’, where players would perform incantations in order to request actions that would then be realised by invisible stagehands.
2. We assigned different crew members different responsibilities, and provided them with ways of monitoring events in virtual and physical spaces and of communicating with one another.
3. We provided dedicated interfaces to allow crew members to manipulate virtual objects and grant and revoke special powers to the players.

The following sections focus on each of these three elements in detail.

2.2.1 Granting capabilities and improvising magic

At the start of Avatar Farm, the players were able to perform only a few basic actions with their avatars. These were: moving around at floor level; talking; picking up an object, waving it about and putting it down; carrying an object while moving; and replaying one of ten pre-recorded gestures. At times, the story required the ability to limit even these basic actions; for example, particular players might occasionally be captured and frozen to the spot or disallowed from picking up certain objects.

Central to the story was the way in which the players gained additional capabilities and learned how to invoke various magical effects. These included: flying up to a fixed height; becoming invisible; changing appearance between a number of pre-determined avatars; moving through the portals that linked the four worlds together; snooping on other players’ distant conversations; becoming immune to the powers of particular gods; and triggering a time-rift. In terms of the story, these capabilities and magical effects might be granted by other characters, especially the gods, or might arise from the correct use of particular objects.

One approach to supporting these effects would have been to program them directly as part of the Avatar Farm application software. However, we were concerned to make the structure of Avatar Farm as open as possible to improvisation so as to be able to take advantage of interesting, yet unforeseen, interactions between participants. To allow flexibility for this, we wanted to be able to choose at any moment whether it would be appropriate to grant an effect and if so, exactly how and when it should be realised. In this way, the timing of an effect could be controlled to fit in with ongoing interaction between participants. Indeed, whether an effect is granted at all could also become a dramatic element. Some effects would also require coordinating multiple players. For example, a time-rift would be a major moment in the story and it would make sense to gather several players together to witness it. However, this would involve finding these players (who might be engaged in activities elsewhere) and persuading them to move to the location of the time-rift. It would be difficult to predict how long this might take. Finally, we were well aware of the possibility that the coordination of the narrative might break down – especially with such a multi-threaded structure. We needed an approach which would enable us to repair and recover from breakdowns. Pre-programming objects might have hindered this if, for example, a set of behaviours were to execute autonomously and erroneously.

In light of these considerations, we adopted an alternative strategy – improvising object-interactions. In this case some crew members, subsequently referred to as stagehands, were also present but invisible within the worlds. They had the ability to manipulate objects and avatars and to directly trigger special effects such as replaying pre-recorded scenes; constraining players; making players and objects appear, disappear and change appearance; and granting and revoking permissions for players to pick up certain objects and move through portals. The invisible stagehands followed the players around, monitoring their activities and triggering effects in response. Improvising a single logical action from the point of view of the players would often involve the stagehand in a quite complex sequence of actions, and actions involving more than one player might require the coordinated effort of several stagehands.

We decided to dress up the process of improvising interactions in the metaphor of magic. The players would act out elaborate rituals, gathering objects, placing them in key locations and making extended sequences of movements, gestures and utterances, in order to invoke an effect. Even where the players gained new abilities such as flying, these would be granted in a magical way – as a gift bestowed from the gods or as a result of an incantation. We anticipated two key benefits from this approach.

1. The result and timing of any request could be left open – everyone knows that magic is dangerous, unpredictable and liable to go wrong if the magician makes only the slightest mistake. We hoped that the metaphor of magic would enable the players to accept and work around delays and failures.
2. Extended sequences of actions would be more visible, predictable, dramatic and therefore interesting to watch for viewers. Stagehands would have sufficient time to spot that a request was being made, to marshal the necessary resources and to plan their response.

The following sections describe the organisation of people and technology that supported this approach of collaboratively improvising magic within Avatar Farm.

2.3 Behind-the-scenes of Avatar Farm

There were two broad categories of people involved in Avatar Farm: the cast (four players and seven actors) and the production crew. The latter consisted of the following:

- **Story director** – responsible for directing the story, including monitoring the progress of the event as a whole, deciding on the course of the plot, and instructing actors and crew members accordingly.
- **Director’s assistant** – supported the story director and coached the actors.
- **Software manager** – assumed overall charge of the MASSIVE-3 system.
- **Stagehands** – four invisible helpers who were charged with the task of improvising interactions as described above. Each was assigned to follow a different player, although the story director might assign them other duties. One was also responsible for cueing and replaying the pre-recorded time-shifts.
- **World-manager** – a further invisible helper who was responsible for granting and revoking access controls on portals, thereby controlling which characters could move into which worlds at which times.
- **Actor helpers** – two crew members who helped two actors who were using immersive VR interfaces, for example, helping to put on and take off the equipment and also holding their microphones.
- **Player helpers** – two crew members to support the four players.

These people were located in a shared studio space as shown in Figure 1. Several features of this arrangement are relevant to this paper. First, the only partitions in the space were black curtains. As a result, the players could not see the behind-the-scenes production areas when the world was live, but there was some potential for audio overspill, which also meant that the production crew had to be careful not to talk loudly or make other noises. Second, the space was designed to encourage mutual awareness among key members of the crew. In particular, the story director, their assistant, world manager, helper responsible for temporal links and camera crew were arranged facing across a shared table so that they could peripherally monitor each other’s affairs and floor manager and actor and player helpers could move freely around the space.

In addition to the physical design of the studio space, the MASSIVE-3 software was also configured to allow different roles to oversee events in the virtual world and to communicate. The story director and assistant were invisibly present within the worlds and they and the stagehands, world-manger, camera operators, camera director and software manager could monitor the conversation between the actors and the players in the part of the world where they were currently located. A separate audio talk-back system allowed the story director to speak directly to any combination of stagehands, actors or to the world manager in order to pass out instructions. In other words,

there were many opportunities, both in terms of on-line communication and the design of the physical studio space for the various crew members to monitor events within the world and to communicate with one another.



Figure 1: views from the studio:

- (a) the story director (foreground) with the row of actors and two helpers and curtain separating the players beyond (background)
- (b) looking over a player's shoulder (foreground) with the curtain open towards the actors (midground) with the remaining crew in the background
- (c) the central table with helper and world manager (foreground), story director, and assistant (background)

Figure 1: views of the inhabited television studio for Avatar Farm

2.3.1 The stagehand and world manager orchestration tools

The stagehands and world manager were provided with dedicated orchestration tools. The stagehand interface consisted of two windows, one containing controls for manipulating entities (objects or avatars) as shown in Figure 2 and a second offering a view of the relevant virtual world.

A stagehand would select an entity to be managed from the list in the lower part of the interface. Upon selection their view of the world would be moved to centre on this entity. The stagehand could zoom and rotate this viewpoint while focused on this entity using the camera controls at the bottom right of the control panel. The associated virtual camera would lock onto and track the entity as it moved. Once selected, the entity could be managed.

The constraint control manoeuvred the entity around the world with the speed of movement being governed by a slider on the control panel. Depending on the type of entity selected, different properties could then be altered using the controls at the top-right. If the entity was an object, the stagehand could select whether it was visible or not. For an avatar they could:

- set whether it was visible or invisible;
- select its appearance from a among a pre-defined selection of geometries;
- alter the scale factor of its geometry (making it grow and shrink);
- grant or revoke its ability to fly, control its own visibility and appearance.

The world manager interface was similarly split into two parts. They could select a world to view and could position their viewpoint either relative to the origin of the world or to a specific entity. To change the access control on either a portal or an object, they would select the portal or object from a list, select an avatar from a second list, and then set whether this avatar had access.

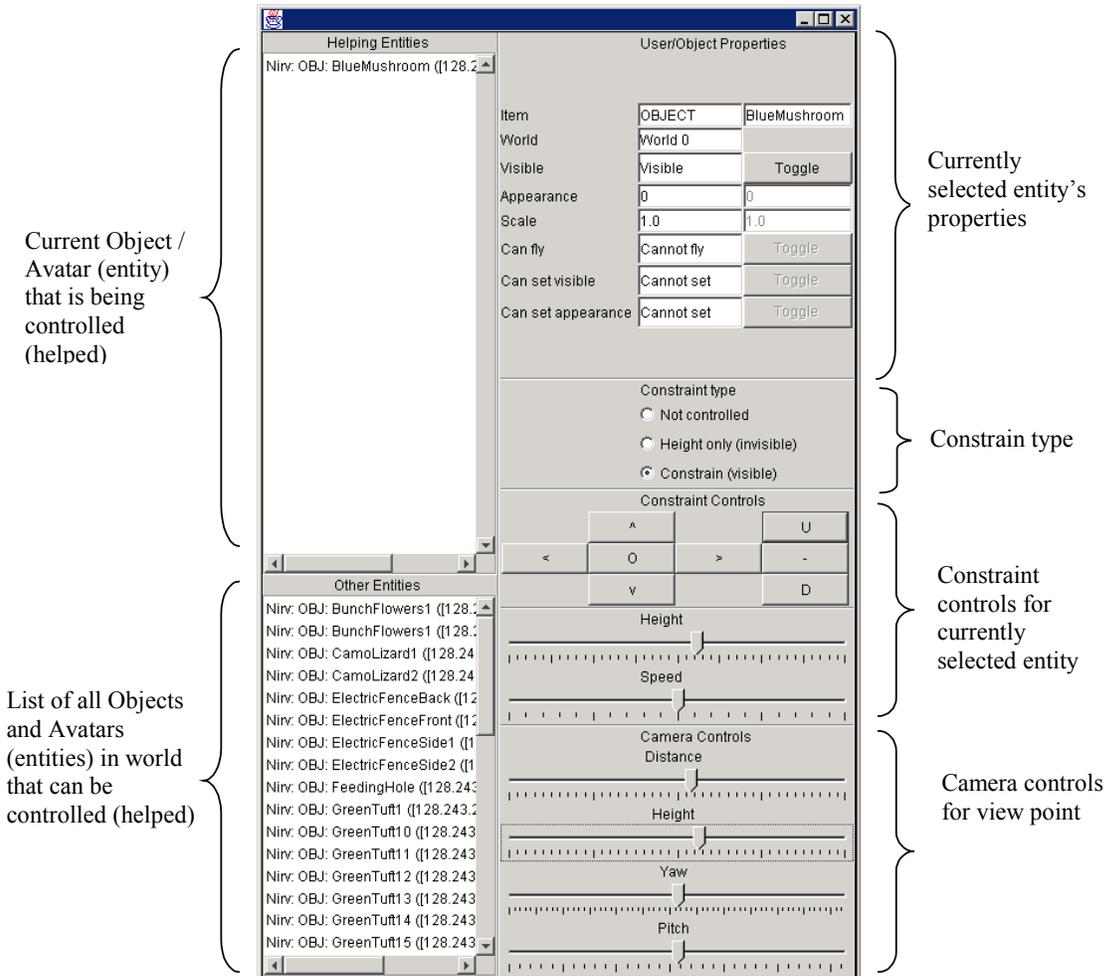


Figure 2: the stagehand interface

2.4 Avatar Farm in practice: the management of magic

To appraise our approach to the orchestration of Avatar Farm, we now focus on a specific example of the process of improvising magic at work. The example is taken from chapter 4 of Avatar Farm. It takes place within a child-like world called Kindergarten and involves three main characters: the player Maple, Squeaky Clean (a scheming trickster) and Botchov (an authoritarian butler). We focus on a sequence that lasts for approximately fifteen minutes in which Squeaky persuades Maple to play a trick on Botchov, causing him to be eaten by the world-serpent. This is one of the more complex sequences in the drama in terms of orchestration.

2.4.1 Using 3D Record and Replay to Analyse Avatar Farm

The following analysis of our chosen sequence has utilised a 3D record and replay mechanism that we have implemented within MASSIVE-3 (Greenhalgh, Purbrick, Benford *et al.*, 2000). This logs and timestamps every action (movement, object interaction, and speech) of every avatar within a locale (region of a world). A log file can then be replayed into a live locale at a later time so that the recorded action appears to be recreated within the live world and mixed with the live action. Live participants can fly around the recording, viewing it from any angle, following different characters and listening to the audio.

The whole two hours of Avatar Farm was captured as four 3D recordings, one for each chapter. To support our analysis, we altered the playback of these 3D recordings so that the stagehands were now made visible, enabling us to see their actions alongside those of the players and actors. We could then repeatedly view and hear the process of improvising magic from the perspectives of different participants. We also had access to a video recording of the physical studio space during this sequence. The virtual world screenshots in this paper were also generated from the 3D recordings.

2.4.2 How Maple and Squeaky Clean Tricked Botchov

The story director originally planned our chosen sequence to unfold as follows. The player Maple, aided by Squeaky, would feed a purple-tuft (virtual object) to the world-serpent (a mythical creature that lives under the ground and is never actually seen), the result of which (given a suitable incantation) would be to trigger a time rift. This would show them both a scene from the past in which Botchov was playing with his long lost sister Octavia and discussing various family secrets, especially “the secret of the green teapot”. Squeaky would then teach Maple how to change appearance by picking up and licking a camouflage lizard and performing a second incantation. Maple would then take on the guise of Octavia in order to trick Botchov. Maple, disguised as Octavia, would encounter Botchov. He would use the secret of the green teapot (that surely only Octavia could know) in order to convince Botchov that he is indeed Octavia. He would tell Botchov that he has to stand on the feeding hole if he wants to join Octavia in the land where she now dwells. Botchov would then be eaten by the world serpent, but not before Maple and Squeaky, had first revealed how he had been duped!!

In fact, as the following description shows, enacting this complex and interwoven sequence of events is far from straightforward due to various unforeseen circumstances involving other threads of the narrative that are taking place elsewhere as well as various local technical and interactional difficulties. However, Maple, Squeaky and Botchov do eventually manage to successfully improvise a version of the trick on Botchov, supported by the stagehands. The listing below summarises some of the key moments in this sequence of events along with our commentary as to what is happening behind-the-scenes. Figure 3 shows key moments from this scene. Maple appears as a red humanoid figure with brown hair. Squeaky has a green body, no legs and long ears and horns. Botchov is wearing a butler’s uniform. Octavia (Maple in disguise) appears as a girl with pigtails. Finally, the stagehands appear as cameras pointing at the character that they are currently controlling (though, remember, stagehands are rendered from the 3D recording here, they were not originally seen by participants). White rings around a character’s head indicate the current volume of their speech.

Event	Commentary
Squeaky Clean briefs Maple about the trick – see Figure 3 (a).	
Maple, guided by Squeaky Clean, begins the incantation to cause a time-rift.	Initially there is one stagehand on Maple, but they are soon joined by a second. Squeaky Clean’s dialogue at this point makes various explicit references to how Maple should use their

Squeaky Clean places the tuft on the feeding hole. Maple finishes the incantation. However, no time-rift occurs.

Squeaky Clean: “Try the incantation once more.”

Squeaky Clean continues guiding Maple through the second incantation.

Maple: “Has it worked?” (at the end of the second incantation)

Squeaky Clean: “No ... there is a time distortion already in place which is breaking up the equilibrium of the world”

Squeaky Clean now suggests that Maple use the camouflage lizard in order to turn into Octavia. He explains the necessary actions and incantations.

Maple grasps the camouflage lizard

Maple completes the relevant incantation.

Squeaky Clean: “... and you’ve changed into ...”

Squeaky Clean: “... Desmond! ...try again!”

Squeaky Clean: “... No ... Tock-Tock!! ...”

Squeaky Clean: “... No ... Octavia!!!. You are an incredibly powerful avatar to have gone through so many changes”

Maple: “Wow” (laughs).

Squeaky Clean now instructs Maple about using the secret of the green teapot to convince Botchov that he is indeed Octavia.

Botchov arrives. Squeaky Clean now makes himself invisible.

Maple: “I am the spirit of Octavia ...”

Squeaky Clean: (evil cackle)

Maple continues to act out the trick on Botchov. He reveals the secret of the green teapot and lures Botchov onto the feeding-hole.

computer to trigger the gestures that are required by the incantation (e.g., “press key 9”).

Squeaky Clean and Maple do not know that a 3D replay cannot be triggered at this point because another one is already taking place elsewhere as part of a parallel thread of the story. The Avatar Farm set-up does not include the resources required to replay more than one recording at a time.

A third stagehand has now become concerned with the progress of events and has attached themselves to Maple. All three can be seen in Figure 3(b). It should be noted that, in the performance itself, the stagehands are invisible to one another as well as to Maple and Squeaky Clean.

The second incantation hasn’t worked either as the parallel time-rift in Nirvana is still playing out.

Squeaky Clean has now heard (probably from the story director over the talk-back system) that there is a problem triggering the replay and is providing an account of this in terms of the narrative.

Squeaky is pressing on anyway. Again, his dialogue contains more references to pressing particular keys on the keyboard.

At this point one of the stagehands leaves the scene (see Figure 3 (c)), shortly followed by another, leaving just one stagehand on Maple. They change him into Desmond, the wrong character (he is supposed to become Octavia).

The stagehand now changes Maple into Tock-Tock (still the wrong character). The exclamation “try again”, of course, can be heard by Maple and the stagehand and understood as an instruction to both.

The stagehand now changes Maple into Otavia, correctly (Figure 3 (d)).

Squeaky Clean is once again accounting for technical problems in terms that make sense to the narrative at this point

Two more stagehands attach to Maple.

He needs to do this through dialogue as the replay that would have given this vital information was never seen. He also reminds Maple to talk in a high pitched voice. One of the stagehands leaves so there are now two in attendance.

Note: As an actor, Squeaky does not need a stagehand to do this for him.

Forgetting his instructions, Maple begins in a low voice but then switches to a higher register.

Squeaky Clean is using a stage-whisper here. Botchov and Maple will be able to hear this.

One stagehand now moves over to be on Botchov while the other remains on Maple.

The other stagehand now moves from Maple to Botchov (being unaware that a stagehand is already there). There is now no stagehand on Maple. The stagehands move Botchov downwards so that he is halfway into the hole

Squeaky Clean explains to Maple that he needs to reveal himself as Maple

Maple: “I am Maple not Octavia”

Squeaky Clean makes himself visible (Figure 3 (f))

Botchov acts surprised.

Maple: “Bye Bye”

Squeaky Clean: “Well done Maple”

Maple: “I think I got everything in there”

Squeaky Clean now engages Maple in conversation for a couple of minutes, recapping recent events. He then asks him whether he has ever seen the maze in this world.

Squeaky Clean: “Soon we will be going to Trade and Power”

(Figure 3 (e)).

Squeaky Clean again uses a stage whisper. However, Maple’s microphone connection has temporarily failed and it takes several such whispers before it resumes and Maple is heard to respond. Botchov acts as if he doesn’t hear, but Squeaky’s whispers help him to understand that there is a technical problem with Maple.

This is the cue to change appearance. One stagehand now moves back onto Maple. He switches Maple’s appearance back to being his normal self.

The stagehand on Botchov drags him entirely down through the hole and out of sight.

Squeaky is delaying Maple at this point. The main action will soon move to the world Trade and Power. However, the portals between worlds are all closed right now and so they cannot currently go there.

Squeaky Clean has received instructions from the story-director over talkback that he now needs to get Maple to the world Trade and Power where the cast is being assembled for the final climatic scene.

This sequence of events shows how the cast and stagehands struggle to overcome various unforeseen circumstances to more or less successfully improvise a version of the planned scene. Maple certainly manages to pull off a complicated trick on Botchov, albeit with extensive support from Squeaky Clean. However, this sequence of events also sheds light onto the ways in which improvisation occurs and the relationship between the work that is taking place “on stage” and the orchestration work that is taking place behind-the-scenes.

2.4.3 Making a drama out of a crisis

Previous studies of CVEs and other CSCW technologies have noted how participants often account and compensate for technical difficulties through their talk (Bowers *et al.*, 1996; Hindmarsh *et al.*, 1998, Hindmarsh *et al.*, 2001). In our example however, Squeaky Clean not only provides such accounts but carefully embeds them into the context of the drama. When a stagehand struggles to find the correct new appearance for Maple we hear that this is because he is: “an incredibly powerful avatar to have gone through so many changes”. Earlier, when it was not possible to replay a 3D flashback we heard that “there is a time distortion already in place which is breaking up the equilibrium of the world”. Indeed, it is often possible for an actor or player to formulate their contribution in such a way that it can be heard (by another actor or player) and overheard (by a stagehand or other crew member) simultaneously (e.g. Squeaky Clean’s instruction to “try again” above). We suggest that the approach of improvising magic provides skilled actors with plenty of room for manoeuvre when it comes to improvising such accounts and we are sceptical whether this would be so easy if interaction were more mechanical and immediate.

We also see examples of the (non-professional) players engaging in such creative accounting of events. In a scene a few minutes after our description ends, we see Maple improvising an account for another delay in a purple tuft triggering a time-rift. Squeaky Clean observes: “The purple tufts sometimes take a long time to work” to which Maple retorts “I should imagine so especially after he [the serpent] has had a long meal”. Such accounts, even or perhaps especially when ironic, allow the participants to maintain their engagement with the story while

providing improvised content which an actor could further develop, all the while covering a delay while the production crew troubleshoots a problem.



(a) Squeaky Clean (left) briefs Maple (right)



(b) Maple incants with 3 stagehands present



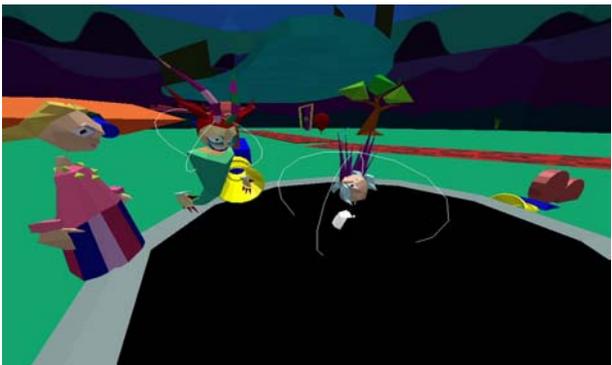
(c) Maple grasps the camouflage lizard



(d) Maple is changed into Octavia



(e) Botchov is lowered into the hole



(f) Squeaky Clean reveals himself

Figure 3: snapshots from the trick on Botchov

2.4.4 Coordinating help

At various times, up to three stagehands can be seen in the vicinity ready to help out with events. This testifies to the ability and willingness of the stagehands to monitor unfolding events, check up on them and be ready to help. However, in the example given, we see several moments where two or more stagehands have attached themselves to a single character, even though only one is necessary to bring about the desired effect. This suggests that

coordination among the stagehands was a significant problem. This is confirmed by noting that, at the crucial moment in the trick on Botchov, we see the two helpers on Maple both detach and move over to Botchov, when what was required was for one to change Maple back to his normal appearance as the other lowered Botchov into the hole. Such coordination involved the physical as well as the virtual environment. Our video recordings showed examples of people physically leaving their workstations to go over and talk to stagehands or check behind the curtain separating the players from the production crew, and so forth.

2.4.5 Delaying and hurrying tactics

As noted previously, Avatar Farm's branching, non-linear structure posed a number of challenges to its producers. In particular, a scene could be affected by parallel scenes that were happening elsewhere. This might be for technical reasons such as in our example when the time-rift cannot be triggered and the portals between worlds cannot be used because another time-rift is already happening. It might be for dramatic reasons, such as when a major scene takes place that requires the players to gather together. For whatever reason, local plans will often have to accommodate events elsewhere.

The approach of improvising magic provides some flexibility for managing the timing of local scenes. It is easy to prevent participants from triggering effects that would be dangerous or inappropriate and actors can employ various tactics to delay players or hurry them along. We see Squeaky Clean employ such tactics towards the end of our chosen sequence where he first reiterates the events which occurred, thereby delaying Maple from engaging in any subsequent activity while also making it clear to anyone who might be overhearing that the essential elements of the scene are completed. Squeaky Clean then changes pace, hurrying Maple along, after instructions have been received from the story director. Throughout, Squeaky Clean and Maple remain in character, improvising appropriate talk as they go.

2.4.6 Technological interaction within improvised talk

Our approach has been one of organising a narrative so that its improvised enactment contains adequate cues for behind-the-scenes personnel to realise that certain actions of technical significance need to be performed. For the most part this works implicitly in that talk about a purple tuft or a blue mushroom will be heard while this cues associated technical changes to be actioned. That is, typically, it is not necessary to directly refer to technical arrangements to get technical consequences. However, some exceptions to this worked rather inelegantly. In the above example, on several occasions we hear Squeaky referring to specific key presses. Players are required to use the keyboard to trigger up to ten different pre-canned gestures in particular sequences as part of making incantations. Such references stand out awkwardly against the general flow of the dialogue. Rather than refer to particular gestures by name (e.g. 'bow'), Squeaky explicitly instructs Maple in the key presses the player should use. There are several reasons for this. Although the association of keys to gestures was fixed throughout Avatar Farm, both actors and players had trouble fluently remembering it. To ensure correct performance, actors kept extensive notes close to hand and took to mentioning keys by their names rather than risk a misunderstanding. Hence, inelegant mentions of "key9" and so forth intrude the dialogue. In this and other ways, it seems that the complexity of the gesturing hindered the process of improvising. While extended rituals may be useful in slowing down the pace of interaction both for viewers and behind-the-scenes crew, they need to be designed carefully.

2.4.7 Now you see me, now you don't

A key feature of Avatar Farm is the way in which various participants and crew were made invisible and/or inaudible. Invisibility featured in the story itself, for example, when Squeaky Clean made himself invisible during

the trick on Botchov. It was also used to hide the stagehands from the players. However, these manipulations of visibility raised some interesting issues.

The stagehands were invisible to the actors, players and viewers front-stage, but also to one another and to the story director and other crew members. Even though they had been assigned to follow different players before the event, they appeared to find it difficult to coordinate their actions, as noted above. Naturally the fact that they were invisible to each other in the virtual world did not help in this.

The relationship between audibility and visibility also requires deeper consideration. Squeaky Clean's various stage whispers provide a good example. He cannot be seen, but his whispers can be heard by anyone nearby including Maple and Botchov. Botchov benefits from hearing them (even though the story says that he can't) as presumably they help him determine his own reaction to events. Maple can also hear them. Does he believe that Botchov cannot (it can be difficult to judge who can hear whom in a virtual world) or is he going along with the convention of the stage whisper? If these matters are ambiguous for us as analysts, then it is likely they were unclear to at least some of the participants too.

This concludes our discussion of Avatar Farm for the time being. In the following section we turn our attention to the orchestration of a second real-time interactive experience called Desert Rain.

3. Desert Rain – a touring mixed reality performance

Desert Rain was a professionally touring mixed reality performance that emerged from collaboration between the artists group Blast Theory and The University of Nottingham. It was a combination of performance, installation and computer game that took place across a mixture of virtual and physical environments.

Six players at a time were sent on a mission into a virtual world to find six targets (people who had had different experiences of the Gulf War). They explored motels, deserts and underground bunkers, communicating with each other through a live audio link. Once in the virtual world, they were given twenty minutes to find their allocated targets and complete their mission. They then left the virtual world to find their way to a final physical room where the identities of the targets were revealed. The virtual world was projected onto six rain curtains, screens made of water through which performers and players could physically pass.

The central artistic concern of Desert Rain was virtual warfare, the blurring of the boundaries between real and virtual events, especially with regard to the portrayal of warfare on television news, in Hollywood's films and in computer games. Inspired by Jean Baudrillard's assertion that the Gulf War did not actually take place because it was in fact a virtual event, both the content and the form of Desert Rain was designed to provoke participants to reevaluate the boundaries between reality and fiction, and between the real and the virtual.

3.1 A brief introduction to Desert Rain

Perhaps the best way to understand Desert Rain is to follow a player's experience from start to finish.

The whole experience lasted for approximately forty minutes. It began in the physical world. Players bought tickets in advance and usually gathered in their groups of six at the venue, although in one location, they gathered in the city centre and were then bussed out to the venue – a disused warehouse on the outskirts of the city. A performer lead them into a bare physical antechamber (figure 4) where they were asked to remove their outer clothing, deposit mobile phones and similar possessions in a box under their seat. They donned a uniform, an anorak, and were briefed as to their mission. The briefing introduced the six targets by name and photograph, explained how to use the footpads, and stressed the time critical and cooperative nature of the mission. There was no opportunity for questions.

Each player was then led in turn by a performer to a fabric cubicle and was zipped inside (figure 5). There they stood on a footpad and put on a combination headphone/microphone headset. When all six players were in place,

the water was switched on. Each player was then facing their own personal rain-curtain – a large screen, roughly two meters tall by two and a half wide, composed of falling water, onto which an image of a virtual motel room was back projected (figure 6). Each footpad acted as a giant joystick; by shifting their weight on its surface the player could move their virtual viewpoint forwards and backwards and could rotate it clockwise and anti-clockwise. Behind the six curtains, unseen by the players, lurked two performers. The asymmetric nature of visibility through the rain curtain means that these performers could observe the players, without being observed in return.

The action now switched to the virtual world. The six players began their journey through this world isolated from one another, each in a separate virtual motel room where a virtual TV set played back a short TV clip of CNN's Gulf War news coverage (as a video texture). Eventually the player left their motel room, passed into an open desert landscape beyond, and headed towards the centre of the world.

As the six players drew closer, they found that they could hear one another through a live audio link, mixed in with an ambient soundtrack. They could also hear voices (those of the performers) advising them where to go if they were lost or what to do if they were experiencing difficulties. When they met, the players saw one another represented as avatars with text labels (e.g., “player 1”) and the texture mapped image of the relevant target for that player on the front. Each player eventually located the virtual doorway that was labeled with the name of their target. On crossing this doorway, they found that they were standing inside a rotating white virtual cylinder, facing a sign that said “wait here”.

The action now swapped back to the physical world. One of the performers who had been observing from behind the rain curtain, physically stepped through the curtain (figure 7), slowly approached the player on the footpad, gave them a plastic swipe-card and without speaking or otherwise acknowledging their presence, turned away and walked back through the curtain. Given that the players have been concentrating hard on the virtual world and that they were likely to have been feeling somewhat disorientated, this was usually experienced as a highly dramatic, even shocking, event by the players.

The action now swapped back to the virtual world. The players were encouraged to find the entrance to an underground bunker. Inside they found a maze of narrow corridors, similar in style to many contemporary computer games. Together they had to find the exit before their thirty minutes were up. Once found, the exit would only open if all of the players had first found their targets. If the team passed through the open exit within the allowed thirty minutes the performers appeared again to lead each player forward, passing through the rain curtain – the reward for success. If they failed, the water was switched off before they were led forward to the next stage.

For the final time, the action swapped back to the physical world. The players passed along a narrow corridor and climbed over an enclosed ramp covered in sand (figure 8) into a physical recreation of the virtual motel room – but one created by pasting wallpaper sized photographs onto the walls (figure 9).

Here they found a real television set. Using the swipe card given to them earlier, each player activated a video clip containing an interview with their target. The six clips offered different reflections on the Gulf War, revealed through interviews with the six targets: a soldier who served in the gulf war, driving a personnel carrier that collected Iraqi casualties; a soldier who was bedridden at the time of the war and watched it on TV; a peacemaker who helped establish a peace camp on the Iraqi-Saudi border in December 1990; a journalist who was in Baghdad on the night the air war started; an actor who played a soldier in a TV drama about the Gulf War; and an actor who was on holiday in Egypt at the time.

Finally, the participants changed back into their original clothing and emerged from the experience. Sometime later they found that a small bag of sand containing an estimated 100,000 grains – the estimated number of casualties in the Gulf War – had been left in their pocket.



Figure 4: The briefing room

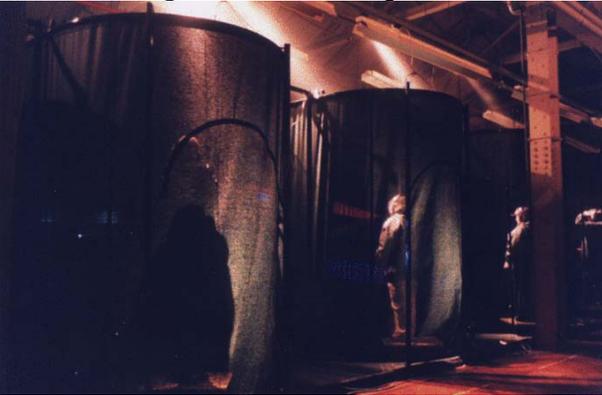


Figure 5: The players zipped into the fabric cubicles

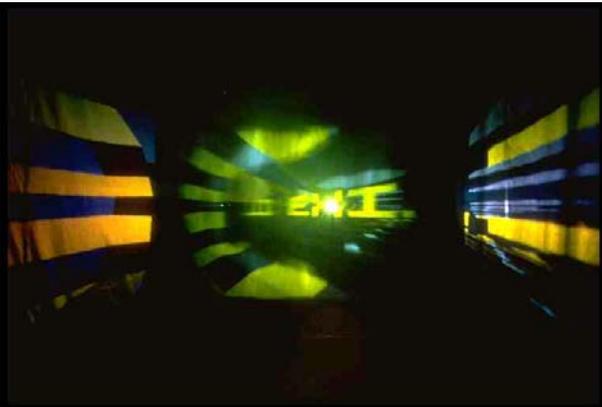


Figure 6: the virtual world as seen in the rain curtain



Figure 9: the final motel room



Figure 7: A performer crossing the rain curtain

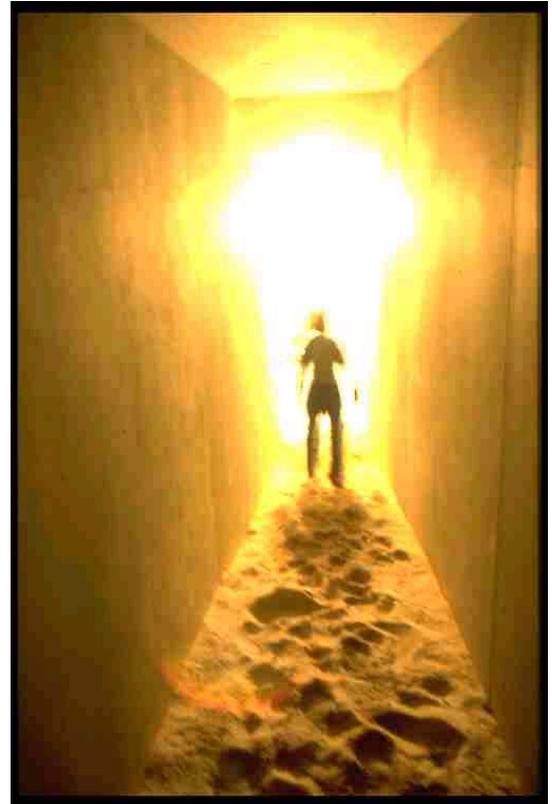


Figure 8: climbing the hill of sand

3.2 Behind the scenes of Desert Rain

Technically, Desert Rain was a practical exploration of using mixed reality technology to create a coherent and engaging public experience. The specific mixed reality technology in question was that of traversable interfaces. These establish the illusion that a physical space is joined to an adjacent virtual space and that participants physically pass from one to the other (appearing to dematerialise from physical space and rematerialise in virtual space or vice versa) (Koleva, 2000). Early prototypes of traversable interfaces employed walk-through projection surfaces to create the illusion of physically stepping into or out of the image of a virtual world. One of these prototypes was the rain curtain, a fine water spray into which images can be back projected. Desert Rain took the rain-curtain technology and used it to create a full-scale public performance. The rain curtain was chosen for its aesthetic qualities, both in terms of its striking visual image and sound, its asymmetric transparency (discussed below), and not least, due to the artistic association of projecting a virtual desert into a curtain of water.

Desert Rain was staged on an extensive physical and virtual set that was home to a complex array of technology. The virtual world was implemented in the MASSIVE-2 system. Figure 10 provides an overview of the physical set for Desert Rain (precise positions vary according to the physical constraints of the venue).

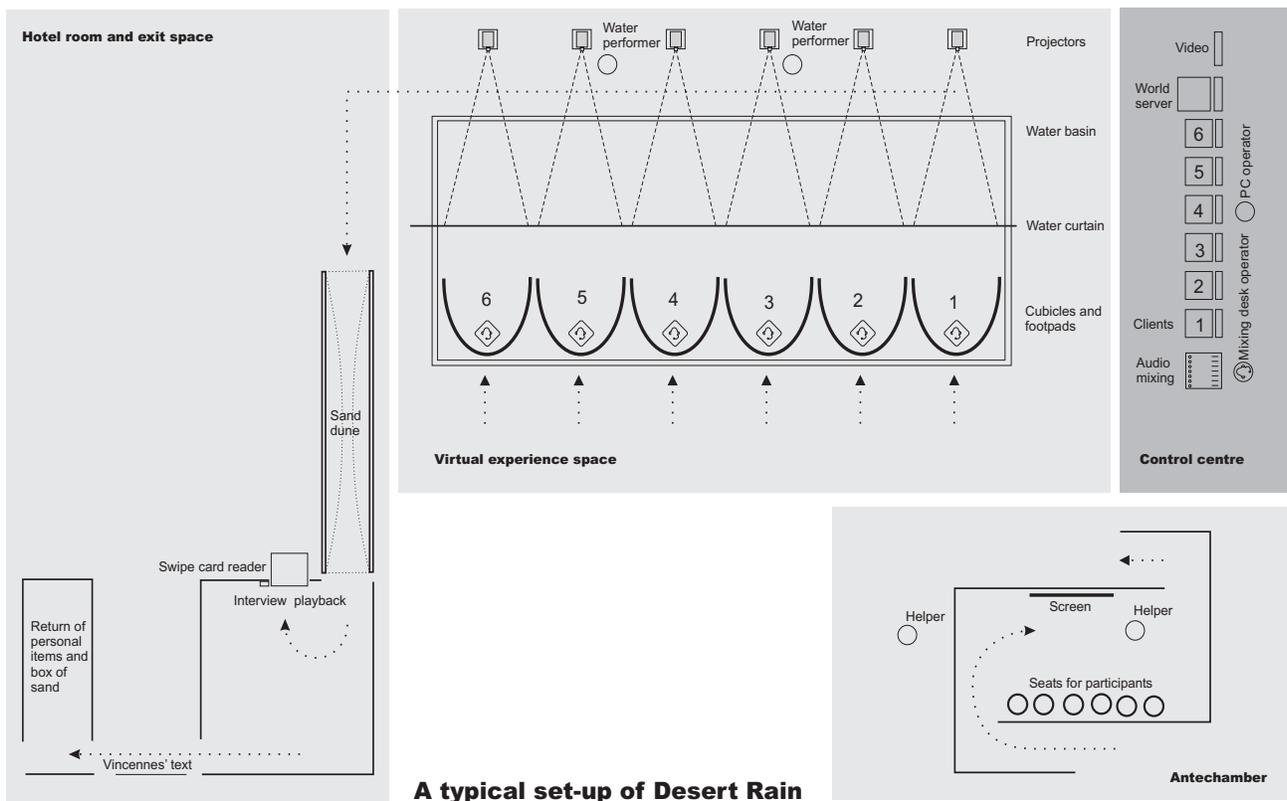


Figure 10: The set for Desert Rain

This physical set was divided into four main areas:

- The *antechamber* where the initial briefing took place (figure 4).
- The *virtual experience space* consisting of the six cubicles, each with an associated rain curtain, footpad and rear projector. Figure 5 shows this as seen from the players' side.

- The *hotel room and exit space*, consisting of the corridor of sand (figure 8) and the hotel room (figure 9).
- The *control centre* (figure 12). This housed the computers and related networking, audio and video technology for the virtual world.

Three aspects of this set-up are worth particular note as they provided the performers with different ways of monitoring and intervening in the participants' experiences.

- Six computer displays located in the control centre displayed the viewpoints of the six players in the virtual world (as projected onto the rain curtains). These monitors enabled the performers to follow all six players as they progressed through the virtual world. Using the arrow keys on the keyboards, the performers could also directly influence the movements of the players' avatars, nudging them in particular directions.
- Using an audio mixing desk that was also located in the control centre, the performers could listen in to the conversations of the players. They could also make additional direct audio connections between pairs of players no matter where they were in the virtual world (audio connections between players were established according to virtual proximity). Finally, they could speak directly and privately to particular combinations of players, giving them assistance without the others hearing.
- Performers could also position themselves in the corridor behind the six rain curtains (figure 11) from where they could surreptitiously monitor the players. This was made possible by the asymmetric nature of visibility through the rain curtain. Looking back through a rain curtain, a performer could observe the physical player standing on the footpad, and could also see a (reverse) image of the player's view of the virtual world projected onto their body and the surrounding fabric.



Figure 11: the corridor behind the projectors

The following table summarises the typical division of responsibility between the performers.

Performer 1 (antechamber)	Met and briefed the players in the antechamber
Performers 2 and 3 (behind the rain curtains)	Led players to cubicles and zipped them in. Observed players from behind the rain curtains. Crossed curtains at the correct moment and handed over swipe card. Led players to the corridor of sand.
Performer 4 (control centre)	Monitored players' progress through the virtual world and their conversations. Talked to players over audio.
Performer 5 (control centre)	Controlled and monitored the MASSIVE-2 software. Helped steer players' avatars using the keyboards.

These individuals communicated with one another by whispering when co-located in the control centre or in the area behind the six projectors. They also employed a simple signaling system using flashlights for communication between these two spaces – for example indicating how many players were present in a particular performance (not always the maximum possible six).

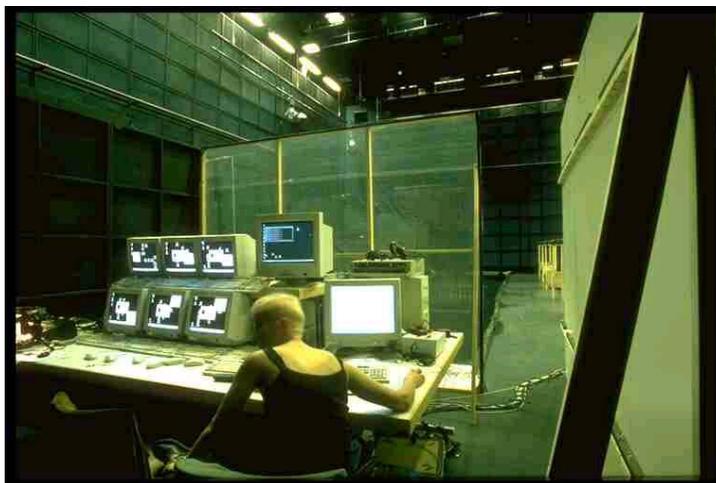


Figure 12: the control area

3.3 Desert Rain in practice: establishing and maintaining engagement

Desert Rain has toured extensively since 1999, having visited Nottingham, Karlsruhe, London, Bristol, Glasgow, Manchester, Middlesborough, Stockholm, Rotterdam, Prague and Sydney. Early performances generated considerable interest among art critics:

“... is possibly the most technologically ambitious art installation ever made” and “Sombre as its aims may be, Desert Rain is exhilarating to experience first hand.” [The Times (UK), May 10th 2000]

“... the experience does recreate some of the fear and disorientation that those on the ground during the Gulf War must have felt” and “part of a growing trend in performance and installation to blur the line between spectator and participant” [The Guardian (UK), May 18th 2000]

This paper takes Desert Rain as a rare opportunity to study how a professional theater group stages a touring real-time interactive experience. The particular design of Desert Rain, especially its physical scale, required a range of different methods of data collection. Following a preliminary investigation into Desert Rain while it was being prototyped at Karlsruhe in January 1999, we used a combination of direct and video-based observation, participant interviews and questionnaires to collect data at five further locations where Desert Rain was actually staged and performed. The following table summarises these various studies.

Location/time	Focus of study and methods
Karlsruhe, January 1999, 3 days	Preliminary study. Initial understanding of goals and technology. Observation and informal interviews with design team and visitors.
Nottingham, October 1999, 5 days	Premiere of finished work. Field observations and ten intensive interviews with players before, and seven after, the experience.
Karlsruhe, November 1999, 5 days	Observations and interviews as at Nottingham
London, May 2000, 5 days	Focus on players experiences. Questionnaires on entry and five intensive interviews on exit.
Bristol, June 2000, 5 days	Focus on social interaction in virtual world. Video recordings of two players’ viewpoints in the virtual world.
Glasgow, July 2000, 5 days	Focus on interaction amongst performers. Interviews with performers and observations of players.

Our study of Desert Rain has shed light on the practices by which performers orchestrate the players' experiences during all stages of the performance. In the following, we will see how Blast Theory developed a range of tacit working practices and procedures for engaging players with the content – setting their expectations and enabling the illusion of traversal – and for monitoring ongoing events and intervening if necessary, with minimal disruption to their engagement.

3.3.1 *Setting expectations*

Our observations suggest that, to a large extent, participants believed (or rather willingly suspended disbelief) when they participated in Desert Rain. The crossing of the rain-curtain by the performer created surprise and excitement and helped to enhance the participants' engagement. In order to establish this illusion of traversal, Blast Theory conducted carefully designed actions that led participants into and out of the different phases of the performance. At several points in the performance Blast Theory deliberately attempted to set the players' expectations, encouraging them to willingly suspend disbelief.

In the antechamber – participants were introduced to Desert Rain. As they entered the installation guided by a performer, participants were instructed to change their outer-clothing. In a command-like style they were briefed about their mission in the virtual world and engaged with the game. Before they were isolated from the rest of the group and guided one-by-one to their cubicles and zipped inside, they were advised to make use of the communication facilities in the virtual world. The importance of presenting a backstory prior to an interactive experience is well understood by the imagineers who design theme-park rides (Dodsworth, 1997; Kamberg, 1998). However, Blast Theory took this a step further. It can be argued that through a range of well prepared stages, participants were gradually desocialized like an inpatient in an asylum who “finds himself cleanly stripped of his many of his accustomed affirmations, satisfactions and defences” (Goffman, 1987), and were subjected to a set of discomfiting experiences. As the participants normally followed the performers' instructions quietly without showing signs of resistance towards this desocialisation procedure, it seems that the performers' carefully conducted actions did indeed set the players' expectations and enable the illusion of the game.

Crossing the curtain – the moment when a player finds her target and a performer crosses the rain-curtain was the climax of the dramatic performance. After the players had navigated the virtual world for some time they had become familiar with the game and were ready for new experiences to help to maintain their engagement. Blast Theory designed the players' discoveries of their targets in a particular way that created surprise and excitement and made possible that their engagement with the game was strengthened by the crossing of the curtain. However, the crossing is a moment that was fraught with danger for the performers. The timing of their movement with the players' actions in the virtual world was of crucial importance. When players saw the rotating drum for a long time they became disoriented and questioned its part in the performance. They sometimes assumed that the system had crashed and waited for a continuance of the program. If the performer walked too quickly through the rain-curtain the illusion of her emergence from the virtual world did not work. If the timing of the crossing was not right, it disengaged the players from the game, and they attempted to involve the performer in conversation.

Leaving the cubicles – the game ended when the players found the exit to the virtual world and left the cubicles. The performance, however, continued as the players were led to the motel room. This is perhaps the point in Desert Rain when it was most difficult to sustain engagement. It seems that players often assumed that with the end of the virtual game the performance came to an end as well. In early performances, they would meet their co-players in front of the cubicles where all of them would take off their anoraks, and then discuss with each other their experiences in the virtual world. Sometimes they would attempt to return the swipecards to the performers. Only when they were asked to climb the sandhill did the tension build up again. However, it then

collapsed again only a few seconds later after they have gathered in the motel room. Their excitement about the experiences in the virtual world appeared to take over and they began discussions with each other. They disengaged from Desert Rain prematurely.

In response to these observations, the performers altered their orchestration of this part of the performance. The players no longer removed their anoraks until after the final motel room. Furthermore, the performers carefully planned the order in which to take the players from their cubicles, so that the players spent the minimum possible time together before moving on up the sandhill, and also so that the performers were best positioned to shepherd them on. This shows the level of detail that has to be considered when planning and executing a performance. The interactions with the players are meticulously planned and repeatedly rehearsed, including dialogue, inflexions, gestures and speed of movement. Potential problems were identified in advance and responses are rehearsed, with a particular focus on how they could be woven into the experience.

3.3.2 Monitoring and Intervening

Our observations indicate that the performers largely managed to lead the participants into the installation and engage them with the game. However, to ensure that the players' engagement with Desert Rain was maintained throughout their journey, performers continually monitored events in the virtual as well as in the physical world. If the players' engagement with the game seemed to be endangered at any point performers had to hand a range of prepared actions through which they could intervene in events. The players were never really isolated in their cubicle as the design of the installation allowed the performers asymmetric access to the players, both in the virtual and physical worlds (through the computer monitors in the control area and the asymmetric nature of the rain curtain respectively). Monitoring and intervening were therefore largely accomplished without the players noticing it. We have identified three styles of intervention that were employed.

Dramatic interventions were conducted by the performer at the control-centre to advise players about actions in the virtual world. Occasionally, the performer also attempted to influence the player's movement on the footpad. These interventions were produced by means of the audio link to talk to a player and to advise her on which directions to take in the virtual world or on the use of the footpad. They could not be conducted without the players noticing them. However, performers used a specially designed, dramatic voice and almost always managed to embed the intervention within the game, thus avoiding the players from becoming distracted from their actions within the virtual world.

As with other recent digital media artworks (Büscher, 2000), Desert Rain was designed to promote collaboration among participants. Another opportunity for dramatic intervention was therefore encouraging the players to help each other, for example, suggesting that players who had already found their targets go back and help their team members who had not. Given that we found communication between the players to be an especially enjoyable and engaging aspect of Desert Rain, this provided an ideal way of intervening without breaking engagement, indeed possibly even enhancing it.

Invisible interventions were carried out from the control-centre by virtue of the arrow keys on the computers. They were very carefully conducted so that the players did not notice them. For example, when a performer observed that a player had been circling her target for some time without crossing it, she focused her observation on this player's movement in relation to her target. As the player came very close to her target the performer pressed an arrow key to push the player's avatar through the target and thus trigger the rotating drum. The performer made her decision about the exact moment when she pushed the arrow key with great care so that it was neatly timed with the player's movement in the virtual world. She moved the avatar only a tiny bit, thus making sure the push remained unnoticed for the player. Virtual interventions were closely timed with the

player's movement in the virtual world. They could be carried out without disengaging the player from the game. Indeed, they ensured that players did not get frustrated as they circled their target.

Face-to-Face interventions were carried out by a performer who directly approached the player's cubicle to give her practical advice on using the footpad. This form of intervention was very intrusive and always resulted in an interruption of the player's engagement in the game. Therefore, it was only employed on very rare occasions. A few players had been observed who despite having received advice through the audio-link, could still not use the footpad; in order to enable them to engage in the game a performer accessed their cubicle from behind to give them hands-on support. Face-to-face interventions were only employed when dramatic interventions failed and were normally preceded by communication between the control centre and a performer behind the rain-curtain.

This concludes our discussion of Desert Rain for the time being. In the following section we turn our attention to the orchestration of a final real-time interactive experience called Can You See Me Now?.

4. Can You See Me Now? – a mobile mixed reality game

Can You See Me Now? was a mobile mixed reality game in which up to twenty on-line players were chased across a map of the UK city of Sheffield by three performers who were running through its streets. In contrast to Avatar Farm and Desert Rain, the experience involved participants who were physically dispersed and some of whom were mobile. The broad aim of Can You See Me Now? was to explore the nature of possible relationships between on-line and mobile players, building on several previous experiences with wireless and mobile games such as Pirates! (Bjork, 2001), Bot-Fighters (Botfighters, 2002), Geo-Caching (Geocaching, 2002), AR2 Hockey and RV-Border Guards (Singletarry, 2000). Can You See me Now? was staged as a public event in December 2001 and, like Desert Rain, was a collaboration between the University of Nottingham and Blast Theory.

4.1 A brief overview of Can You See Me Now?

Central to Can You See Me Now? was a relationship between up to twenty on-line *players* (members of the public using the Internet) who were moving across a map of Sheffield, and three *runners* (members of Blast Theory) who were moving through the streets of Sheffield. The runners chased the players. The players avoided being 'viewed' (caught). Everyone, runners and players, saw the position of everyone else on a shared map. Players sent text messages to each other, which were also received by the runners. The runners talked to one another over a shared walkie-talkie radio channel that was also overheard by the players. The game took place over an area of Sheffield that was roughly half a mile square and that consisted of a mixture of open spaces and narrow streets lined with tall buildings.

A player's experience began at the Can You See Me Now? homepage (www.canyouseemenow.co.uk) where they entered a name for themselves in response to the enigmatic prompt "is there some one you haven't seen for a long time?" They then joined the game queue, and from there were eventually dropped into the map of Sheffield. They used the arrow keys on their keyboard to move around this map. They were unable to enter solid buildings and other restricted areas.

Figure 13 shows the player interface. A player was represented by a pair of icons on the map. A simple white icon showed their current position according to the client software running on their local computer, providing immediate feedback as to their movement. A blue icon showed their position according to the game server, and this trailed behind the white icon with a lag of a few seconds (due to the communication delay over the Internet and the time taken to process players' movements at the server). Other players were represented as blue icons. The runners were shown as orange icons.

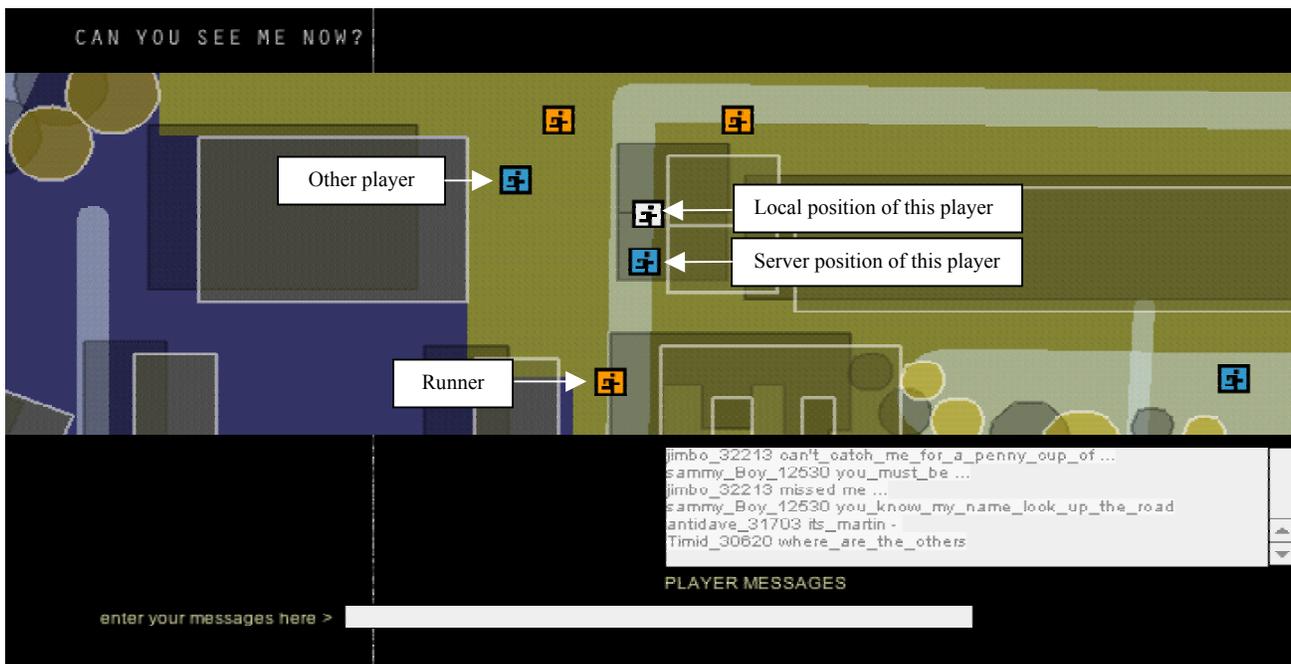


Figure 13: The player interface

The players continued to move and text until a runner got sufficiently close to them that they were viewed. At this point they were removed from the game and offered a chance to re-enter the queue.

The runners also saw the map of Sheffield showing their positions as well as the players' positions and text messages. Unlike the players, their map allowed them to zoom between a global view and a close-up local view centred on their current position. This interface was delivered on a Compaq iPAQ from a server in a nearby building over a 802.11b local area network. A GPS receiver plugged into the serial port of the iPAQ registered the runner's position as they moved through the streets and this was sent back to the server over the wireless network. The iPAQ and GPS receiver combination was attached to a wooden board that could be placed in a plastic bag to improve ruggedness, ease of carrying, and to provide some weatherproofing. The runners used walkie-talkies with earpieces and a head-mounted microphone. Finally, they carried digital cameras so that they could take pictures of the physical locations where the players were viewed. These pictures subsequently appeared on an archive web site. Figure 14 shows one of the runners kitted out (left) and the equipment that they carried (right).

4.2 Behind the scenes of Can You See Me Now?

Can You See Me Now? required position updates to be exchanged between players and runners; text messages to be exchanged between the players and also transmitted to the runners; and audio information to be transmitted from the runners' walkie-talkies to the players. This was achieved by two separate subsystems, one dealing with position updates and text messages and the other dealing with audio. These subsystems were spread over four locations: the streets of Sheffield, a temporary control room that was established in Sheffield, our laboratory back in Nottingham, and the Internet at large (connecting to the players' locations).

4.2.1 Position and text subsystem

Figure 15 shows the position and text subsystem. Players initially contacted a public HTTP server from which they downloaded their game web client (a Macromedia Shockwave program) and were then placed into a queue. A Fuselight multiuser server, hosted at Nottingham, managed admission to the game, ensuring that no more than twenty players could play at a time. Once admitted, a player's client contacted the main game server

(implemented in Macromedia Director) that was hosted in the control room in Sheffield. The client's game events (position updates and text messages) were sent to the server. In return, this sent back game events from the other players, as well as the positions of the runners.

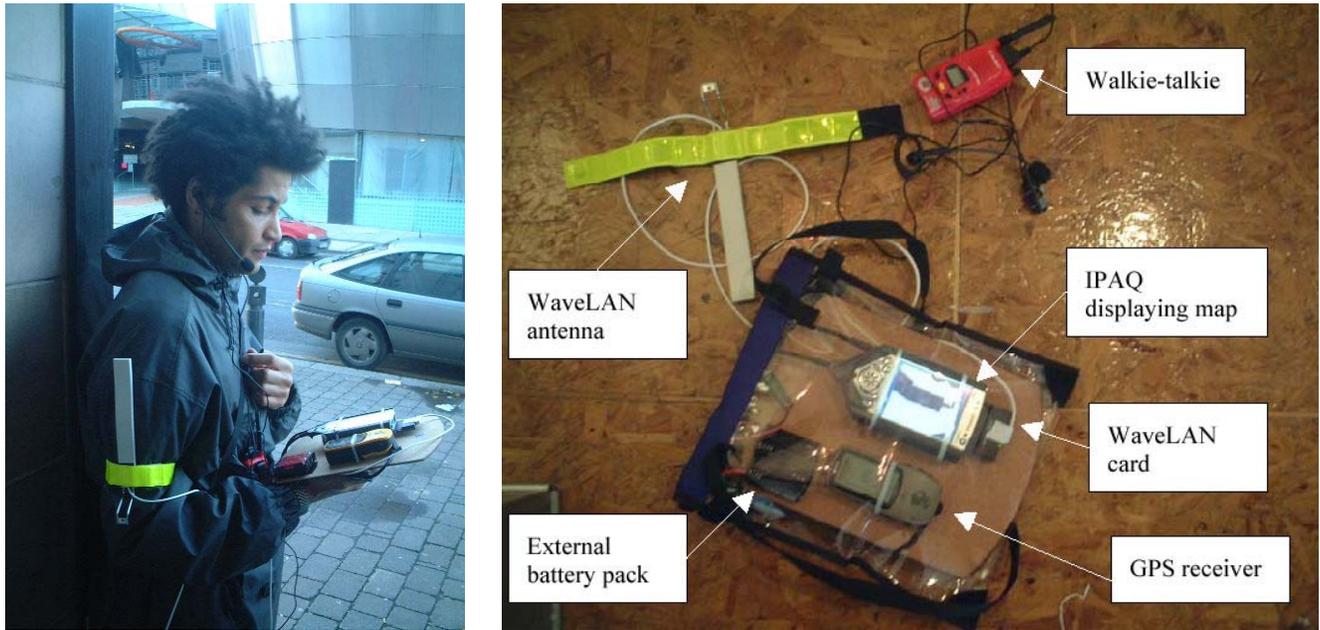


Figure 14: A runner (left) and their equipment (right)

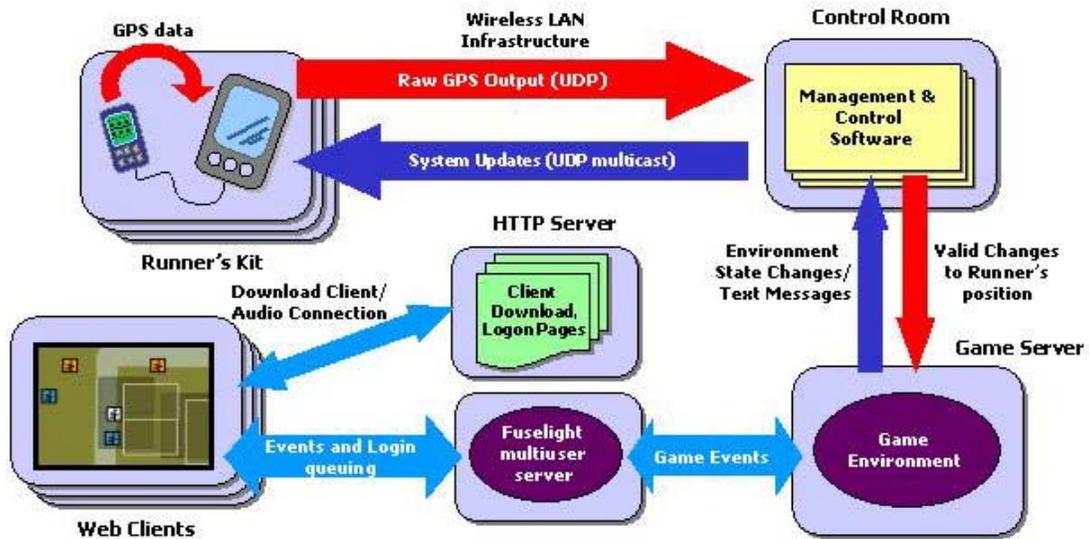


Figure 15: Position and text subsystem

Each runner on the streets transmitted position updates from their GPS receiver via their iPAQ to a proxy server that was running in the control room in Sheffield. These updates were unicast over the 802.11b network using UDP. The proxy server converted the GPS coordinates from latitude and longitude to metric units based upon a known reference point in Sheffield. It then transmitted them to the main game server (via TCP). In return, the

players' positions and text messages were transmitted from the game server to a second proxy server which then sent them on to the runners' iPAQs over the 802.11b network as (multicast) UDP messages.

4.2.2 Audio subsystem

Figure 16 summarises the audio subsystem. The runners communicated using radio walkie-talkies. An additional walkie-talkie in the control room also received this communication. This was wired into a local computer, enabling the audio to be encoded into a digital audio stream (using Sorenson Broadcaster). From here it was transmitted to a Darwin audio streamer that was hosted in Nottingham and then made available to the players over the Internet. A useful feature of this set-up was that the walkie-talkie in the control room could also be used to talk back to the runners (e.g., to give them guidance and instructions) without being overheard on the on the public audio stream

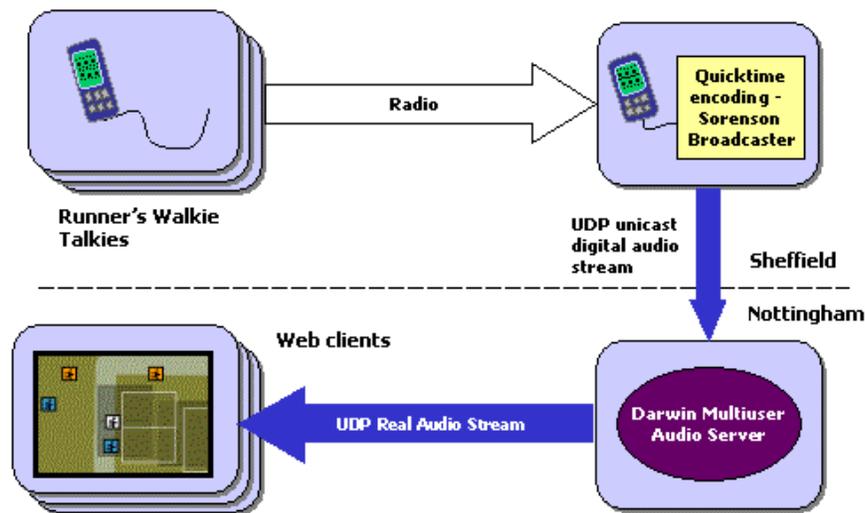


Figure 16: Audio subsystem

4.2.3 Wireless network

We invested considerable effort in establishing an 802.11b network with sufficient range. Two advance trips to Sheffield were undertaken in order to test out 802.11b and GPS coverage. These enabled us to establish a sense of the area within which the game would be playable (determined by a combination of physical accessibility, 802.11b signal strength and GPS accuracy).

Our final 802.11b network involved deploying a high-power omni-directional antenna mounted on an eight-meter mast on the roof of the building where the control room was located to give longer-range coverage. This was supplemented with a smaller lower power omni aerial to fill in coverage on the street immediately below the control room.

4.2.4 The control room

We established a control room in Sheffield to support the orchestration of Can You See Me Now? This was home to several monitoring tools:

- An application that monitored the signal quality of the 802.11b connection for each runner.

- An application that monitored the GPS data feeding back from each runner so as to provide an indication of tracking accuracy.
- A management interface that displayed all participants' positions, text messages, and supported management functions such as removing particular players.
- A laptop running a player's online interface so that staff in the control room could join in as a player to carry out live testing of the game.
- An application that tested connections to the main servers at Nottingham and to the 802.11b router the roof of the building.

4.3 Can You See me Now? in practice - orchestrating a distributed experience

Can You See Me Now? was live for 6.5 hours during the weekend of Friday 30th November and Saturday 1st December 2001. 214 players took part over the Internet, some of whom were attending a nearby conference where there were public terminals. 135 players were viewed, 76 logged off and 3 were never viewed. The best 'score' (time without being viewed) was 50 minutes. The worst was 13 seconds.

As for Avatar Farm and Desert Rain, we made ethnographic observations (utilising video and field notes) of the activities of the different participants, including runners, some players, and especially the behind-the-scenes production crew. We also gathered offline feedback from players via email and our web-site and held debriefing meetings with the project team in order to solicit the opinions of different participants. Finally, we instrumented the underlying system to log all movement updates and text messages that passed through the game server. Statistical and manual analysis of these logs revealed broad patterns of user activity to support or contradict other observations. Our analyses shed light on how runners, players, and crew worked together to play the game.

4.3.1 *Using tactics to control the pace of the game*

The runners' tactics changed significantly over the course of the event. The first session on the Friday saw them running frantically through the streets. Following a debriefing meeting on the Friday night, it was agreed that the players needed to be slowed down. However, this technical fix turned out to be impossible to implement in time, and so the Saturday sessions began with the players at the same speed. The runners now changed their tactics in several significant ways. First, they slowed down in order to lure the players closer, and then suddenly sprinted to catch them. Second, they learned to exploit areas of good GPS coverage where having accurate updates would make it easier to catch the players. Third, they collaborated more closely, essentially hunting as a pack, and cornering chosen victims. This deliberate tactical approach is perhaps best summarised by the following fragment of conversation that was observed in control room during a rest period:

- Steve: So your tactics: slow down, reel them in, and get them?
 Runner: If they're in a place that I know it's really hard to catch them, I walk around a little bit and wait till they're heading somewhere where I can catch them.
 Steve: Ambush!
 Runner: Yeah, ambush.
 Steve: What defines a good place to catch them?
 Runner: A big open space, with good GPS coverage, where you can get quick update because then every move you make is updated when you're heading towards them; because one of the problems is if you're running towards them and you're in a place where it slowly updates, you jump past them, and that's really frustrating. So you've got to worry about the GPS as much as catching them.

These accounts are backed up by analysis of system logs of the runners' movements. Plots of the normal distribution of runner speeds showed that the runners went from running at a relatively fixed speed (2 meters per

second) at the beginning of the event, to utilizing a broader distribution of speeds at the end. Also, as the game progressed the mean distance between runners decreased to around 75 meters. There were also many examples of players sending text messages to taunt the runners and goad them into action. This suggests why the runners' tactic of slowing down were so effective; essentially, there was little else to do in the game and so the players were naturally drawn into flirting with danger.

In summary, pace was a critical issue for playing *Can You See Me Now?*, particularly with regard to the runners being able to catch the players. What at first felt like a technical issue turned out to be a human issue, being addressed through the runners' choice of appropriate tactics, although this required familiarity with the technical performance of the system (i.e., knowledge of areas of good GPS coverage).

4.3.2 *Exploiting local knowledge*

The runners made extensive use of their local knowledge of Sheffield to coordinate their actions. Collaboration between runners was primarily achieved through the use of the walkie-talkies. The following sequence shows the work involved in coordinating runner-runner interaction.

Runner 1: (on walkie-talkie) I need a runner at the glowing mushroom. I need a runner at the glowing mushroom.
Runner 2: (on walkie-talkie) I'm thirty seconds away.
Runner 1: I need another runner to meet me at the glowing mushroom.
Runner 2: I'm ten seconds away.
Runner 1: Where are you?
Runner 2: I'm going round to your right.
Runner 1: Okay.

As the sequence makes clear, the use of the walkie-talkie in the coordination of the actions of the runners relied on, and was accomplished through, the use of local reference points. The runners were familiar with the topographical features of the built environment in which the game took place. Runners knew the location of structures that made up the built environment and were aware of the spatial relationship that buildings had to other structures (pavements, roads, walls, etc) together with the contours of the landscape (inclines, slopes, and hills). This meant that the runners shared common knowledge of the physical landscape, which was embodied in locally formulated names (e.g. 'the glowing mushrooms'). These names provided shared points of reference in the physical terrain that the runners oriented to and employed to coordinate their actions.

However, these named landmarks were not reflected in the digital domain (e.g., as labels on the shared map) and so, on the whole, the players did not share the same knowledge of the game environment as the runners. Consequently, a particular player might not be aware that he was being targeted, or how far off a particular runner was, or what direction the runner was approaching from, or where blind spots were in the game, and so might not take evasive action until it was too late, or alternatively might take evasive action when none was required.

One solution would be to label common landmarks and key locations on the map. This might be done in advance of the game. More interestingly however, the game might also allow participants to create and share their own annotations (after all, the label "glowing mushroom" does not appear on any conventional map of Sheffield).

Local knowledge also extends beyond labels, to include other features of the environment. For example, there was a hill in one part of the playing area. Repeatedly running up this hill was tiring for the runners (and hence would have been a good tactic for the players). Similarly, traffic on busy roads would have hindered the runners and so could have been exploited by the players had they been more aware of it. A few players did figure out such tactics. As one put it in an email sent after the event:

"I figured out pretty quickly what was uphill and downhill. I also figured out which was the main road to cross"

However, many players apparently did not, and might have benefited from techniques to enrich their local knowledge. Features such as hills, traffic and other obstacles might be represented on the map. They might also be

more subtly incorporated into the gameplay. For example, players might be slowed down when moving uphill. Finally, video cameras might be used to provide the players with live views of the city streets, an increasingly feasible idea given the growth of traffic and other cameras in public spaces and their availability over the web. Feedback from one of the players after the game clearly identified the potential of such views. This player revealed that they had been playing from a machine in a nearby building next to a window that enabled them to look out onto the physical game area. They reported thoroughly enjoying moving across the online map so that could watch a runner physically chasing past a few seconds later.

4.3.3 *Performing orchestration*

The runners used the walkie-talkie system to coordinate their actions. Much of this orchestration discussion took place over the public channel that was streamed to the players. In other words, the process of orchestration was a deliberately performed aspect of the game, just as it was with the off-face interventions of Desert Rain, although in this case the players were given the sense that they were eavesdropping rather than being directly addressed. As noted above, messages from the crew in the control room to the runners were not publicly available (i.e., were excluded from the performance) as were messages on the second private walkie-talkie channel. We observed several ways in which the deliberate performance of orchestration contributed to the game.

First, it provided a further mechanism for conveying local knowledge. The runners would often make references in their talk to local features including landmarks (as already mentioned), traffic conditions (e.g., “I’m waiting for a Green Man”, meaning I’m waiting at a road crossing), being tired, or the state of the technology (e.g., references to batteries being low and GPS accuracy). Although many players may have failed to pick up on these cues (possibly due to some problems receiving the audio stream over the Internet), providing descriptions of conditions on the streets through talk offers a potentially powerful way of better connecting the physical and on-line environments.

Second, hearing their name mentioned by the runners could be an exciting moment for a player. As our previous player put it in the same email:

“I only managed to get on to the map once for about 15 minutes. I can’t remember the name I used, but it was pretty unnerving first hearing my name said”

In effect, the runners could suddenly heighten the tension for a player by indicating that they were being hunted,

Third, the runners would use the audio channel to make ritualized and formal announcements of successful viewings. For example:

Runner 3: I’ve seen Max. Runner 3 has seen Max ...in the stairwell ... next to the spiral staircase ... next to the L shape.

Once again, these announcements were superficially addressed to the other runners, but were actually also addressed to the players as a way of sustaining tension (it would be clear that a colleague had been viewed and of course, that someone else was now going to become a target).

As with Desert Rain, it was important to plan such performances. The runners’ discussions before the event considered how to use the audio channel, especially defining appropriate language (e.g., using terms such as ‘seen’ and ‘viewed’ as opposed to ‘caught’ and ‘killed’). In practice, in the heat of the moment (and possibly because of the relatively unrehearsed nature of the experience), it was not always easy to maintain this discipline and there were some examples of more traditional gaming terms making their way onto the public audio channel.

4.3.4 *Joining and leaving the game*

Joining and leaving were key points in the game. Joining involved being admitted to the game (as discussed previously) and also familiarization with its workings and with the other players.

On joining the game, players were dropped directly into main the game area at one of ten randomly chosen locations. Beyond some basic instructions on the initial web page, they had received no training and had been given no opportunity to practice before this point. Engagement with the game thus relied upon players' familiarity with and expectations about the technology (e.g. that one may move around by using the arrows keys, that the avatar in my view is mine, that one may communicate with others via text messages, etc.). This working knowledge was often not sufficient for establishing a sense of the game's workings. As a result, players would often instruct each other over the text channel as to how to play.

The text channel also supported the coordination of collaboration between players, where players would explore the game-play environment and interact with the runners together. Examples included trying to meet at a common location, reporting when they were being chased or were about to be caught, and exchanging encouragement and tips. A key aspect of this coordination was establishing the identity of other players, an activity that took some work given their identical iconic representations.

Again, we see that encouraging players to help one another can be a useful orchestration strategy. This might have been further encouraged by initially dropping the players into a safe zone away from the main game play. Here they could become familiar with the game, orientate themselves and form relationships with other players before progressing to the main play area to deal with the runners.

Another possibility that was originally proposed (but not implemented due to lack of development time) was to allow players waiting in the queue to watch the game being played by others in advance of their turn. Such a public broadcast might support training and familiarity, although at the cost of reducing surprise.

Exiting the game also raised issues. As noted above, the runners had established a common ritual for when a player was caught: they would take a digital photograph of the site of capture (to appear on the web archive after the event) and would then report the viewing over the audio channel. The game server would also generate a text notification that was sent to all participants. It was also common for players to use text messages to notify others of imminent capture or to support others in trouble as in the following exchange between Nanny and Scott:

```
nanny    they've_got_me_in_a_pincer_movement
scott    nanny_good_luck
nanny    am_doomed
scott    no_way_to_fight_them
scott    bye_nanny
scott    come_back_reincarnated
```

There were even some reports of players coming back into the game just to say goodbye to others. However, there was no ritual that allowed players to mark the moment of being viewed. Perhaps the game could have been extended to allow players to mark the moment of their passing, perhaps with a final text message that would have been sent to other players and posted to the archive web site.

4.3.5 *Work in and around the control room*

Control room staff spent a great deal of their time working together to monitor the game as shown by the following example of shoptalk:

```
Steve:    (playing the game) How many players have we got?
Martin:   (looking at the global view) 21 players so far.
Steve:    Altogether or at the moment?
Martin:   You've got 21.
Steve:    Is it just two at the moment?
Martin:   What, players or runners?
Steve:    Runners.
Martin:   All three are out there. Dave's joined.
```

Steve: Just two players?
Martin: Two players.

This fragment illustrates how the crew had to work to establish a sense of the overall state of the game. What was being asked for when querying the population and its make up (of runners and players) was not so much a head count, but whether or not the game was working properly? Being able to see, via the global monitor, that players were engaged in the game told the controllers that the online system was working. Similarly, being able to see that a number of runners were actively engaged in the game told the controllers that the runners' gear was working. And, taken together, the global view on the game told the controllers that there were no prima facie technical problems. However, our monitoring tools were not ideal for this:

- the GPS and 802.11b monitors provided very low level information that was difficult to read at a glance and that required expert interpretation.
- the information was spread across several monitors in the control room.
- some vital information was not available. Dead batteries turned out to be a particular problem for Can You See me Now? A combination of the power management system and battery status reporting on the iPAQs made it difficult to predict when the runners' devices would fail and, as a result, they tended to fail mid-game. Orchestration would have benefited from accurate telemetry data concerning battery status.
- another class of missing information concerned the status of the players. For example, there was no visible indication of the length of the queue and no easy way of detecting problems with specific players (such as the failure to receive the audio stream).
- monitoring information was associated with a particular device (i.e., iPAQ). However, a scheduled change over of runners (the runners worked in overlapping shifts with three on and one off at any time) or the complete failure of a device (necessitating the introduction of a spare) meant that a runner would swap devices. It proved difficult to keep track of which runner was using which device (a problem when talking to the runners over the walkie-talkies).

These observations imply that more sophisticated monitoring interfaces are required. These should make it easy to get a quick overall sense of the status of the game or of particular participants while also supporting drilling down to obtain more complete detailed information on a particular participant.

Orchestration also spilled out of the control room and onto the streets. Coordination between controllers and runners over such practical matters as getting new batteries and establishing whether or not this or that was the 'problem', was facilitated via the mechanisms described previously: the walkie-talkies, the game overview, the GPS monitor and the 802.11b monitor. However, resolving these problems required us to deploy a full time member of the production crew outside of the control room, at street level, so that they could directly service the runners (e.g., changing batteries or trouble shooting with the iPAQs and GPS receivers). In turn, this raised the issue of how this person monitored the game (what kind of interface was available to them) and how they communicated with the runners and the other crew in the control room. A further problem was that it still took the runners approximately five minutes to reach this person whenever they needed help – a major disruption to the game. Perhaps this crew-member should themselves have been mobile?

5. Discussion: orchestration issues, techniques and challenges

This section reflects more broadly on our three experiences and draws out a framework of issues and available techniques to guide experience designers.

Avatar Farm, Desert Rain, and Can You See Me Now? are broadly representative of a range of participatory experiences. They include aspects of on-line gaming, role-play in virtual worlds, mobile gaming, performances, and artistic installations. They also differ in interesting ways:

- Physical-virtual mix – the content in Avatar Farm was primarily virtual, Desert Rain and Can You See me Now? mixed physical and virtual content.
- Co-location of participants – participants in Avatar Farm and Desert Rain were co-located (although sometimes segregated by boundaries), whereas in Can You See me Now? they were dispersed.
- Temporal structures and rhythms – each instance of Desert Rain involved set participants and had a fixed duration; Avatar Farm also involved set participants, but spanned four distinct episodes, each of flexible duration; Can You See me Now had a more fluid structure, with participants joining and leaving at different times over a longer time period.
- The role of external viewers – Avatar Farm explicitly addressed external viewers (via a live webcast and 3D recordings), whereas Desert Rain and Can You See Me Now? addressed direct participants only.

Although their approaches to orchestration differed, a comparison of the three identifies a common set of orchestration issues. In chronological order, from starting to finishing an experience, these are:

- Admission - when and how do participants join an experience?
- Training and familiarization – how do they learn how to participate?
- Establishing engagement – how do they engage with content and willingly suspend disbelief?
- Avoiding fractures in engagement – how do fractures happen and how can they be reduced or mitigated?
- Monitoring – how do performers and crew know what is happening in the experience?
- Intervening – what techniques are available to allow performers and crew to intervene in events?
- Managing pace and timing – how to sustain an appropriate pace?
- Coordinating – how do performers and crew work together to plan and execute what is to be done?
- Closing an experience – how does the experience end and what follow-through is there?

The following table summarises how these issues were reflected in our three experiences.

Orchestration issue	Avatar Farm	Desert Rain	Can You See Me Now?
Admission	Participants pre-selected by producers. Admission by invitation only. Limited to four players at a time. No joining and leaving during.	Players self-selecting. Traditional booking system with fixed time-slots. Six players at a time. No joining or leaving during.	Players self-selecting. Limited to twenty at a time. Admission by web interface with queuing system. Players can join and leave during.
Training and familiarization	Briefings and try-out sessions beforehand. Helpers on hand and mutual support among players. Carry over from previous experience in an on-line community.	Short rehearsed briefing beforehand, but no opportunity to try out and no obvious helper on hand. Players were encouraged to help one another.	Short instructions on website, but no other briefing. Players physically isolated. Mutual training through text chat. Planned to support watching while queuing (not realized).
Establishing engagement	Interest carried over from previous experience. Introduction and briefing meetings.	Carefully designed entrance space. Military-style formal briefing (including changing clothing) to create suspense.	Introductory web page with enigmatic question. Mention players' names over audio channel to create suspense.
Avoiding fractures in engagement	Participants wove fictional accounts of problems and delays into their talk. Some counter examples of fractures where talk directly	Early performances suffered from a key fracture as players regrouped. Performers changed approach to collecting	No particular techniques used. Players left to fend for themselves.

	referred to technology.	players to mitigate.	
Monitoring	Physical monitoring through co-location. Virtual camera interface used by director, assistant and stagehands to track players. Invisibility allowed stagehands to follow players without being seen.	Semi-transparent fabric enclosures and asymmetric nature of rain curtain supported unobtrusive monitoring of physical activity. Six displays showed players' virtual viewpoints.	Interfaces in the control room to monitor game state, GPS data and 802.11b signal strength. Problems with level of detail and integration of these. Inability to physically monitor runners and players.
Intervening	Stagehands invisibly manipulate virtual objects and avatars. Performers 'manipulate' players through talk. Physical intervention to deal with technical problems.	Performers intervene through talk and also invisibly directly manipulate players' avatars. Physical intervention a last resort.	No way of invisibly manipulating players or even engaging in one-to-one talk. Players helping one another through text chat..
Coordinating	Definition of specialised roles (director, assistant, stagehand, performer etc). Audio talkback system from director to crew and physical running of messages.	Coordination between two performers to time entrance through the rain curtain. Back-stage whispers. Use of flashlights as signaling devices behind-the-scenes.	Talk and running messages within and to/from the control room. Use of private walkie-talkie channel between runners and control room.
Managing pace and timing	Approach of improvising magic – extended incantations with responses from invisible stagehands – allowed open-ended timing of responses.	Performers use audio channel to keep the players under time pressure and invisible interventions to speed them up or slow them down as necessary.	Pace largely determined by the runners' tactics (e.g., waiting for players to come to them) as much as by the technology.
Closing an experience	No special mechanisms used. Endings of each chapter were sudden and abrupt.	Carefully designed exit from the experience. Small box of sand left in each player's pocket as a reminder.	Ritualised announcement and digital photo from runners when players viewed. Extensive archive website after the event.

We now revisit each issue in turn, breaking it down into more detailed design considerations and proposing guidelines for supporting tools and techniques.

5.1.1 Admission

Admission is concerned with how participants join an experience.

First, how are they selected? Options range from by the producers (Avatar Farm) through to self-selection (Desert Rain and Can You See Me Now?). An intermediate option is selection as a result of success in previous or related experiences, such as in trials or previous rounds of a game (we see this in Avatar Farm's selection according to participants commitment to a previous on-line community).

When do participants join? They might only be able to join at the start of each experience (Avatar Farm and Desert Rain) or might be able to join at any point (Can You See me Now?). Alternatively, there might be several defined key moments at which new participants can join.

Where do participants join? There may be carefully chosen entry points into the experience. Portals between worlds in Avatar Farm dropped participants into the center of the world. Players in Can You See Me Now? were also dropped into the thick of the action. With hindsight, these entry points might have been more carefully chosen. Important factors to consider include: the transitional experience for the new participant (what is their initial view of the action? do they get an informative overview of a new location when they first arrive? is there

time for them to get their bearings?), and the effect on other participants (will they appear in the middle of an ongoing scene, potentially disrupting it?).

Are the numbers of participants limited? If so, how is this achieved? One option for on-line events is a queuing system (Can You See me Now?). In this case participants need to be aware of the operation of the system and the current queue length and/or waiting time. If not, designers need to think carefully about the implications of potentially unlimited numbers of participants on the experience and on technical resources.

Do participants join as groups or as individuals? Some experiences (Avatar farm and Desert Rain) admit groups of participants. Designers need to consider where these are formed and what happens if there aren't enough participants to make up a group (can the experience work with a smaller group or do performers step in?). Also, are members of the group expected to already be familiar with one another and if not, is there an opportunity for bonding beforehand?

What experience do participants have before the main event (especially where queuing is involved)? They might see previews or reviews of previous experiences. They might meet other participants beforehand in an on-line 'green room' (an approach that is already used in on-line gaming). They might also watch the live experience while waiting. Alternatively, it may be important that the main event is a surprise to them in which case they will need to be carefully isolated beforehand

What infrastructure do participants require (hardware and software) and who is responsible for providing, installing, configuring and testing this? If they provide their own infrastructure, is any technical support available?

How do participants find out about the experience in the first place and how do they know when to join? Options here include websites, email and advertising. For time critical events, a notification system can count down the time to their entry and confirm that they are still available and fully prepared.

What is the fallback strategy if a participant will not or cannot join? Perhaps other queuing participants can be dropped in at short notice or recently finished participants can be invited back again. Alternatively, a performer may have to take their place.

5.1.2 Training and familiarisation

Training and familiarisation are concerned with how participants learn about the technology (e.g., how to use the interface) and the content (e.g., any necessary backstory and about their own and others' expected roles). Again, there are many possible approaches. Our three experiences demonstrate: transfer from previous experience (Avatar Farm), briefings (Avatar Farm and Desert Rain), on-line instructions (Can You See Me Now?), the physical presence of helpers (Avatar Farm and Desert Rain), and mutual training among participants (all three to some extent). Other strategies are learning by watching (e.g., while queuing) and providing dedicated practice areas (e.g., in an on-line green room).

5.1.3 Establishing engagement

A key aspect of the transition into the experience proper is engaging the participant with the content. The aim here is to generate a combination of excitement, suspense and tension and to focus their attention away from the technology (the focus of training and familiarisation) and onto content. More generally, the producers need to mark a point at which participants recognise that they now taking part in a performance and that it is appropriate and safe to willingly suspend disbelief.

Several techniques are available to support this transition. These include purpose designed entrance spaces (Desert Rain) as well as changes in lighting and sound that are commonly used in conventional theatres. Entrance spaces can also be virtual. For example, some presence experiments in virtual environments have adopted the

technique of first introducing participants to a virtual recreation of the physical environment in which the experiment is taking place, before allowing them to step into another less familiar virtual environment (Slater, 1998). Another possibility is to use a traversable interface so that participants physically step through a public external image of a virtual environment in order to access more immersive equipment (Koleva, 2000).

Performances and ritual are also an important part of establishing engagement. This is most clearly seen in Desert Rain where a carefully rehearsed military style briefing, complete with changes of costume, is used to unsettle participants and to convey the seriousness of the moment.

Finally, it is important to recognize techniques for suddenly increasing tension or engagement during the experience. The most notable example here is mentioning an individual players' name over the public audio channel in Can You See Me Now?

5.1.4 Avoiding fractures in engagement

The experience of Desert Rain suggests that there are key moments in an experience when engagement is particularly susceptible to fracture. These include sudden regroupings of participants following periods of separation (in which case there is a natural tendency to chat), movements between scenes and locations, and the completion of particular sub-objectives. Orchestration needs to minimise the risk of fracture (e.g., by collecting and moving participants in a particular way) or if unavoidable, must recognize that participants will have to be re-engaged before progressing further (using the techniques listed previously).

Our experiences also suggest that talk can play an important role in causing and avoiding fractures. References to technical terms in talk (as we saw in Avatar Farm) immediately focus attention back on the technology. These must be avoided through careful rehearsal and also through designers providing a common language within the content to refer to technical effects (e.g., it should be easy for participants to refer to the effects of actions without needing to discuss particular key presses and mouse clicks). Designers must make sure that participants encounter this language as part of familiarisation. On the other hand, Avatar Farm also showed how players and performers would weave accounts of technical problems into the content itself (e.g., explaining away technical problems in terms of magical effects). Although often clearly consciously ironic, this still serves to repair fractures in engagement and so is a strategy that performers in particular should be aware of and encourage.

5.1.5 Monitoring physical and virtual activity and system state

Being able to monitor the activities of different participants is a prerequisite to orchestrating them. Our experiences suggest that designers need to consider three broad aspects of monitoring.

First, it is necessary to monitor participants' virtual activities, (e.g., to track the actions and talk of their avatars in virtual environments). Where participants are local and their numbers are small, this can be achieved by attaching secondary displays to the computers that host their clients (Desert Rain). However, when they are physically dispersed or their numbers are greater, it becomes necessary to use separate virtual cameras that can lock onto and follow different participants and that can also offer a range of different viewpoints from close ups of detailed activity to broad overviews of an entire scene (Avatar Farm). It may also be useful to combine these with map interfaces that show the locations of participants in a virtual world (Can You See Me Now?).

Second, it is necessary to monitor participants' physical activities, i.e., to observe them using the technology and to be aware of the impact of other activities within their local physical environments. This is easier when there are local. However, even then, techniques such as semi-transparent partitions and asymmetric screens (Desert Rain) can facilitate more detailed and yet unobtrusive monitoring. Where they are dispersed, technologies such as video, audio, measures of keyboard activity, and even physiological data, may need to be used.

Third, it is necessary to monitor the performance of the underlying technology such as the state of software components (clients, servers, proxies etc), machines (CPU and memory usage), network (connections and traffic), and tracking (accuracy).

5.1.6 *Intervening to deal with problems*

Problems will arise, maybe as a result of technical failure or perhaps due to participants' lack of familiarity or skill. We identify three broad strategies for intervening to deal with them:

- Off-face interventions involve carefully embedded instructions within the performance, usually as spoken by performers. These need careful design and rehearsal to be effective.
- Invisible interventions involve manipulations in the virtual world such as steering participants' avatars without them knowing (Desert Rain) or manipulating objects while invisible (Avatar Farm).
- Face-to-face interventions are usually the last resort. Even then, they might be embedded into the experience to some extent. If this is not possible, designers should consider whether to isolate them from other participants, as there is no point fracturing someone else's engagement when they are not suffering a problem. The reverse strategy is to get participants to help one another.

At this point, we briefly raise the additional issue of governance. In some situations, especially those involving regulated public broadcasts, producers will be contractually bound to ensure that content is decent and legal. This may require relatively drastic mechanisms for intervening, such as the ability to quickly gag, constrain or remove participants. Governance may also utilise mechanisms such as delay-lines that impose a fixed delay between live action and its broadcast.

5.1.7 *Coordinating to achieve complex interactions*

Orchestrating complex interactions involving several participants and objects may require a team of performers and crew, in which case designers need to consider how they coordinate their activities.

It may be useful to identify different roles within the production team. In Avatar Farm, we saw the roles of stagehand (responsible for detailed object manipulations), performer (responsible for dealing with the players in the virtual world), the director and his assistant (responsible for overseeing the action, determining the overall plan of events and passing out instructions to the other roles), a floor manager (responsible for the management of the physical environment), actors' helpers (responsible for overseeing the use of immersive VR equipment), players' helpers (responsible for looking after the players in the physical environment) and a technical manager (responsible for overseeing the software).

Such teams will require dedicated communication facilities. These include private audio talkback systems ranging from open but private channels (e.g., the use of a separate walkie-talkie channel in Can You See me Now?) through to more flexible solutions that route audio between different combinations of roles (such as the software based audio switching system in Avatar Farm that allowed the director to route audio to individuals, selected groups or to broadcast to all). Silent communication mechanisms are also useful (e.g., the use of flashlights in Desert Rain). Finally, we would recommend the use of text chat tools as these are silent and can provide a useful fallback in distributed experiences when other media such as audio are problematic (e.g., due to network congestion).

The use of private backchannels is one example of making the activities of the crew invisible to players and viewers. Avatar Farm extended this approach by exploiting a potentially powerful feature of virtual environments, the ability to make selected participants invisible to others. However, Avatar Farm also showed how this can introduce further problems as a result of crew being invisible to one another. We return to this issue later when discussing the idea of a 'virtual backstage'.

5.1.8 *Managing pace and timing*

Sustaining engagement and tension involves setting the right pace of events. In *Desert Rain* it is important that the players' feel that they can achieve their goal and yet also feel under the pressure of time. In *Can You See Me Now?* the game has to be playable, i.e., it has to be possible for runners to catch players. Orchestration has to address several aspects of timing including speeding up slow players, slowing down fast players and getting a group of independent players to a common location at the right time. It may also be necessary to generate enough time to be able to marshal resources to carry out a requested interaction.

A key technique for managing pace and timing is to avoid direct interaction in favour of indirect indirection. In *Avatar Farm*, players request that interactions occur by undertaking potentially long-winded incantations rather than directly manipulating objects themselves. This buys enough time for performers and crew to coordinate an appropriate response. A further observation is that managing pace is as much a human as a technical issue. In *Can You See Me Now?* it transpired that the runners could largely control the pace of the game through their tactics.

5.1.9 *Closing the experience*

Our final issue concerns the closing of the experience. How do participants complete the experience? How do they know that the experience is over? Are they stimulated to reflect on it, discuss it or return to it at a later time? We see several different approaches in our three examples. *Desert Rain* introduced a carefully designed exit space. *Can You See Me Now?* relied on rituals, with performers reporting the viewings of participants over the public audio channel in a formal and stylized manner. Both *Desert Rain* and *Can You See Me Now?* introduced momentos of the experience to encourage reflection (a box of sand and photos respectively). It is interesting to consider the role of participants in creating such momentos (e.g., of players in *Can You See Me Now?* leaving their own final words for others to see). Finally, there is the technique of creating post-produced accounts of an experience, for example videos or archive web sites.

5.2 **Future challenges for orchestration**

We see two major trends for the future of participatory experiences, both of which raise significant challenges for orchestration.

- Scale in terms of increasing numbers of active participants. Fuelled by the success of very large-scale persistent role-playing experiences such as *Everquest* [X] and *Ultima Online* [X], the games industry is beginning to focus on massively multiplayer gaming in which hundreds or thousands of players are simultaneously active in a persistent game environment. These players will also be physically dispersed, accessing the game from homes and offices over the Internet.
- Mobility in terms of participants playing games using handheld or wearable devices that connect over wireless networks. There is already significant interest in gaming using mobile phones and researchers are increasingly focusing on ubiquitous, wireless and augmented reality gaming.

There is also the strong possibility that future experiences will combine these two trends to involve very large numbers of on-line and mobile participants. We identify three challenges for orchestrating these kinds of future experiences.

5.2.1 *Where is backstage?*

The idea of separating front-of-house from back stage is central to many (although certainly not to all) forms of theatre. Perhaps the best-known examples are theatres in which a seated audience faces a stage that is framed by proscenium arch. Like such theatres, the orchestration of *Desert Rain* and *Avatar Farm* relied on physical backstage areas. Their participants were physically co-located and their spaces included segregated backstage

spaces from which crew and offstage performers could monitor activity and intervene (even if the arrangement was not the same as for a proscenium stage). In the case of *Desert Rain*, there were also well-defined entrances to front of house – the rain curtains – behind which performers could wait and monitor the players in order to time their entrances.

As scale and mobility increase, it will become increasingly difficult to define a physical backstage. The spaces of both online and mobile players will be separated by large distances and will be physically inaccessible to performers and crew alike. Some other notion of backstage will be required. The solution may be to create a virtual backstage through parts of the virtual world that are hidden from but connected to the physical and virtual front of house. This might exploit several existing mechanisms in CVEs.

A virtual backstage might be defined as separate regions of the virtual world. This could exploit CVE mechanisms such as ‘locales’ (Barrus, 1996; Greenhalgh, Purbrick and Snowdon, 2000) and ‘third-party objects’ (Greenhalgh and Benford, 1999) that support one-way visibility between different virtual regions. It could also exploit spatial access control mechanisms that constrain which participants can pass through a virtual boundary (Bullock, 1997).

Alternatively, virtual backstage might be defined using the technique of invisibility, as was the case with *Avatar Farm*. In this case, members of the production crew and offstage (i.e., invisible) performers could move among the players without being seen. However, we also saw that simple treatments of invisibility (such as where individual avatars are ‘switched off’) can cause problems where crew members also cannot see one another and so cannot coordinate their actions. The solution may lie in more general techniques for creating subjective views in CVEs (Smith, 1997) such as the mechanism of ‘aspects’ (Greenhalgh, Purbrick and Snowdon, 2000) where participants subscribe to different layers of information in a virtual world. Crew and performers would subscribe to the front-of-house and backstage layers (the latter containing representations of themselves and other orchestration information), whereas regular players would only subscribe to the former.

An alternative approach is to make the production crew visible and to rely on players’ and viewers’ suspension of disbelief to hide them. This technique is used in conventional theatre (for example, where stagehands can sometimes be seen on stage) and in some forms of puppetry, where puppeteers are visible to the audience.

Finally, mobile experiences may involve creating multiple and/or mobile back-stage areas throughout the physical world. This could involve orchestration from vehicles or possibly from individuals (who might be disguised so as to act covertly) intervening through mobile orchestration interfaces.

5.2.2 *Monitoring activity*

It will become increasingly difficult to monitor experiences as their scale and mobility increase.

A lack of physical co-location will increase the difficulty of monitoring players’ physical activities. One potential solution is to make use of video monitoring, either from players’ own webcams that they set up (also providing opportunities for player feedback shots where experiences are being broadcast to a viewing audience) or from public surveillance cameras that are increasingly being made accessible over the Internet. The technique of texturing multiple live video streams onto virtual objects (Reynard, 1998) can be used to integrate these various video streams into the virtual world. Video textures might be connected to players’ avatars or to virtual walls (to create mixed reality boundaries (Benford *et al*, 1998)), providing production crew with some knowledge of their local environments.

A second problem demonstrated by *Can You See Me Now*, is the need to better integrate multiple sources of information (e.g., device status, network status, sensor data etc.) into a shared display. One solution might be to make use of larger shared displays that provide several team members with an overview of a situation. Even then, there is the problem of scale, i.e., of monitoring many players who are spread across many virtual worlds. This

may be addressed through activity logging and visualization tools. CHI and CSCW research has produced a range of techniques for visualising patterns and histories of communication among people, including PeopleGarden's data portraits of message board use (Xiong, 1999), visualisations of message threads (Smith, 2001) and Chat Circles' visualisation of synchronous conversations (Viegas, 1999). In our own work, we have recently developed a mechanism to extract information about key scenes (encounters among participants) from live or recorded 3D environments. Our technique applies a clustering algorithm to participants' positions in order to extract momentary groupings and then applies a further set of rules to map these onto longer-term scenes. The resulting scene events can be visualized or used to trigger changes in monitoring interfaces (e.g., highlighting the location of a new scene on a map). Figure 17 shows a prototype tabletop interface for monitoring an on-line experience. This consists of a horizontal map view of a virtual world, within which a virtual camera can be positioned to control a vertical detailed view. Ongoing scenes (as determined by our scene extraction technique) are then highlighted on the map using additional graphics.

5.2.3 Human versus automated orchestration

The focus on this paper has been on how humans orchestrate experiences, supported by appropriate technologies. Human orchestration brings the advantages of being responsive, adaptive and intelligent. For these reasons, it is an appropriate approach to orchestrating experiences that are exploratory (such as our three examples). However, it does not scale well. Human expertise is a scarce and expensive resource. For economic reasons alone, the approaches explored in this paper will probably not directly scale to support massively multiplayer games. We see two broad approaches to this challenge.

The first is to extend the approach of participants helping one another, i.e., to push the work of orchestration down onto the players. Our three experiences have already shown examples of this happening and perhaps more importantly, suggested that the resulting collaboration between players can heighten their engagement with the experience. The idea of participants taking on responsibility for technical control of an online experience is familiar from current large-scale persistent environments (Mynatt, 1997). Examples include MUDS and MOOS where committed and successful members of a community may progress to become 'wizards' and so become responsible for running the world. It is easy to imagine how this approach could be applied to a larger-scale version of Avatar Farm, with successful players gaining the ability to grant further abilities to other players.

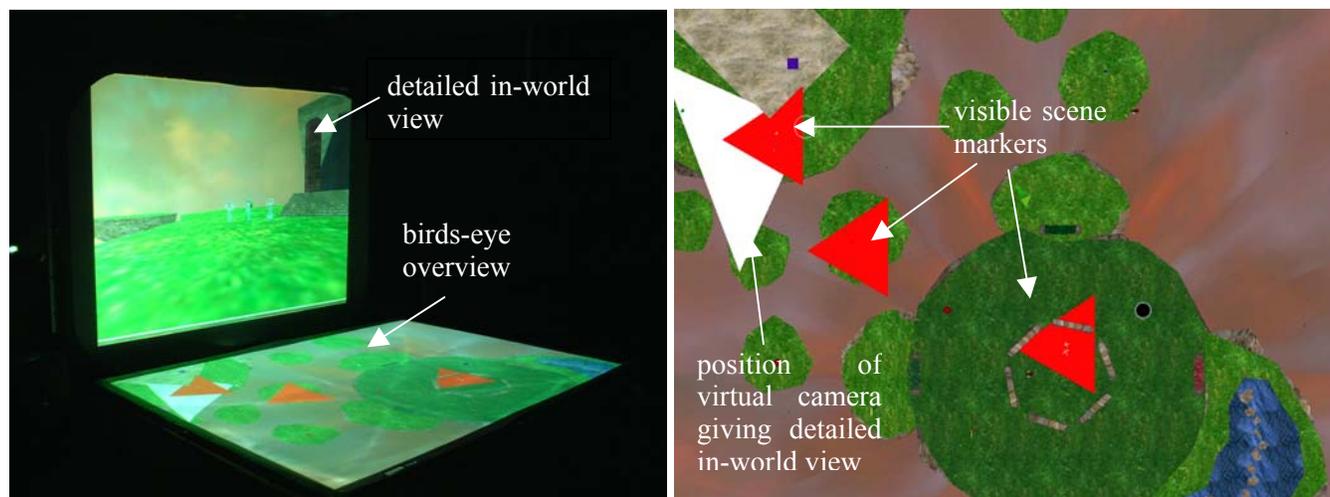


Figure 17: (left) overview of the table monitoring interface and (right) detail of the map component

The second is to integrate pre-programmed interaction and human orchestration. In multiplayer online games, the action is largely predetermined and preprogrammed. These games appear to work well without the need for large

teams of human orchestrators (although the nature of the human orchestration actually required by the players would be worthy of a further separate study). On the other hand, there is relatively little scope for improvisation and much less focus on the role of audio dialogue and role-play compared to direct action. The solution may be to combine both approaches. Large-scale persistent experiences may involve a high degree of automated game-play, but then introduce human orchestration to achieve particularly complex scenes and to resolve exceptional difficulties. One can imagine SWOT teams of orchestrators who are monitoring the experience and who are able to drop in to key locations when required (perhaps by assuming control of normally pre-programmed non-player characters). One can also see human and automated orchestration as supporting different points in the design process. Human automation may be useful for testing out and refining new interactions (in an improvised ‘Wizard of Oz’ manner). Once perfected, these may then be automated. In a large-scale persistent environment, this would be an ongoing process within the experience. Human orchestration would support the introduction of new features that might eventually make their way into the underlying implementation.

6. Summary

The key point of this paper has been to raise the profile of orchestration as an important issue for real-time participatory experiences. In the three examples studied this has been a real-time process involving collaboration between performers and a behind-the-scenes production crew. Although the importance of orchestration is widely recognised for conventional theatrical and musical performances, it has not been considered widely in the design of interactive software where the dominant focus has been on plug-and-play experiences, often for individual participants. This paper argues that support for orchestration needs to be high up the agenda of software and experience designers.

Through the study of three contrasting participatory experiences, Avatar Farm, Desert Rain, and Can You See me Now? we have uncovered a variety of orchestration issues that need to be considered by experience designers including: admitting participants to an experience, training and familiarisation, establishing engagement, avoiding fractures in engagement, monitoring physical and virtual activity and system state, intervening to deal with problems; coordinating complex actions, managing pace and timing, and closing an experience. We have also demonstrated new technologies to deal with these issues including tools for invisible stage hands to track and manipulate players and objects in a virtual world, and the use of semi-transparent partitions and asymmetric screens to monitor their activities in the physical world. We have also analysed different work practices that successfully integrate these technologies into the smooth running of an experience. Finally, we have considered future challenges for orchestration as the scale and mobility of participatory experiences increase.

We hope that our observations and proposals will be of benefit to the experience designers in areas such as games, performances, museum and art installations and theme park rides as well as to the designers of supporting software and hardware platforms. We also hope that they may be of relevance to other applications and fields that involve complex collaborative interactions, such as simulation, training, education and conferencing. Finally, we hope that our observations add to the body of studies of work practice with technology that underpin the field of CSCW.

As a final note, we plan to take up the challenge of orchestrating large-scale distributed events through a further collaboration with Blast Theory; an experience provisionally entitled Citywide that extends Can You See me Now? towards a more narrative-oriented form and that involves groups of participants spread across an entire city interacting with other groups who are online.

References

- Barrus, J. W., Waters, R. C. and Anderson, D. B. (1996): Locales: supporting large multiuser virtual environments, *IEEE Computer Graphics and Applications*, 16 (6). pp 50-57, Nov 1997, IEEE Computer Society.
- Benford, S., Greenhalgh, C., Craven, C., Walker, G., Regan T., Morphett, J., and Wyver, J. (2000): Inhabited Television: broadcasting interaction from within collaborative virtual environments, *ACM Transactions on Computer Human Interaction*, ACM Press, December 2000.
- Benford, S., Greenhalgh, C., Reynard, G., Brown, C., and Koleva, B. (1998): Understanding and Constructing Shared Spaces with Mixed Reality Boundaries, *ACM Transactions on Computer Human Interaction*, 5 (3), 185-223, ACM Press.
- Bjork, S., Falk, J., Hansson, R., Ljungstrand, P. (2001): "Pirates! Using the Physical World as a Game Board", *Proceedings of Interact 2001*.
- BotFighters (2002): web site (In Swedish) <http://www.teliamobile.se/botfighters/>, News article on BotFighters http://on.magazine.se/pdf/6_2000/Games1.On_6_2000.pdf
- Bowers, J., Pycok, J. and O'Brien, J., (1996): Talk and Embodiment in Collaborative Virtual Environments, *Proc. ACM CHI'96*, ACM Press, 1996.
- Bullock, A. and Benford S. (1997): Access Control in Virtual Environments, in *Symposium on Virtual Reality Software and Technology (VRST'97)*, Lausanne, Switzerland, 1997, ACM Press.
- Büscher, M., O'Brien, J., Rodden, T., and Trevor, J. (2000): He's behind you: the experience of presence in shared virtual environments, *Collaborative Virtual Environments* (E. Churchill, D. Snowdon & A. Munro), London 2000
- Craven, M., Benford, S., Greenhalgh, C., Wyver, J., Brazier, C. J., Oldroyd, A. & Regan, T. (2000): Ages of Avatar: Community Building for Inhabited Television, *ACM CVE'2000*, San Francisco, September 2000.
- Dodsworth Jr., C. (1997): *Digital Illusions: entertaining the future with high technology*, Addison Wesley.
- Drozd, A., Bowers J., Benford S., Greenhalgh, C. and Fraser, M. (2001): Collaboratively Improvising Magic: An Approach to Managing Participation in an On-Line Drama, *Proc. ECSCW 2001*, Bonn, Germany, Kluwer, September 2001.
- Geocaching web site <http://www.geocaching.com/> (2002)
- Goffman, E. (1987): *Asylums. Essays on the Social Situation of Mental Patients and Other Inmates*. London.
- Greenhalgh, C. M. and Benford, S. D. (1999): Supporting Rich And Dynamic Communication in Large Scale Collaborative Virtual Environments, *Presence: Teleoperators and Virtual Environments* , 8 (1), pp.14-35, February 1999, MIT Press
- Greenhalgh, C., Purbrick, J. and Snowdon, D. (2000): Inside MASSIVE-3: Flexible Support for Data Consistency and World Structuring, in *Proceedings of the Third ACM Conference on Collaborative Virtual Environments (CVE 2000)*, San Francisco, CA, USA, September 2000, pp. 119-127, ACM Press
- Greenhalgh, C., Purbrick, J., Benford, S., Craven, M., Drozd, A. and Taylor, I. (2000): Temporal links: recording and replaying virtual environments, *Proc. ACM Multimedia 2000*, L.A., Oct 2000.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S. and Greenhalgh, C., (1998): Fragmented interaction: establishing mutual orientation in virtual environments, *Proc. ACM Conference on Computer Supported Cooperative Work (CSCW'98)*, 14-18 November, 1998, Seattle, WA, AC, New York, pp 217-226.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S., and Greenhalgh, C. (2001): Fragmented Interaction: Sharing objects in Collaborative Virtual Environments, *ACM Transactions on Computer Human Interaction*, 2001.
- Kamberg, M. (1998): Seafari: An expedition into motion base ride filmmaking, in virtual reality, *Digital Illusion: entertaining the future with high technology*, (Dodsworth, C.), Addison-Wesley.
- Koleva, B., Schnädelbach H., Benford, S. & Greenhalgh, C. (2000): Traversable Interfaces Between Real and Virtual Worlds, *CHI'2000*, 2000.
- Koleva, B., Taylor, I., Benford, S., Fraser, M., Greenhalgh, C., Schnadelbach, H., vom Lehn, D., Row-Farr, J., Adams, M. (2001): Orchestrating a Mixed Reality Performance, *Proc. CHI'2001*, Seattle, USA, April 2001.
- Laurel, B. (1992): *Computers as Theatre*, Addison-Wesley.
- Mynatt, E. F., Adler, A., Ito, M., O'Day, V. (1997): Design For Network Communities, *Proc. CHI 97*, 210-217, ACM Press.
- Pausch, R., Snoddy, J., Taylor, R., Watson, S. & Haseltine, E. (1998): Disney's Aladdin: first steps towards storytelling in virtual reality, *Digital Illusion: entertaining the future with high technology*, (Dodsworth, C.), 357-372, Addison-Wesley.
- Reynard, G., Benford, S., and Greenhalgh, C. (1998): Awareness Driven Video Quality of Service in Collaborative Virtual Environments, *Proc. ACM Conference on Human Factors in Computing Systems (CHI'98)*, Los Angeles, March 1998.

- Singletary, B., Lyons, K., Gandy, M. & Pair, J. (2000a): "Towards Augmented Reality Gaming", *Proc. IMAGINA 2000*.
- Slater, M., Steed, A., McCarthy, J., and Marinelli, F., (1998): "The Influence of Body Movement on Presence in Virtual Environments", *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 40(3), Sept 1998, pp469-477
- Smith, G. and Mariani, J. (1997): Using subjective views to enhance 3D applications, *Proceedings of the 1997 ACM symposium on Virtual reality software and technology (VRST 97)*, 1997, Pages 139 - 146
- Smith, M. A. And Fiore, A. T. (2001): Visualization Components for Persistent Conversations, *Proc. CHI 2001, CHI Letters*, 3 (1), 136-143, ACM Press.
- Viegas, F. B. and Donath, J. S. (1999): Chat Circles, *Proc ACM CHI 99*, 9-16, ACM Press, 1999.
- Xiong, R. And Donath, J. (1999): Peoplegarden: Creating Data Portraits of Users, *Proc. UIST'99, CHI Letters*, 1 (1), 37-44, ACM Press

Acknowledgements

We gratefully acknowledge the support of the UK's Engineering and Physical Sciences Research Council (EPSRC) through the Equator Interdisciplinary Research Collaboration (www.equator.ac.uk). We also gratefully acknowledge the support of the European commission for the SHAPE project, part of the Disappearing Computer initiative within the European Framework V Future and Emerging Technologies programme (www.nada.kth.se/shape).

We are grateful to John Wyver and his team at Illuminations Television, to Jason Morphett and his team at British Telecom Laboratories, and to Alex Butterworth for their collaboration on Avatar Farm. Thanks to Andrew Chetty and his team at the NOW festival, Jeffry Shaw and his team at the ZKM and to the Arts Council of England for collaborating on and supporting the development of Desert Rain. Finally, thanks to The Arts Council of England, BBC Online and b.tv for commissioning Can You See Me Now? as part of the Shooting Live Artists programme.

Additional web resources

For further information, including colour images and video clips:

Blast Theory: www.blasttheory.co.uk

The Mixed Reality Laboratory: www.mrl.nott.ac.uk

The Equator Interdisciplinary Research Collaboration: www.equator.ac.uk

Desert Rain: www.blasttheory.co.uk

Can You See Me Now?: www.canyouseemenow.co.uk

Avatar Farm www.crg.cs.nott.ac.uk

Out of This World: www.crg.cs.nott.ac.uk