

# The CaveUT System: Immersive Entertainment Based on a Game Engine

Jeffrey Jacobson(1), Marc Le Renard (2), Jean-Luc Lugin (3) and Marc Cavazza (3)

(1) Department of Information Sciences, University of Pittsburgh 135, North Bellefield, PA 15260

Jeff@planetjeff.net

(2) CLARTE, 4 Rue de l'Ermitage, 53000 Laval, France.

Lerenard@esiea-ouest.fr

(3) School of Computing, University of Teesside, TS1 3BA Middlesbrough, United Kingdom.

j-l.lugin@tees.ac.uk, m.o.cavazza@tees.ac.uk

## ABSTRACT

We describe the development of the CaveUT system, which is a software supporting immersive virtual reality installations based on the Unreal Tournament game engine. CaveUT implements several high-end VR features such as real-time stereoscopy with head and hand tracking. We demonstrate the use of CaveUT in the SAS Cube™, a PC-based CAVE™-like immersive four-screen display. One of the main advantages of the system is to support full immersive VR while retaining the advanced features of game engines in terms of interaction and inclusion of behavioural (or AI) systems. We illustrate the use of CaveUT on two installations: an artistic VR installation and an immersive interactive storytelling system.

## Categories and Subject Descriptors

H5.1 [Multimedia Information Systems] Artificial, Augmented and Virtual Reality - Virtual Reality for Art and Entertainment.

## General Terms

Design and Experimentation.

## Keywords

Game engine, Immersive Displays, Digital Arts, and Intelligent Virtual Environments.

## 1. INTRODUCTION

Sensory immersion combined with real-time interaction (Virtual Reality, or "VR") has always had great promise for innovative game design, but application development software and libraries for VR research are intended for very diverse applications. They generally do not provide the animation support, optimized graphics and real-time Physics that most game engines incorporate. CaveUT<sup>1</sup> solves this problem by adding a VR interface to the Unreal Engine<sup>2</sup>, which preserves the engine's built-in advantages, allows for the re-use of existing game content, and allows for the creation of new content using standard methods. CaveUT, which follows the principles described in the original CAVE™ system [6], supports a variety of immersive display strategies, from low-tech to fully stereoscopic multi-screen display, which we describe in this paper. CaveUT is part of a general trend among researcher to take advantage of new technologies developed by the game industry [11], including support for immersive displays. Large, single-screen displays offer simplicity and compatibility with a variety of game engines, and the most advanced provide stereographic imaging with tracking<sup>3</sup> (see Figure 1). Several approaches to the immersive visualization of game engines output have been described. Graphics "tiling" software allows games engines to display on very large composite screen displays<sup>4</sup>. With specialized graphics drivers, and the correct settings, a game engine can display in the new all-digital dome displays<sup>5</sup>. A game engine can display in a CAVE™-like [6] enclosure through direct modification of the engine itself<sup>6</sup> with full tracking and stereographic imaging. Finally, a game engine can display in CAVE™-like enclosures

---

<sup>1</sup> CaveUT. (2004). CaveUT, <http://planetjeff.net/ut/CaveUT.html>

<sup>2</sup> EpicGames.(2004).UnrealTournament, <http://www.unrealtournament.com>

<sup>3</sup> Vizbox. (2004). Vizbox, Inc., <http://www.visbox.com/>

<sup>4</sup> Chromium.(2004).Chromimum,<http://chromium.sourceforge.net/>

<sup>5</sup> DomeUT. (2004). DomeUT, <http://planetjeff.net/#DomeUT>

<sup>6</sup> CaveQuake.(2004).CaveQuakeII, <http://brighton.ncsa.uiuc.edu/~prajlich/caveQuake/>, VRizer. (2004). VRizer, <http://futurelab.aec.at/vrizer/>

through the combination of modified graphics drivers<sup>7</sup> and top-level code to synchronize the views. The latter approach has been adopted for the development of CaveUT. It avoids many licensing and distribution issues, and allows straightforward upgrades when new versions of the game engine are released.

CaveUT 2004<sup>(TM)</sup> is a set of open-source freeware modifications, which allows the player to interact with Unreal Tournament, where s/he sees a unified view across multiple screens which can be in any orientation to the user. It has been available to the public since 2001 ([7], [8], [9]) and is detailed at its online distribution site<sup>1</sup>. The latest version is CaveUT 2004 which takes advantage of the latest release of Unreal Tournament and provides better synchronization between screens.

Recently, in collaboration with members of the ALTERNE project [3], CaveUT has been extended to support stereoscopic display and real-time tracking of head and hand position. These capabilities have been incorporated into CaveUT 2003 and we will soon add them to CaveUT 2004. CaveUT 2003 v2.0 is available by request (jeff@planetjeff.net) and will be available as part of the Alterne<sup>TM</sup> software platform for VR Art<sup>8</sup> [3]. In further sections, we will illustrate the use of CaveUT through some of the artistic VR installations developed as part of the ALTERNE project (figure 2).



Figure 1: A Immersive Display: The SAS-CUBE<sup>TM</sup>

## 2. A CaveUT Primer

The development of CaveUT was made possible by the fact that Unreal Tournament (UT) is partially open-source. UT uses the proprietary Unreal Engine, which handles graphics rendering, animation, physics, networking and a byte-code interpreter which supports Unreal Script, a Java-like programming language. Much of the game, itself, is written in Unreal Script, and all of it is open source. A large community of players, game designers and researchers constantly produce new, open-source code and content, and CaveUT initially followed this trend. CaveUT is a package of original code written in Unreal Script with extensive documentation for its proper use. The original (core) CaveUT

supported only monoscopic imaging, no rendering synchronization and no tracking capability. However, it can still be useful to develop low-cost immersive displays, straightforward to set up, and is currently employed in several research installations.

A multi-screen display based on CaveUT requires a server computer connected by a standard LAN to a number of client computers, at least one for each screen in the display. The operator begins a multiplayer game of Unreal Tournament with one normal player on the server and one “spectator” player on each of the clients. Each spectator initially duplicates the view seen by the player on the server. On each client, the CaveUT code rotates the view according to parameters defined in a configuration file, so that each screen is showing the part of the composite view it is supposed to. For example, the view in the SAS-Cube<sup>TM</sup> shown in figure 2 is produced in part by having one client “look” forward, one look ninety-degrees left, another look ninety-degrees right and the last one look down. CaveUT preserves these rotations relative to the server player’s view, so the operator can navigate using standard game controls attached to the server.

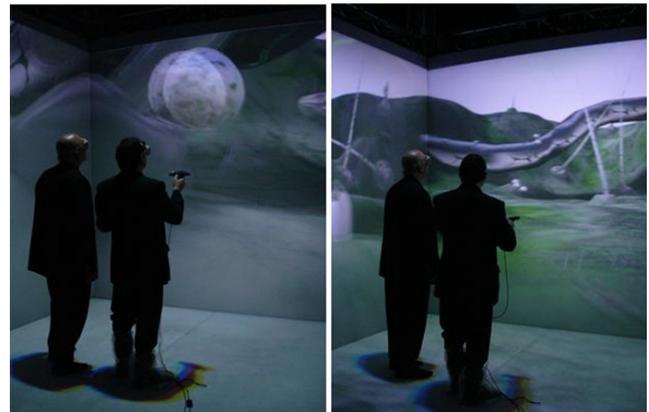


Figure 2: Immersive Visualisation In the SAS-CUBE<sup>TM</sup>

However, the perspective correction on each screen must be adjusted so that the ideal viewing point for each screen is located at the same point in physical space. For a static (not using any tracker) CaveUT installation, this creates single, ideal, viewing location for the whole display. As long as the player keeps her head at or very near this point, she will see a unified and undistorted view. Interestingly, the screens can rest in any orientation to the viewer and they do not even have to be contiguous. For example, CaveUT could be used for a driving simulator, where CaveUT provides the view for the front and back windshields, each side window and the rear-view mirror. The perspective correction is introduced by a modified OpenGL wrapper library, called VRGL (by Willem de Jonge), which rests “between” OpenGL and the Unreal engine. VRGL and CaveUT share the same configuration file, which provides the parameters needed for the perspective correction. Unlike the CaveUT code, VRGL is an independent package and could be used for other applications (for a more detailed explanation, see the CaveUT distribution site<sup>1</sup>).

<sup>7</sup> VRGL. (2004). VRGL, <http://planetjeff.net/ut/CUTVRGL.html>

<sup>8</sup> <http://www.alterne.info>

### 3. Rendering Synchronization

In CaveUT 1.0, rendering on the component display screens was not synchronized. This is not a problem for monoscopic viewing if all of the clients are able to render at thirty frames per second or faster, but that is not always going to be possible, even with identical machines. Load will fall unevenly on each of the client computers, depending on which part of the scene it is rendering. CaveUT 2.0 solves this problem with the addition of a simple swaplock server, which runs on the server computer in parallel with the UT game server. The server starts by broadcasting a “ready” message to the client computers. The signal instructs the VRGL on each computer to wait for a “render” message before displaying the current rendered frame. Each computer then sends the swaplock server a “ready” signal. The server will wait until it has received a signal from all of the clients, before sending them the “render” signal. This insures that all screens will render at the same time and at the speed of the slowest client. Although this could in theory decrease the rendering rate below acceptability thresholds, in practice this has never been observed in the various implementations with which we have experimented.

### 4. Tracking

CaveUT now supports real-time tracking in physical space, using the Intersense™ IS900 system<sup>9</sup> or any similar devices. Attaching sensors to the user's head or a hand-held controller creates many opportunities for user interactions with the virtual environment. Tracking the player's head allows CaveUT to generate a stable view of the virtual world, while the player is free to move around inside the display (which has the size of a traditional CAVE™). This provides automatic adjustment for the user's height, allows her to use parallax effects for better depth perception of the scene, and permits a wider range of interactions with the application. Hand-tracking allows the UT-based application to be aware of the location of user's hand or some hand-held controller at all times. This allows much more natural control designs, such as being able to physically interact with a virtual object by simply pointing the controller at it and pressing one of the control buttons. This potentially supports a wide range of tracker-based interactions (for a detailed taxonomy of interaction with VR applications and tracked controls, see [1]).

From a system integration perspective, CaveUT uses another freeware package, Virtual Reality Peripheral Network (released by the Department of Computer at the University of North Carolina at Chapel Hill) to handle input from all control peripherals such as joysticks, buttons, gamepads and the tracking system itself. All controllers are physically attached to the server machine, and data from the peripherals are collected by the VRPN server, which runs in parallel to the UT game server. The VRPN server converts data from the control peripherals into a generic normalized form and sends it to the CaveUT code in the UT game server, via a UDP port. The modified UT game server uses this information to update the user's location in the virtual world from the head tracker and to process commands from the other control peripherals.

With respect to the head tracker, this means that when the user takes a step in physical space, the server moves his corresponding viewpoint (and avatar) in virtual space. The VRPN server also

broadcasts the user's new location to each one of the UT clients, and the information is received by a VRPN client. Then, the VRPN client sends the tracking information via another UDP port to the VRGL code attached to the UT Client. VRGL uses this information to adjust the perspective correction, in real-time, to preserve the perspective depth illusion. The overall result is that the user's view into the virtual world looks stable to him and the correspondence between the virtual world and the real one is maintained.

Data from the hand-held controller can be used to select and manipulate objects in virtual space. For example, the player could select and “grab” an object by pointing the controller at it pressing a button on the controller. This supports various forms of object interaction, such as the triggering of object behavior or the real-time manipulation of the object position and orientation. Tracking is also a prerequisite for binocular display using active stereo imaging in immersive virtual reality systems.

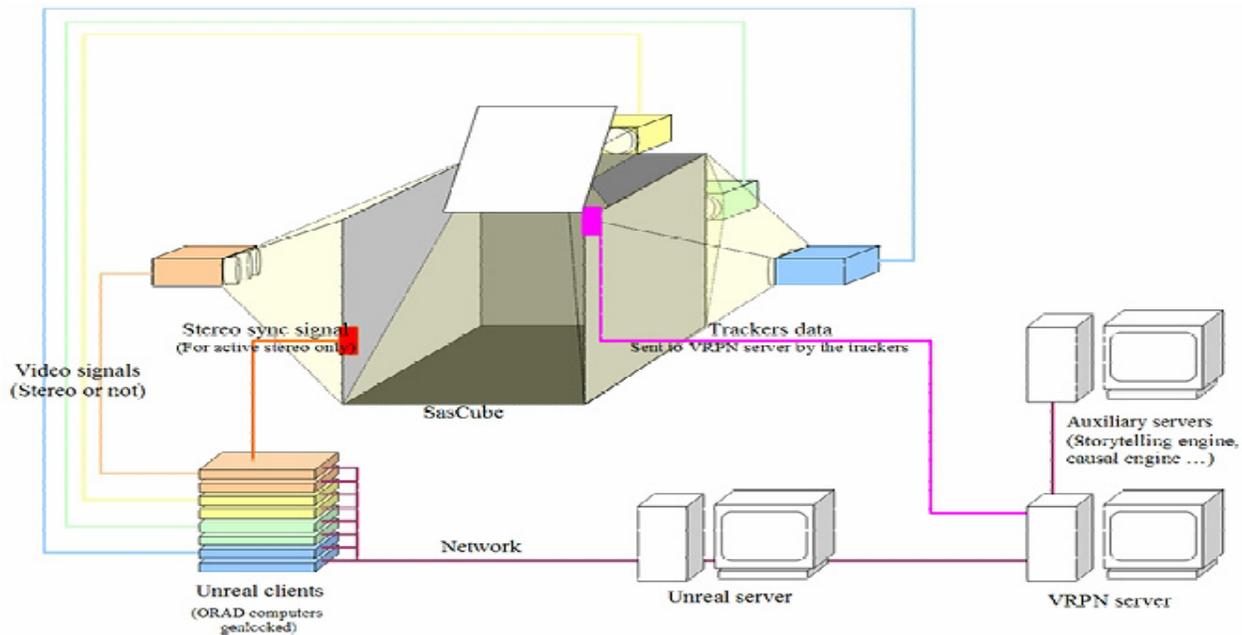
### 5. Stereographic Display

CaveUT 2.0 supports stereographic display by using two computers per screen, one to render the left eye view and one to render the right eye view, with an average frame rate of 60 frames/sec per eye in most experiments reported here. The camera view can be offset from the viewer's default configuration by a set value equal to half the inter-pupillary distance. If no tracking is used, then the separation is horizontal (left and right) and the illusion will hold as long as the player keeps his or her head level, which makes floor or ceiling screens impractical [10]; however, with tracking, the view for each eye remains centered on the designated eye, regardless of head position and orientation. This is the solution implemented for use in the SAS Cube™, which includes a floor screen. For active stereo, (as well as for color-based stereo) this allows the user to look in any direction and at any angle. CaveUT supports various approaches to stereoscopic display, including passive and active stereoscopy. For passive stereo with two projectors per screen, a linear polarizing filter should be placed in front of each projector lens and in front of each of the viewer's eyes. The orientations of the filters on the projector lenses should differ by ninety-degrees. As long as the player keeps his or her head level, the illusion is excellent and the colors are vivid. This approach requires that the projectors do not produce light which is already polarized as some LCD projectors do. This should not be a problem with DLP projectors. For color-based stereo, the left and right projectors should be set to produce only red and green respectively, and red-green filter glasses are to be worn by the player. This approach is almost never used, because the resulting image is in an unattractive monochrome, but it is easy to implement and robust.

Active stereo requires a single stereographic projector which will alternate between the left and right eye views at 120 frames per second. The player wears “shutter glasses”, on where each lens alternate between black and clear, also at 120 frames per second. The glasses switch in time with the display, and the result is that each eye gets the view it is supposed to at 60fps—the left view for the left eye and the right view for the right eye. All of the screens in the composite display must also switch view at exactly the same time, a desirable state called “genlock”.

The original CaveUT system does not have software support for generating an active stereographic image from a single machine.

<sup>9</sup> Intersense. (2004). <http://www.isense.com/products/prec/is900>



**Figure 3: CaveUT 2.0 Architecture**

The CaveUT 2.0 installation in the ALTERNE SAS-Cube™ platform uses two computers for each screen, one for each eye view, and uses the DVG (video) cards in their ORAD<sup>10</sup> (PC) cluster to mix the two and send it to a single stereographic projector. The DVG cards also handle the genlock synchronization across all screens of the composite display (see figure 3).

## 6. Interaction Design With CaveUT

The interaction design for Unreal Tournament and most other computer games assumes a standard keyboard, mouse and monitor interface. However, optimal interaction designs for immersive displays are quite different. While CaveUT can be used with almost any preexisting game written for Unreal Tournament, important features of the game's interaction design may no longer make sense in an immersive display, especially if the player is standing up, using a completely hand-held control such as a gamepad. Unlike a desktop display, an immersive display must present a view of virtual world which matches the real world in scale. This means that the player must perceive a meter in the virtual world to be the same as a meter in the real world, as though could simply walk through the display into the virtual environment.

CaveUT handles these issues automatically, but the resulting illusion may make the virtual world look and feel larger or smaller than the designer wants it to be. The game designer can solve most of the problem by applying a scale factor to the virtual environment display and to the collision handling. Finally, another important self-location cue in a visualization engine derived from a first-person-shooter game, is the location of the user's "hand", which can be derived from the display of the players' weapon in the original game. This is important whenever the user has to interact precisely with stereo objects at short

distance. The best solution in our experience consists in associating such position to the tracker-equipped gamepad held by the user.

With CaveUT, the player can use any game peripheral (joystick, gamepad, etc.) or more advanced versions of these devices, equipped with trackers, which is the solution we adopted. There are many complex issues related to controls in immersive environments [1]. Two issues stand out. First, most FPS game players move very quickly in the virtual environment, which is disorienting in an immersive display and does not correspond to any reasonable walking speed (notwithstanding the fact that it can generate cybersickness). The "gain" on the controls must be carefully adjusted to match the experience intended for the current design (immersive gaming, VR art, etc.). Second, most FPS games require the user to center his view on an object to shoot or select it, which requires rotating the view quickly and often. In an immersive display this can be disorienting and fails to take full advantage of the wide field of view afforded by the display. Ideally, the user should be able to select something by simply pointing at it. Generally, immersive game design must take advantage of the wide field of view afforded by the display to be worthwhile. If everything interesting is always at the center of the display, then the player gets no benefit from the display.

## 7. CaveUT in action: Artistic Installations

The stereoscopic version of CaveUT described in this paper has been developed in the course of the ALTERNE project, which is a VR Art project [3] aiming at developing a re-usable platform for the creation of VR Art installation. The main objective of ALTERNE is to define a technological platform for the design of "alternative reality" environments, i.e. virtual worlds whose fundamental behavior can be entirely re-defined in terms, for instance, of laws of Physics. This objective brings specific demands on the visualization engine that should support the system. Game engines, and in particular the Unreal Tournament 2003™ engine, appear as a natural choice, as they include

<sup>10</sup> <http://www.orad.tv>

sophisticated event-based systems, which serve as an integration layer for the re-definition of environment behavior. The ALTERNE software platform is thus organized around three main components: i) the UT 2003™ engine for visualization and basic object interaction mechanisms, ii) CaveUT to support immersive stereoscopic visualization and tracking and iii) the “alternative reality” engine, which contains qualitative simulation systems overriding the native Physics engine for specific object categories. ALTERNE installations have provided the first test beds for the stereoscopic version of CaveUT described in this paper. Several artistic briefs have been implemented, which define immersive virtual worlds in which the user experiences novel forms of interaction with the environment.

The first example illustrating the use of CaveUT is an artistic installation developed as part of the ALTERNE project, “Ego.Geo.Graphies”, by Alok Nandi [3]. In this installation, the user navigates in an organic world populated by spheres which originate in determinate areas of the environment. The spheres’ behaviour depends on the perceived “empathy” of the user, which is a function of her navigation patterns, unknown to her. This behavior manifests itself essentially through the effects that follow collision between spheres, which range from soft sphere merging to explosions propagating to the environment. These effects are under the control of the alternative reality engine, which intercepts collision events and computes alternative forms of causality.

This brief makes use of most of the features supported by CaveUT, from tracking and object interaction to stereoscopic visualization. User navigation brings her in close vicinity to geometrical structures which acquire their full dimension as real stereo 3D objects, prompting the user to adopt appropriate navigation patterns around or under such objects. The spheres themselves can traverse the SAS Cube™ volume as floating 3D objects, conferring a high level of realism to the user interaction. In addition, the ultrasonic tracking implemented in CaveUT supports direct physical interaction with the spheres through the SAS Cube™ gamepad, which can be attracted or pushed back by the user (Figure 4).



Figure 4: "Ego.geo.Graphies" in The SAS-CUBE™

More recently, we have investigated how work developed with UT 2003 could be ported to an immersive context using CaveUT. Interactive storytelling is one of the applications that epitomize what future entertainment systems could look like. In particular, the “Holodeck™” system popularized by the Star Trek series has become a model for research in future immersive interactive storytelling system [12] [13]. Such a system is characterized by the full immersion of the user in a 3D stage, populated by virtual actors, in which the plot unfolds around the user herself. It has thus seemed a natural extension of our research in Interactive Storytelling [2] to adapt it to a fully immersive platform to explore the “Holodeck™” concept. Our storytelling system is developed on top of the Unreal Tournament engine, which made CaveUT a natural choice to port it to a fully immersive context. In the next sections, we briefly summarize the main features of the storytelling system, and discuss how the immersive implementation affects basic interactive storytelling concepts. We then comment upon the technical issues which had to be solved in the context of our CaveUT implementation.

We have named our approach “character-based interactive storytelling” [2] to reflect the specific stance taken with respect to relations between characters and plot. The baseline plot for the interactive narrative (in the example supporting our experiment, the plot consists in a Sitcom-like episode about a group of friends organizing a party) is projected onto individual roles for the virtual actors, formalized as HTN plans. Each virtual actor will thus act independently according to its baseline role in the virtual world. The dynamic interaction between actors is the key principle behind the generation of multiple narratives from a basic storyline. The technical basis for dynamic story generation is that each actor’s role is formalized as a plan to be executed in the virtual stage; several actors will be competing for resources shared in the same environment, these resources being objects or other actors. These conflicts for resources result in plan failure and re-planning, hence creating humorous situations and driving the narrative forward. The same mechanisms support various forms of user intervention in the narrative, which will affect the characters’ actions by failing or modifying some of their goals. Detailed technical descriptions can be found in [5].

An immersive implementation affects the paradigm for user involvement in the interactive narrative. We have originally developed our system for an “interactive TV” philosophy in which the user would influence the story from a god-mode perspective rather than as an actor. We then explored a mixed-reality approach in which the user was simultaneously actor and spectator [4]. In the “Holodeck™” paradigm implemented here through CaveUT (figure 5), the user interacts from a first-person perspective, as a member of the cast. The stereoscopic display confers an increased feeling of realism, as 3D actors traverse the SAS Cube™ space, avoiding the user whose tracked position generates a bounding box. Full immersion and hand/wand tracking supports improved physical intervention on the virtual stage, i.e. the user can influence the story unfolding by removing/hiding key narrative objects (in addition to other forms of interaction, such as influencing actors through speech recognition).



Figure 5: The “Holodeck™” in The SAS-CUBE™

## 8. Conclusions

Through the implementation of CaveUT, we have shown that the advanced features provided by game engines for both visualization and interaction could be adapted to the context and specific requirements of immersive virtual environments. This has significant implications for entertainment technologies, as it opens the way for the exploration of gameplay in VR and perhaps the search for game genres better suited for the VR setting. One of these new genres is undoubtedly interactive storytelling, whose immersive form endeavors to implement the Holodeck concept. Existing forms of digital entertainment such as VR Art, whose installations used to be developed through custom-made, high-cost, VR systems, can greatly benefit from this approach, which combines sophisticated features within greater accessibility of the software platform to developers. Finally, immersive VR based on game engines also constitutes yet another example of the increasing use of game technologies in non-gaming applications.

## 9. ACKNOWLEDGMENTS

Part of the work presented here has been funded by the European Commission through the ALTERNE (IST-38575) Project. The “Ego.Geo Graphics” installation has been authored by Alok Nandi as part of the same project. The authors acknowledge the assistance of Epic Ltd through “PublicVR”. The opinions expressed in this article are the authors’ only and this does not imply any endorsement by Epic.

## 10. REFERENCES

- [1] Bowman, D. A. Principles for the Design of Performance-oriented Interaction Techniques. In K. M. Stanney (Ed.), *Handbook of Virtual Environments*, Mahwa, New Jersey: Lawrence Erlbaum Associates, Inc., Publishers, 2002.
- [2] Cavazza, M., Charles, F. and Mead, S.J.. Character-based Interactive Storytelling. *IEEE Intelligent Systems*, 17, 4, 2002, 17-24.
- [3] Cavazza, M., Lugin, J. L., Hartly, S., Libardi, P., Barnes, M. J., LeBras, M., Le Renard, M., Bec, L., Nandi, A. New Ways of Worldmaking: the Alterne Platform for VR Art. *ACM Multimedia 2004*, New York, USA, 2004.
- [4] Charles F., Martin O., Cavazza M., Mead S.J., Nandi A. and Marichal X. Compelling Experiences in Mixed Reality Interactive Storytelling. *International Conference on Advances in Computer Entertainment Technology, ACE 2004*, ACM Press, Singapore, 2004, 32-41.
- [5] Charles, F. and Cavazza, M. Exploring the Scalability of Character-based Storytelling. *Third ACM Joint Conference on Autonomous Agents and Multi-Agent Systems*, ACM Press, New York, USA, 2004, 872-879.
- [6] Cruz-Neira, C., Sandin, D.J. and DeFanti, T.A. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. *Proceedings of the ACM SIGGRAPH 1993 Conference*, 1993, 135-142.
- [7] Jacobson, J. Configuring Multiscreen Immersive Displays With Existing Computer Equipment. *Human Factors and Ergonomics Society 46th Annual Meeting*, Baltimore, 2002.
- [8] Jacobson, J. Using CaveUT to Build Immersive Displays With the Unreal Tournament Engine and a PC Cluster, *ACM SIGGRAPH 2003 Symposium on Interactive 3D Graphics*, Monterey, California, (April 2003).
- [9] Jacobson, J., and Hwang, Z. Unreal Tournament for Immersive Interactive Theater, *Communications of the ACM*, 45, 2002, 39-42.
- [10] Jacobson, J., Redfern, M. S., Furman, J. M., Whiney, L. W., Sparto, P. J., Wilson, J. B., Hodges, L. F.. *Balance NAVE: A Virtual Reality Facility for Research and Rehabilitation of Balance Disorders*. ACM Virtual Reality Software and Technology (VRST), Banff, Canada, 2001.
- [11] Lewis, M., and Jacobson, J. Game Engines In Scientific Research. *Communications of the ACM*, 45, 2002, 27-31.
- [12] Murray, J.H. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*, MIT Press, Cambridge, 1997.
- [13] Swartout, W., Hill, R., Gratch, J., Johnson, W.L., Kyriakakis, C., LaBore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thiebaut, M., Tuch, L., Whitney, R., Douglas, J. Toward the Holodeck: Integrating graphics, sound, character and story. *Proceedings 5th International Conference on Autonomous Agents (Agents-01)*, ACM Press, New York, 2001, 40