

Applications of Temporal Links: Recording and Replaying Virtual Environments

Chris Greenhalgh, Martin Flintham, Jim Purbrick, Steve Benford
School of Computer Science and IT
University of Nottingham
Email: {cmg, mxj, jcp, sdb}@cs.nott.ac.uk

Abstract

Temporal links allow recordings of multi-user sessions to be dynamically inserted into current virtual worlds in a flexible and principled way. This paper explores key applications of temporal links, showing how they can add new content to virtual worlds, support usability and system evaluation, and link VR to other media such as film and television. These applications illustrate just some of the possibilities of a ubiquitous and flexible record and replay facility such as temporal links. Building on our experience of implementing temporal links in the MASSIVE-3 system, we identify key requirements for system architectures that wish to support an equivalent mechanism.

1. Introduction

Virtual reality has been dominated by spatial concerns; how to create 3D geometries, how to arrange these to form a virtual environment, how to partition an environment into manageable chunks, and how to link an environment to other environments. Relatively little attention has been paid to the temporal structure of environments; they remain resolutely focused on real-time experiences of the present.

This paper explores how virtual reality applications can be given a flexible temporal structure in which action in the present is combined with action from the past. Specifically, it explores different applications of a mechanism called *temporal links* that allows real-time virtual environments to be linked to recordings of prior virtual environments so that the two appear to be overlaid [7]. These applications show how temporal links can be used to create new content for virtual environments, how they can support VR researchers in conducting usability and system evaluation, and how they can allow VR to be integrated with other media such as film and television. This paper also discusses key requirements for system architectures that wish to support an equivalent

mechanism to temporal links. We begin with a general introduction to temporal links as realized in the MASSIVE-3 system.

2. Introducing temporal links

MASSIVE-3 [6] is a multi-user VR (Collaborative Virtual Environment, or CVE) system that supports populated and interactive virtual worlds combining 3D graphics, real-time audio and streamed video. Following the SPLINE system [1], MASSIVE-3 allows its virtual worlds to be spatially structured as multiple linked ‘locales’, each of which can be an arbitrary virtual space (e.g. a room, building, open region) with its own Cartesian coordinate system.

MASSIVE-3 extends locales by allowing current locales to be linked to *recordings* of other locales. A recording captures all movements, object interactions, and speech within a locale. A live user in a current locale can then experience the recorded virtual action happening around them and can move to view and hear it from any perspective. However, in our current realization they cannot directly interact with the recorded material (e.g., colliding with a recorded object). We call this kind of inter-locale link a ‘temporal link’, because it has temporal properties (*when* it links to), as well as the spatial and presentational (appearance-specifying) properties common to all inter-locale links.

The key feature of temporal links is their flexibility with regard to how current locales can be related to recorded locales. This arises from a number of configurable properties that are associated with each temporal link. Full details of these along with their implementations can be found in [7]. The following is a brief summary.

Temporal properties – specify the relationship between live time in the current locale and recorded time in a recorded locale. This is as specified by an offset that

relates current time to a specific moment in recorded time and a rate that sets the speed and direction in which recorded time unfolds relative to current time. Time in the recorded locale may be speeded up or slowed down relative to live time and may run forwards or backwards.

Spatial properties – specify the spatial relationship between the live and recorded locales. Recorded material may appear to be directly overlaid on, or may be scaled, rotated and sheared relative to the live material. For example, a recording may appear as a world in miniature.

Presentational properties – control the appearance of recorded material, for example it may be rendered translucent to make it ghostlike, have its color removed to make it shadow-like, or may be indistinguishable from the live material.

Each live locale may be linked to many recorded locales so that several scenes from the past appear to intermingle. A live locale can also be linked to a recording of itself in the past. Furthermore, temporal links can be nested. A live locale can be linked to a recorded locale that is itself linked to further recorded locales. This enables complex layered recordings to be created.

Finally, temporal links can be created and deleted on the fly in the same way as normal (non-temporal) inter-locale links. Their properties can also be modified on the fly. For example, a user can dynamically manipulate the rate temporal property in order to wind forwards and backwards through a recording as they would a videotape.

3. Related Work

Several existing CVEs support some form of record and replay facility. Fielded DIS systems include tools for recording and recreating entire sessions in terms of all of the network messages exchanged [9].

In the CAVERNsoft system [12], recording of an avatar's movements and audio is possible as part of general support for persistent virtual environments. A related facility has been used to create the Vmail system [10], a form of VR email, and to create guided tours within tele-immersion applications.

As part of the COVEN project, the DIVE system was extended with event logging facilities that could completely record an entire virtual environment and the activity within it [4]. Although initially implemented to support the statistical analysis of patterns of user activity in relation to network traffic, this recording facility was subsequently extended with a replay facility to allow a previous session to be recreated (although not within another live virtual

environment). The dVS system supports a similar facility for recording and then replaying a virtual environment.

Multi-player 3D games also use record and replay techniques to show highlights of previous game-play, examples of which include FIFA soccer from Electronic Arts and the automobile game Driver from GT Interactive Software, the latter allowing players to edit together their own movies from recordings of their own actions.

In a related area, recent work on 'virtualized reality' has developed techniques for capturing live action within a physical environment by synthesizing recordings from multiple video cameras to produce a 3D graphical simulation [11]. Furthermore, a script language VRML History has been proposed to facilitate time navigation in WWW browser-based 3D-worlds [13].

Clearly, recording and replaying virtual worlds is not unique to our work, but we believe that it is currently under-utilised by the VR community. We also argue that the 'temporal links' metaphor and mechanism provide a very simple, elegant and flexible way to use and control recordings of virtual activity, making this technology more accessible. The following sections describe different applications of temporal links to support these assertions.

4. New content for virtual environments

Our first applications of temporal links focus on new ways of creating content for virtual environments.

4.1. Rapid authoring by enacting and recording

Temporal links can bridge the gap between pre-programmed and live content for virtual environments. Actors can enter an environment and quickly enact and record complex scenes that are then played out many times.

As an example of this approach, we were recently involved in a project to recreate Nottingham's 12th century medieval castle (which is linked to the legend of Robin Hood) on the site of today's 17th century Ducal Palace. The castle museum already employs a variety of techniques to help visitors visualise how the medieval castle would have been including physical models, artists impressions, information boards at key viewpoints, and actors who wander the castle grounds playing medieval characters.

We obtained a 3D model of the medieval castle and imported it into the MASSIVE-3 system. We then hired one of the actors (who had not previously used a VR system) to play their character, a medieval guard, as an avatar within this model. Over the course of five hours we recorded approximately forty minutes of content in which this

medieval guard gave a series of guided tours through the ancient castle, talking about key features as they went. The 3D model and recorded scenes were subsequently demonstrated at the Castle. Figure 1 shows a moment from one of the recorded tours. The castle model is part of the live locale and the guard avatar part of a recorded locale.



Figure 1. A scene from a recorded guided tour

This example demonstrates the most straightforward configuration of temporal links in which the recorded locale is directly overlaid on the live locale, has a fixed temporal offset, plays forwards at the same rate as live time, and is not altered in its presentation.

One can imagine extending this example with multiple characters and more complex scenes. Of course, not being able to interact with 3D recordings limits the use of this technique – it will not replace pre-programmed content, but it can certainly serve as a useful additional authoring tool.

4.2. 3D flashbacks

Our second application of temporal links involves creating flashbacks within a story in a virtual environment.

Our example is taken from an event called Avatar Farm, an experiment in inhabited television [5] that sought to create a sophisticated non-linear drama in four connected virtual worlds and that was implemented in MASSIVE-3 [5]. Avatar Farm was a relatively rich and complex virtual event and provided the raw material for several of the examples in this paper. More detail on Avatar Farm and these examples can be found in [2].

Avatar Farm was primarily improvised, based on a nominal overall structure and a set of stereotyped characters with a common history. This structure and background was initially unknown to the four members of public taking part,

who started the experience from a position of ignorance. During the experience (which comprised four chapters, each lasting about 30 minutes) the back story – the history and tensions inherent between the various characters – was revealed to the public players by means of ‘flashbacks’ to significant past activities within the virtual worlds. The flashbacks appeared as semi-transparent or ‘ghost-like’ avatars and objects re-enacting the past within the current live virtual worlds (see figure 2). They were incorporated into the plot as time-rifts that were triggered (by a member of the production crew from behind the scenes) in response to the actions of the players and actors.



Figure 2. Live players watch a ghostly flashback

Flashbacks were realised by dynamically creating temporal links to scenes that had been improvised by the actors and recorded on the previous day. The temporal links were configured to change the presentation of the recording, making the replayed activity semi-transparent.

One flashback included a further embedded flashback in which several actors could be seen watching and discussing a previous time-rift. This exploited the ability to layer 3D recordings using chains of temporal links.

This example demonstrates the ability to create complex narrative structures within virtual environments that involve reference to past events. Such techniques are an important part of established media such as film, and we anticipate that they will become increasingly important for VR as it matures as a medium for drama and storytelling.

4.3. Virtual messaging

Our third example further explores the potential of creating layered recordings.

Inspired by the Vmail system [10], we have used temporal links to support asynchronous communication in multi-user virtual worlds through virtual messaging.

To facilitate this we built a simple interface that allowed users in a virtual world to dynamically create and replay small recordings. These recordings are independent of the user that created them, allowing all users to access the recorded events at any time.

As shown in figure 3, our interface was embodied as a virtual peach tree, with each peach containing a different recorded message. A user could interact with an existing peach in order to view a message and could interact with the tree to obtain a new peach and so record a new message.

The properties of temporal links can be related to two familiar features of messaging systems. First, embedding multiple temporal links in the live locale enables access to many messages in the same world. Second, layered recordings can be used to implement reply and forwarding functions. For example, a user could play an existing message and at the same time record a response or annotation as a new message. Figure 3 shows an example in which the avatar on the left is playing out a message in which the avatar on the right is responding to an original message created by the avatar on the right.



Figure 3. Nested virtual messages

5. Supporting the VR researcher

Our next set of applications focus on how temporal links can support researchers in carrying out experiments.

5.1. Reviewing and analysing user interaction

There is already a tradition of using ethnographic techniques in which experts observe, record and analyse social interaction among users in the study and design of CVEs (for examples, see [8,5]). Up to now, this has relied on the analysis of videotaped output from different participants' machines. However, video-based approaches soon become infeasible, even with moderate numbers of participants, due to the number of viewpoints to be recorded and the amount of video data to be cross-referenced.

Temporal links allow experts to repeatedly review events in virtual environments from any perspective. We have constructed a viewing 'table' specifically to support small group exploration and discussion of recordings (see [2]); this combines a zoomable 'live' map view of a virtual world that shows avatars and objects (horizontal projection), with a steerable in-world camera for detailed viewing (vertical projection) as shown in figure 4. Users can rapidly navigate the virtual space using the map view, and maintain an overview of activities in the virtual world. Using the in-world view (and associated audio) they can follow specific interactions and activities.




Figure 4. Using the 'table' to explore Avatar Farm

We have used the table to conduct ethnographic studies of user behavior. For example, we studied the activities of the invisible 'backstage' helpers in Avatar Farm who were responsible for various behind-the-scenes activities and pieces of 'magic' such as making key props appear and disappear. By making these invisible helpers visible within the recordings we were able to observe their interaction and transcribe their conversations. Our analysis revealed the confusions that arose from these helpers being mutually

invisible, leading to proposals for how our helper technology should be redesigned (see [3] for a full account).

An extension of this approach is to provide tools within virtual environments for participants to review and analyse activities while online. These might support distributed meetings between ethnographers or perhaps more importantly, allow participants themselves to review and discuss their own activities. This could be a useful extension to training and simulation applications of VR.

We have implemented an example of such a tool called the Hologrid. This provides a world-in-miniature view of the virtual world within which it is located (analogous to the physical models that can often be found in the foyers of large buildings). It also provides a series of controls through which avatars can browse forwards and backwards through recordings at different speeds. Figure 5 shows the Hologrid being used in one of the worlds of Avatar Farm called Nirvana. We see the Hologrid in the foreground currently showing all of Nirvana in miniature, with a miniature view of the Temple of Nirvana towards the right. The Hologrid is itself located inside the full-scale Temple of Nirvana. We see a viewer at the far end of the Hologrid, and beyond them we see out through the door of the temple into the Nirvana.




Figure 5. The Hologrid being used in Nirvana

5.2. Repeatable scenarios for system evaluation

Temporal links are also useful for system level evaluation. Researchers can compare the performance of different versions of a system (e.g., measuring network traffic, CPU load etc.) when running exactly the same application scenario. Realistic and complex scenarios can be acquired by making recordings of everyday uses of CVEs. Finally, researchers will be able to build up libraries

of test scenarios and (given a common format for 3D recordings) exchange or publish test data alongside their results.

We have used this approach to support our ongoing research into persistence and extensible infrastructures for CVEs [14]. We began by recording several weeks of use of a persistent editable virtual environment. By replaying these recordings we were able to analyse the activity within the environment from a system perspective. Specifically, we examined the temporal patterns of the use of virtual objects (creation, modification and deletion). We performed this analysis by replaying the recordings as fast as possible (not linked to real-time at all), allowing weeks of activity in the virtual world to be replayed and analysed in hours.

This quantitative analysis [14] revealed significant differences in the life-cycles of the various classes of items in the persistent environment; user embodiments, walls and other objects each had very different characteristic in terms of when and how they were updated. This led us to build a persistent CVE platform that allows differentiation between virtual items at an infrastructural level, for example making only certain items persistent, or using different consistency methods for different items. Having built this platform we are now using the same replay facilities to evaluate it. We are able to replay the same patterns of user and world activity through different versions of the system, comparing the performance of the original system and of our new flexible infrastructure (in various configurations). We are thus able to demonstrate and measure the potential benefits of our infrastructure.

6. Integrating VR with other media

The growth in 3D games and rides show how VR is rapidly maturing as a new entertainment medium. Our final applications of temporal links focus on the broader nature of VR as a medium, especially how it can be integrated with established media such as film and television.

6.1. Links to animation

In addition to introducing 3D recordings into live worlds, temporal links – and the recordings that underlie them – can be used as a bridge between virtual environments and animation. For example, our current implementation allows scenes from recordings to be exported as animation files for Lightwave (version 5) that can then be lit and rendered off line. Figure 6 shows a frame from a Lightwave animation of a scene from Avatar Farm.



Figure 6. Offline rendered image from Avatar Farm

As well as exporting activities exactly as they occurred within the system, recordings can be modified or adjusted. For example, during Avatar Farm we used graphical ‘speech bubbles’ to indicate when people were speaking (determined from audio energy levels). For the animation shown above this was replaced within the recording by a simple animated mouth (using geometry morphing), a capability that was only added after Avatar Farm took place.

CVEs might therefore be used to ‘block out’ an animation (establishing positions and structure). They might even be viewed as a kind of virtual ‘social motion capture’ system for animation, enabling all of the actions, speech and world interactions of a group of avatars to be exported into animation software for post-production and rendering.

6.2. Viewing non-linear stories in VR

Finally, it is interesting to consider VR as a storytelling medium in its own right and to explore the kinds of interfaces that future ‘viewers’ of VR might use. One feature of VR as a medium is its non-linearity. A VR story such as Avatar Farm may be spread over one or more worlds and consist of many parallel and interwoven threads of action. We have already seen two interfaces for viewing VR recordings in the form of the Hologrid (an in-world interface) and the Table (a physical interface). We have recently experimented with some others.

In one experiment, we used the Avatar Farm recordings to stage an on-line ‘promenade performance’ of Avatar Farm (a promenade performance in theatre is where the action takes place among the audience, who move around

an extended stage in order to view it). Using temporal links, the recorded chapters and backstory segments were replayed within a live virtual world. Members of a viewing audience could enter this world over a network (using the normal MASSIVE-3 user applications). They could move freely within the recordings, for example following particular characters, discussing the action with one another.

In a second experiment, we played the Avatar Farm recordings in an immersive cinema facility. The EVE Dome at the ZKM in Karlsruhe, Germany, is a large dome environment that can hold an audience of approximately ten people. We used four networked workstations and a world server to generate four coordinated views of the Avatar Farm recordings. Each of the four views was projected in one quadrant of the dome, together with its audio, giving four different perspectives from the same recording. For much of the time, each view tracked a different player. However, when the players met in the virtual world the views dynamically reconfigured to give a panorama.

Viewers controlled the viewing experience by moving around within the dome and attending to the different views as they chose. Their movements were tracked using an IR camera, which in turn modified the 8 channel audio mix (2 channels each from 4 views) to try to give a more comprehensible audio field (compare the movie Time Code [15], which has four simultaneous tiled views, and in which Mike Figgis carefully mixes the audio in post production so that it makes sense to viewers in a traditional cinema). Figure 7 shows two of the projected views inside the EVE Dome.



Figure 7. Watching Avatar Farm in the EVE Dome

7. System requirements for temporal links

We suggest that the scope of the applications described above shows great potential for the use of 3D record and replay mechanisms in VR, as well as the advantages of wrapping them up in flexible general-purpose abstraction such as temporal links. This section now turns our attention to underlying system architectures. More specifically, it discusses key requirements for systems that wish to support temporal links or an equivalent mechanism.

7.1. Recording

In order to record the activity within a virtual world the system must be able to capture the complete momentary state of the virtual world (at the beginning of the recording) and to monitor all of the changes that are made to the virtual world. It must also be able to externalize this state and these changes to some form of secondary storage with sufficient information to unambiguously recreate them.

In MASSIVE-3, all changes to the virtual world are represented by event objects that move through the system in a well-defined way. To support temporal links each event is written to disk as it is enacted, together with a time stamp. The events are externalized using the system's normal serialization mechanism (which is also used for network communication).

Equivalent approaches to data sharing (e.g. distributed shared memory) could be logged within the communication/support system (provided that they are consistently applied). An API-based system could be instrumented so that it logged all invocations and relevant parameters. It is not normally possible to record activities that bypass these – or equivalent – mechanisms.

7.2. Recontextualised replay

In order to replay past activity it must be possible to re-enact recorded operations. Furthermore, in some situations the recorded information has to be adjusted before replaying. For example, MASSIVE-3 allows a virtual world that is still running to connect to an earlier version of itself. In this case there may be different versions of the same virtual object in existence at the same time, one from the present and one from each recording. Care must be taken to ensure that these remain independent.

Temporal and inter-locale links also have to be changed when they are replayed in order to remain consistent. Consider an inter-locale link in a live locale, A, that points to the live version of another locale, B. When A is replayed

this inter-locale link should now point to a corresponding replay of B (rather than to B itself, which may still be running). Similarly, replaying a temporal link from A to B will now link from the past of A to the *further* past of B.

7.3. Explicit notion of time

If a system supports deterministic behaviours that evolve with time (and that are not represented by explicit frame-by-frame changes in the recording) then the system must also manipulate the replay's notion of time. For example, MASSIVE-3 allows positions to be specified as a function of time, and locally evaluated at each machine with no further explicit events. When these behaviours are re-evaluated in a replay they must be informed of the nominal time within the recording (which may be faster or slower than real-time, or may even be moving backwards) rather than the current time. This also allows a replay to extrapolate to states that were never actually observed, e.g. during a slow motion replay.

7.4. Explicit temporal links and segmented worlds

Our last point concerns efficiency, rather than being an absolute requirement. As noted in section 2, MASSIVE-3 divides virtual worlds into locales and relies on explicit links (temporal and non-temporal) to combine separate locales into a single virtual world. Each locale is recorded independently, together with any links that it contains. Replaying a recorded locale creates a new locale that represents that locale at the requested moment in time.

This approach allows the recording and replay facilities to be more efficient. For example, recording a locale with a temporal link in it just records the temporal link itself and the live (new) activity; when the locale is replayed the original recording associated with the temporal link is used (if required) to generate the nested replay. The alternative – replaying a recording as part of the live activity – means that a replay within a recording will actually be copied into the new recording (rather than referenced).

MASSIVE-3 also uses these explicit links to declaratively specify the spatial and presentational properties of a temporal link, rather than requiring the replay to be modified on the fly (each client performs the spatial and presentational modifications itself on the same shared replay). Explicit linking also allows “impossible” worlds to be constructed, such as a world that connects to its own recent past, without recordings becoming infinitely large (the replay has a link to its past which has a link to its past which...).

8. Summary

We have explored various applications of temporal links and their underlying record and replay mechanisms to extend and support multi-user virtual environments. Three key points emerge from this work.

First, temporal links can enhance virtual environments in many ways, from new kinds of content, through support for experimental research, to new ways of (re)viewing events in virtual worlds. Although the majority of our examples have been based on Avatar Farm, an entertainment-oriented application, we suggest that the same approaches can be applied in other domains such as: education, training and simulation; engineering, industrial and architectural design; and ergonomics and user studies.

Second, wrapping up a 3D record and replay mechanism within the more general abstraction of temporal links is a necessary step to realising its full potential, as this enables users and programmers to manipulate the relationships between live and recorded time and to create layered and inter-linked recordings.

Third, underlying systems that wish to support a flexible approach to 3D record and replay such as temporal links need to exhibit certain key characteristics in terms of how they handle world evolution, time and structure.

We hope that these examples and observations will stimulate other researchers and practitioners to adopt a flexible approach to the treatment of time in virtual environments, and so develop systems and applications systems that provide richer experiences to users.

9. Acknowledgments

We gratefully acknowledge the UK's EPSRC for funding the Equator Interdisciplinary Research Collaboration; Illuminations TV and British Telecom for their collaboration on Avatar Farm; and Nottingham Castle Museum, CIMI and Byte Ltd for use of the 3D model of Nottingham's medieval castle.

10. References

- [1] Barrus, J.W., Waters, R.C. and Anderson, D.B., Locales: Supporting Large Multiuser Virtual Environments, *Computer Graphics and Applications*, 16(6), Nov 1996, pp. 50-57.
- [2] Craven, M., Taylor, I., Drozd, A., Purbrick, J., Greenhalgh, C., Benford, S., Fraser, M., Bowers, J., Jää-Åro, K., Lintermann, B, Hoch, M., Exploiting Interactivity, Influence, Space and Time to Explore Non-linear Drama in Virtual Worlds., *CHI'2001*, Seattle, US, April 2-6, 2001, ACM Press.
- [3] Drozd, A., Bowers, J., Benford, S., Greenhalgh, C., Fraser, M., Collaboratively improvising magic: An Approach to Managing Participation in an On-Line Drama, *Proc 7th European Conference on Computer Supported Cooperative Work (ECSCW'01)*. Bonn, Germany, September 2001, Kluwer Academic Publishers.
- [4] Frécon, E., Greenhalgh, C. and Stenius, M., The DIVE-BONE – an application-level network architecture for Internet-based CVEs, *ACM Virtual Reality Software and Technology (VRST'99)*, London, Dec 1999, pp. 58-85.
- [5] Greenhalgh, C., Benford, S., Taylor, I., Bowers, J., Walker, G. and Wyver, J., Creating a Live Broadcast from a Virtual Environment, *SIGGRAPH '99*, L.A. Aug. 1999, pp. 375-384.
- [6] Greenhalgh, C., Purbrick, J., and Snowdon, D., Inside MASSIVE-3: Flexible Support for Data Consistency and World Structuring, *Proc. Third International Conference on Collaborative Virtual Environments (CVE2000)*, Sept 2000, San Francisco, USA, pp. 119-127, ACM, New York.
- [7] Greenhalgh, C., Purbrick, J., Benford, S., Craven, M., Drozd, A. and Taylor, I., Temporal links: recording and replaying virtual environments, *ACM Multimedia 2000*, L.A., Oct 2000.
- [8] Hindmarsh, J., Fraser, M., Heath, C., Benford S., and Greenhalgh, C., Object-focused interaction in collaborative virtual environments;; *ACM Transactions on Computer-Human Interaction*. 7, 4 (2000), Pages 477 - 509
- [9] IEEE, Institute of Electrical and Electronic Engineers, International Standard, ANSI/IEEE Std 1278-1993, Standard for Information Technology, Protocols for Distributed Interactive Simulation (March 1993).
- [10] Imai, T., Johnson, A., and DeFanti, T., The virtual mail system, *IEEE Virtual Reality (VR'99)*, Houston, USA, March 1999, p. 78.
- [11] Kanade, T., Rander, P., Vedula, S. and Saito, H., Virtualized reality: digitizing a 3D time varying event as is and in real time, in Ohta, Y. and Tamura, H. (eds.), *Mixed Reality, Merging Real and Virtual Worlds*, Springer, 1999, pp. 41-57.
- [12] Leigh, J., Johnson, A., DeFanti, T., Brown, M., et al., A review of tele-immersive applications in the CAVE research network, *IEEE Virtual Reality (VR'99)*, Houston, USA, March 1999, pp. 180-187.
- [13] Luttermann, H. and Grauer, M., VRML History: Storing And Browsing Temporal 3D-Worlds, *ACM Virtual Reality Modeling Language Symposium (VRML'99)*, Paderborn, Germany, February 1999, pp. 153-181.
- [14] Purbrick, J., Greenhalgh, C., "Collaborative Creation of a Persistent Virtual World", in *Proc. Human-Computer Interaction - INTERACT'01*, 9-13 July 2001, Tokyo, Japan, Michitaka Hirose (ed.), IOS Press, Amsterdam.
- [15] <http://www.timecode2000.com/> (verified 27/4/2001)