THE JOURNAL OF VISUALIZATION AND COMPUTER ANIMATION
J. Visual. Comput. Animat. 2001; **12:** 203–214 **(DOI:** 10.1002 / vis.262)

The Caracol Time Travel Project

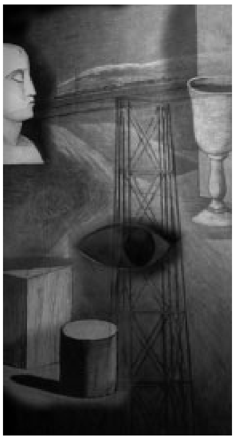

By Charles E. Hughes*, J. Michael Moshell, Dean Reed, Diane Z. Chase and Arlen F. Chase

Virtual drama is based on the use of a shareable virtual world as a stage setting, with avatars controlled by actors and audience members. The Caracol Time Travel Project was an experiment in the use of virtual drama for learning about archaeology. Eighteen undergraduate students at the University of Central Florida used a locally developed Java-based system for sharing VRML worlds. They designed and constructed a virtual drama to teach basic concepts of Mesoamerican archaeology and the cultural history of the ancient Maya for middle schools. This paper presents their story design and details of the system we developed to support interaction in this shared virtual world. We then discuss performance issues, lessons learned and newer features that we did not have available at the time. Copyright © 2001 John Wiley & Sons, Ltd.

Received: 6 June 2001; Revised: 5 July 2001

KEY WORDS: virtual environments; virtual drama; VRML; situated cognition; level-of-detail; streaming media

Introduction

Virtual Drama

Several experiments in virtual drama have been reported in the literature. Often these projects involve interactions between human guests and autonomous characters.[1,2] In the Oz project,[3] experiments were also conducted in which live actors were used. The live actors followed a script, whereas the guests (users of the system) were given only a general sense of the objectives and background of their characters.

The project described here focuses on human–human interaction in virtual worlds. As in the Oz example, we maintain the distinction between 'guests' and 'cast members', but we are not working toward automating the cast members' functions. Rather, we focused on developing the virtual medium as a friendly and convenient place for experimentation with live improvisational drama as an educational tool.

*Correspondence to: Charles E. Hughes, School of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL 32816–2362, USA.
E-mail: ceh@cs.ucf.edu

VE and Education

A number of experiments have been conducted in the use of VE systems for education.[4,5] Only a few of those experiments investigated the educational impact of social interaction within the virtual environment, and these were usually based on the interactions between two explorers pursuing a mutual task (e.g., Johnson *et al.*[6]). Most work on larger-scale social interaction in cyberspace has taken place in the MUD/MOO universe.[7]

The ExploreNet system, originally modeled after the well-known Habitat role-playing environment,[8] has been used for a series of experiments with elementary and middle school students.[9,10] ExploreNet uses simple cartoon scenes and links multiple computers together. Normally, each computer controls one character ('avatar'). Additional non-player controlled characters ('props') populate the world. The behaviors of these characters are determined by avatar interactions and the state of other world objects (avatars, props and scenes). This can make the props appear to exhibit intelligent behavior, but that appearance is only a consequence of the potentially large state space used to select preplanned actions.

While the Caracol Time Travel Project did not actually use the ExploreNet environment, we did employ its drama-based educational framework and the 'Virtual Academy' model[11] developed in conjunction with ExploreNet.

The Virtual Academy Model

Central to the Virtual Academy model is the use of role categories: guests, cast members, world builders and tool builders. Guests are naïve first-time visitors to the virtual world. They have, at most, read some background information and perhaps seen a pre-show. Cast members are persons who have learned enough about the simulated world to play specific roles, in costume (as avatars). Guests encounter them in the world. Cast members need three kinds of skills and knowledge:

- how to move and operate fluently within the virtual world;
- how to act in character, within the role they are assigned; and
- how to move the storyline forward in an interactive drama, so as to facilitate the guests' learning.

World builders define the world's learning objectives, develop scripts and theatrical scenarios, construct the geometric models and program the behaviors of the objects in the world. World builders also serve as the directors, in a theatrical sense, of the actual productions. Tool builders are technically adept persons who modify the software with which world builders work to construct the world.

The ultimate mission was to develop an educational virtual environment suitable for use by middle school students. During the project described in this paper, undergraduate students played all these roles. The 18 participants consisted of 10 computer science majors, three liberal studies majors, two music majors, two graphic design majors and a psychology major. Several anthropology students served as consultants.

Objectives of the Project

The project had as its three main objectives to improve our understanding in the following areas:

- *Interactive virtual drama for learning*: how can scripts be developed which establish the overall direction of an experience, without constraining the actors to a single course of action?
- *World-building*: how can you organize the design and construction of virtual worlds that support improvisational dramatic learning experiences? What do the builders need to know; how should they be organized?
- *Technical infrastructure*: how can virtual worlds be set up to maximize usability and minimize the effects of network bandwidth limitations?

Virtual Drama for Learning

The presentation paradigm of the Caracol Time Travel project is as follows:

- A host school (or some other organization) provides well-trained cast members, and offers to stage performances for other schools with appropriate Internet-accessible computing equipment.
- The guest class reads background material about the Maya civilization before the performance. Two to four members of the class are selected as guests; they serve as the class surrogates in the experience. Each guest sits at a separate computer, with one additional computer connected to a projector.
- The guests participate in three episodes, aggregating about an hour of interaction. Their roles are archaeology student trainees, working on the Caracol Archaeology Project.

A good deal of research on situated cognition[12] has indicated that the cast members—those with primary, predetermined roles to play—will benefit from such an experience. We need to consider the role of the guests. What is their intended role, and what is the benefit of having them in the scenario?

Guests serve three purposes. They provide a visible audience to whom the cast members are playing. In this capacity and with their ability to move freely about and ask questions of all sorts, they present the cast members with many challenges. They also provide the invisible audience (the rest of the class observing the projected image) with a means of interacting with the story. And finally, guests are future cast members in training.

World-Building

The basic hypothesis is that students will learn about the subject matter by constructing a script and a virtual world that conveys this subject matter to others. This project provided an extended opportunity to explore this idea.

World-building must proceed in intimate relationship with script development. It would be desirable to have a mature and stable set of tools for world-building, and a team of designer/constructors who were completely familiar with the tools. It would also be desirable for the students to have substantial prior knowledge of the subject matter. None of this was possible for this project. Nevertheless, the project produced useful virtual worlds, and provided insight into the key questions set out above.

Technical Infrastructure

VRML was the most flexible tool available to this project in the fall of 1998, given that the worlds must be shareable and each object's behaviors must be readily programmable. A system called VRMINet was developed in the course of this project. Users controlled avatars (moving human figures) using a mouse. Each user had a separate computer. The avatars were constructed by several means, including adaptation of avatars found on the Internet. Only limited animation (a walking cycle) was provided for the avatars; their primary role was to indicate the location of the action. Props generally had much richer animated behaviors.

Consideration was given to using other systems such as Blaxxun, AlphaWorlds or Alice. However, at the time of project initiation, neither of the first two provided sufficient user programmability of object behaviors and Alice had no pre-built infrastructure for sharing virtual worlds with multiple interacting avatars.

Learning to Design the World

We began by reading papers about the Caracol project,[13–17] as well as several books about the ancient Maya.[18–20] An essential element of modern Maya storytelling is that the audience actively participates, by asking certain stereotypical questions. The student designers of *Caracol Time Travel* resolved to incorporate this technique into their virtual world.

For an introductory exercise, we used the two-dimensional ExploreNet system[9] to develop and role-play a story about the rain god Chaac. The scenario we built was quite simple, but proved useful to train students to be cast members. We also used it to learn about issues like isolation of avatars from the main scene's action. This arose fairly often in our naïve narrative, leading our students to design a world in which this situation would not be encountered. In particular, it led to the creation of non-player-controlled characters that help keep the troupe of guests together whenever scene transitions are made.

Building the 3D World

As the story developed, we realized that we wanted to have three major episodes. Three (overlapping) teams were assembled, each with a coordinator/scriptwriter who had experience and talent in writing. The spring 1999 semester's additions were 10 computer science majors. We inventoried their skills in art, VRML and Java, and made sure that each of the three teams received at least one student with each skill. Other than the team coordinators, we did not assign any specific roles or titles (such as Technical Director). These roles emerged as the teams began to function. A leadership council consisting of the three team chairs played the essential coordinating role.

We knew that we wanted to personalize the experience of archaeology along two dimensions. We wanted identifiable characters with plausible motivations, doing things real people would do in similar situations. We also wanted to explore both the modern world in which archaeology is taking place, and the ancient world in which the Classic Maya lived.

We also knew that on top of Caana, the largest pyramid at Caracol (Figure 1), a high-status female burial had been found (Figure 2).[14] At other Maya sites, few female Maya burials receive the ceremonial treatment normally afforded to males. The female's name and date of her birth are not yet known (but see refs 20 and 21); her death date is recorded in the tomb as having occurred in AD 634 (ref. 14, p. 27). We decided to incorporate a fictitious version of this lady, named Rainbow, into our story.

The overall storyline was based on time travel. Guests would play the role of student archaeology trainees, coming to Caracol to work on the project. They would interact with characters representing the children of indigenous people such as the camp's cook. The children would reveal magical pathways that led to the mystic Maya past.

Based on the lessons learned while working with Chaac, it was decided to use a 'Wizard of Oz' principle for characters. Each character in the present-time scenario would be matched with a character in the Classic Maya time. They would resemble one another in physical appearance and personality, and would be played by the same cast member. The guests, since

Figure 1. Caana, the largest pyramid at Caracol.

they travel through time, would not need two characters. As the project progressed, we discovered the need to provide for a 'camera avatar' controlled by a cast member on a computer used to render the public view of the performance.

The Caracol Time Travel story incorporates elements of several Maya myths, together with a substantial amount of new dramatic material. The ancient queen Rainbow sacrificed her jade amulet so that Caracol would win the climactic battle with the city and state of Tikal (e.g., ref. 14, pp. 60–62); but this doomed her soul to exile forever. The archaeology students (guests) learn her story while exploring the ruins during three interactive episodes, and reassemble the amulet to free her soul for its journey to the underworld. Figures 3–5 depict some scenes from the adventure.

Figure 2. Burial site at Caana.

The students designed the Caracol world by storyboarding and scripting the action, while examining maps provided by the archaeologists. A large-scale map measuring 16 feet by 18 feet (original scale, 1:1000) was laid out on the floor of the lab using electrical tape, and used for 'blocking'—planning the initial positions of the avatars and of the principal

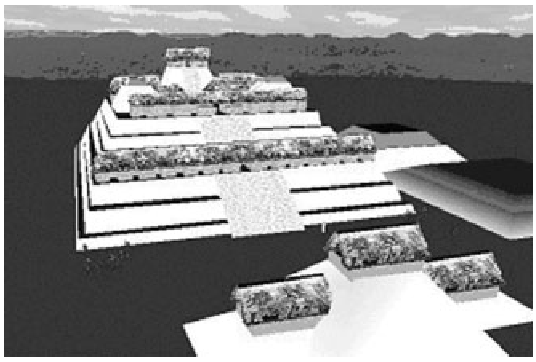

Figure 3. VRML model of Caana.

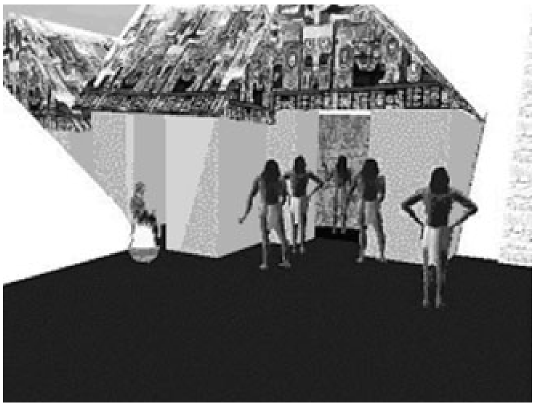

Figure 4. Rainbow's funeral.

activities. The three episodes used different VRML models, built by three different teams.

The sets for the story's three episodes were constructed using Cosmo Worlds, which is a VRML editor developed by Silicon Graphics. This tool is satisfactory for some relatively simple purposes, but has disadvantages when behaviors are attached to VRML objects. In fact, the VRML source produced by Cosmo Worlds generally needed extensive repair to make it useable with the External Authoring Interface that provides communication from Java to the VRML scene graph and browser.

Photographic texture maps were used to establish the overall sense of place. Since several project members had visited the actual archaeological site, the virtual world's evocation of the Belizean jungle could be verified.

To keep the computer rendering speed at acceptable levels, the VRML worlds had to be as small as possible. Consequently, the three episodes' worlds were disjoint.

Figure 5. Examining a map in the cave.

The episodes were separated by intermissions, as in a theatrical production, to allow for changing of the scenery.

Models

Members of several teams contributed to the construction of human figure avatars. Some of the figures were built using Sven Technology's Avatar Maker (a very nice, but discontinued product from a defunct company); others were harvested around the Internet and customized for this project. Photographs were taken in Guatemala of nine-year-old schoolchildren (with the teacher's and parents' permissions) and used to provide faces for a boy and girl avatar.

Episode One's world was based on a high-detail model of Caracol's largest structure, the pyramid called Caana ('Sky house'). This model was initially developed from AutoCAD data that was imported into Arc View, a geographic information system. In modern Caana, all the roofs of the structures on Caana have collapsed. To construct the ancient Caana, the team added roofs. The exterior surface of the ancient pyramid was shown as being covered in white plaster, as it actually was. Atop Caana, in the actual location of one of the tombs found there (below the left side of the north temple structure), the artists located Rainbow's tomb (see Figures 1 and 3). Episode Two was sited in a cave. This was done because one of the students constructed a visually compelling cave, and everyone liked it. While caves do occur at Caracol,[22] they are not very large. Caracol exhibits many more tombs than caves.[23] The cave in the virtual world is, however, much bigger than any of Caracol's tombs. This portion of the story was largely developed before the field trip to Caracol, and represents the largest divergence (in terms of 'set design') from physical reality. In the actual performance, the role played by the cave is museum-like, in that the characters move from artifact to artifact and discuss them.

Complexity of Scenes and System Performance

The complexity of the avatars used in Caracol Time Travel ranged between 1000 and 5000 polygons. Performance was substantially enhanced when low-polygon count models represented most of the characters in a scene.

The computers used were 450 MHz Pentium II systems equipped with 128 MB of RAM and Accel Galaxy graphics accelerators. These systems were capable of rendering the VRML worlds at between 1 and 10 frames per second, depending on the specifics of a world, the avatars in use and the viewpoint and viewing direction. Table 1 represents the world' complexity (including all props).

In general, system performance was not a bottleneck for the dramatic purposes we were pursuing. However, during development the lab computers were in constant demand since most students' home machines could not manage the large VRML databases. This situation would not arise today, given the advances in consumer-grade graphics accelerators. In fact, when we reran this system on a 1 GHz Pentium III equipped with 256 MB of RAM and a 64 MB DDR Nvidia GeForce2 GTS AGP, the performance was over 40 frames per second using the models described above. This performance makes movement much smoother and would have made the experience more enjoyable, but we were limited by the technology available at the time of these experiments.

Behaviors

In Episode One, an animated toucan played an important role. The toucan existed in the modern timeframe and also in the past. It perched on top of the stela. Once the guests touched the Caracol glyph on the stela, the camera's viewpoint was bound to an over-the-shoulder view of the toucan. With an 'Awwk!' cry, the toucan flew upward, following the steps on the front of Caana, to lead the viewers to the site where Rainbow was being buried. The toucan's flight-path was automated so as to provide for precise control of camera position and direction during the flight. It was found that using a live camera operator for this rapid flying scene was very difficult, and usually resulted in a disorienting experience for the viewers.

The principal technical problems to be solved for Episode Two concerned the task of zooming in on an object and viewing it from all sides. The main motion control paradigm that is provided by the Cosmo VRML browser is based on using the mouse. The user presses the button, and then drags the mouse upward to move forward, downward to move back ('zooming'), or to the left and right to rotate the viewpoint about a vertical axis ('panning') in the desired direction. The four cursor keys perform the same functions, and are often used instead of the mouse. This technique could be referred to as a 'unicycle' steering paradigm.

Even if the dollying function were convenient, it would require a coordinated dolly and pan, to go around an artifact and see it from all sides. This proved too challenging for our camera operators. We experimented with pre-programmed camera paths that circumnavigated the artifact, as well as with animated actions that allow the artifact to be rotated about its vertical axis. The latter technique proved somewhat preferable, since circumnavigating the object produces a complex background field motion that distracts the viewer from the artifact itself.

Several nicely textured artifacts were developed for this episode. A Maya vase with a detailed painting of the ceremonial ball game was constructed, and made part of the narrative. However, it proved difficult to actually see and comprehend the figures on the cylindrical surface of the vase object in VRML. This was partially due to the cylinder's curvature and partially due to the difficulty of zooming in on an object in VRML, using the Cosmo control system. It might have been better to include a hyperlink so that the vase's image was delivered on the screen as a flat, unrolled image.

Episode Three was the most ambitious episode, in terms of behaviors programmed into the world. A shovel could be selected from a menu, and was seen to make digging motions. An excavation pit was provided, the bottom of which moved downward as excavation continued. At a succession of depths, each of four fragments of the artifact (amulet) was revealed. The fragments had to be placed into the sifter in order to clean

Episode	Polygons	VRML files	VRML Mbytes	Texture files	Texture Mbytes
1	6085	26	4.15	47	1.04
2	4854	42	3.73	109	8.22
3	7393	135	5.76	119	6.02

Table 1. Scene complexities

them, at which point the jade green amulet fragments became visible. When the fragments were placed close together, they moved to form a single amulet.

Lessons Learned

Direction. Faculty must be ready to exert leadership on occasions when the student leaders' inexperience is seriously impeding the project's development. But too much intervention stifles creativity. Throughout the fall semester, students developed many workable ideas for elements of the script, but no unifying treatment emerged. Ultimately, one of the faculty members took the scattered pieces and wrote a two-page story treatment that established the key character of Rainbow, the idea that she sacrificed a jade amulet to save her people, and that the archaeological project in the modern era could complete the pattern of Rainbow's life by restoring the amulet. Once this thread was established, the students proved quite capable of embellishing the theme (and in fact of changing essential parts of it).

Mixed Teams. Initially we were concerned that the programmers might be headstrong and would only provide lip-service to the humanistic students' goals and scripts. This did not occur; in fact the humanists were intimately involved in building the worlds, and the programmers tried very hard to meet their specifications. The technical students taught some of the humanists how to construct VRML models, and incorporated some of their modeling products into the scenes.

The team leaders developed skills in understanding their members' capabilities, and learned how to motivate their members to meet deadlines.

Bottlenecks. A technical lead emerged in each group, and became a bottleneck. Even though several students might have VRML or Java skills, one person was usually trying to integrate the episode's components. We would often hear team leaders saying 'As soon as Henry ...', referring to the fact that the process was stalled because the technical lead was overloaded.

Role-Playing. A regular scheduled rotating rehearsal of episodes is a very good model. Each team knew that a week or 10 days in the future it would be their turn to present an improvisational session. They would schedule their work to achieve a functional partially completed world at that point, would meet outside class to assign roles and refine the script, and would be ready to go when class began. Teams learned a great deal by participating in other teams' rehearsals. The feed-forward of techniques was often obvious in the next round.

Knowledge. Cast members should be held to an objective standard of knowledge about the subject matter. Guests quickly discovered that an important part of their task was to probe the cast members' knowledge and flexibility, by asking novel questions or wandering away from the expected direction of the performance. The cast members playing authority figures, when asked a novel question such as 'How many people lived on Caana?', would sometimes make wild guesses, so as to keep the enactment going.

Our recommendation is that in the future all members of the enactment teams be given a standardized set of readings, required to study them, and then be given written examinations before they are 'licensed' to serve as cast members and sources of information for guests.

Character Development. Staying in character is difficult when content is dominant. Cast members knew whom they were playing, which characters were supposed to be children and which adults. However, the cast members were focused on conveying the facts about the archaeology site and the back-story, and so they often spoke as adults even when playing children's roles.

Model Building. Make sure your team includes 3D modelers with the right kind of experience and interests. If not, it can take a long time to learn the needed skills. There was only one trained graphic designer on the team, and he was in high demand. Skills taught in animation courses are only partially relevant to low-polygon situations such as this. As a consequence, we now offer a course on modeling for real-time applications. Students who complete this course find themselves in great demand, both for school projects and paid jobs in the simulation community surrounding our campus.

Behaviors. Keep the objects' behaviors as simple as possible. Programming complex interactions between objects was labor-intensive, but was generally successful in achieving the desired behavior.

Real-World Experience. Time spent at the site on which a world is based is very helpful for building compelling virtual models. Before March of 1999, the project's only sources of information about Caracol's geography were maps, photos and two videotapes,[24] provided by the archaeologists. Visiting the virtual worlds built during this period felt like visiting a cardboard model. One's attention is focused on the structures (pyramids and the ruins on top of them) because there was nothing else in the virtual world except terrain. However, when you are at present-day Caracol in Belize, the jungle (rather than the archaeological ruins) is the overwhelming presence. Usually you see jungle, with the occasional glimpse of structures (which mostly look like steep hills of various sizes, within the jungle).

In virtual Caracol, the jungle is primarily represented by a panoramic backdrop that is comprised of photographic images of the jungle, as seen from the archaeology camp. Several layered images provide motion parallax. The backdrop, together with a carefully constructed blue sky with clouds and the principal structures, is quite satisfactory in conveying the overall sense of the environment when the viewpoint is on the ground plane. However, since there are only a few foreground trees, the experience of being atop most of the virtual structures does not resemble the real experience. The exception is Caana, the large pyramid. In this case, one can look downward into the completely cleared courtyard—here the virtual world closely resembles the actual scene in Belize.

Design Rules. Modelers should be given a set of design rules. These would include standard ways of modeling human figures; scale information so that furniture, figures and terrain are of commensurate size; and an up-and-running version control system for archiving work. Fortunately, no important part of the work was lost during the development process. However, the integration team often had to make scaling corrections when including new models in a scenario.

Existing Data. An extensive model of modern Caana existed in AutoCAD format. We cleaned up this model using ArcInfo software and exported it to VRML. This model served as the starting point for the textured modern Caana model, as well as for the geometrically enhanced ancient model. Roofs were placed on the structures, temples were restored based on best available information, and exterior surfaces were colored with white plaster, red plaster (for the apartment and temple complexes) or colored texture maps of Maya art, to represent conjectured images painted on temple roofs.

Non-Participant Viewers. One of our goals for this project was to develop a system that allows non-participants to observe the world as a live audience. We originally thought this would be easily accomplished by having an avatar that acted as the audience's surrogate. In effect, the avatar functions like a camera.

During the first few trial runs of the integrated Caracol Time Travel Project it was discovered that an actor/operator had to look over his shoulder at the main projector screen to determine his avatar's position with respect to the audience view. If a virtual actor needed to position his avatar in the main view of the audience for a dialog, the only way to guarantee that the avatar was visible to the audience was by using the audience view as a frame of reference for position.

A programmer and one of the modelers devised a clever solution to this problem. A virtual studio camera complete with a lens and a red blinking light represents the actual viewpoint of the audience view screen. Using this solution, operators could position their avatars with respect to the audience by looking at the camera avatar. In effect, a virtual actor interacts with the virtual camera just as a real actor would interact with a real camera. An added benefit to the virtual camera avatar is that it removes the need for the actors and the audience to be in the same room. Movements can be synchronized to the audience view by the virtual camera avatar, and voice communication can be handled through Internet telephony. Additionally, the virtual camera avatar added to the feel of being on a virtual set in a theatrical production.

Technology. Your project will only succeed if the technology works reliably and doesn't cost you too much effort. Trying to make something as complex as interactive drama work properly is daunting under any circumstances. To have done so with a substrate that was actively under development was a high-risk venture, but at the time we felt that we had no viable alternative that provided the necessary flexibility in programming the behavior of props (necessary objects in the virtual world).

Our approach of using three teams of programmers, who were effectively in competition to produce the best episode, worked reasonably well. They produced interactive features as called for by their scripts, and then provided these scripts for integration into the distributed system. On most occasions, we were able to get the results to work in time for the next rehearsal.

The major breakdown of the system occurred at the climactic performance, when Episode Three's code would not run—although a 95% complete version had been demonstrated two days previously. Ultimately this problem was traced to an incompatibility between Windows 98 and Windows NT—a texture that loaded flawlessly on Win98 failed to load on NT! While this was unexpected, the Episode Three design team bore some responsibility because of the rapid rate at which they were adding features, right up to the last minute. They simply overwhelmed our ability to integrate and test features. By this experience, the students learned some important lessons:

- Enforce a hard deadline for new features, 72 hours before any important shows.
- Preserve a complete copy of previous successful worlds, so you have a fallback position for demonstrations.

Other Lessons Learned. High-performance graphics cards are essential for building virtual worlds, as well as testing and using them. Students trying to work at home on the development of behaviors found themselves critically handicapped by the slowness of VRML rendering, and often had to work in subset worlds to get the polygon count low enough to achieve anything at all. Fortunately, the performance of current consumer graphics boards such as the Nvidia GeForce exceeds that of the Accel Galaxy boards that we had in our labs. With adequate hardware and technical support, VRML is a useful tool for projects such as this.

Provide an easily visible cue to identify which avatar is speaking at any time. A major source of confusion was that there was no convenient means to tell, while looking at the screen, which of the visible avatars corresponded to a cast member who was speaking. At a minimum, some sort of motion cue should have been provided.

Videotape your performances! It's hard to assemble and train a team to demonstrate a multi-player virtual world. Videotapes of the actual experience are essential for subsequent demonstrations.

Extensions

Navigating Complex Scenes

The models of man-made objects (pyramids, temples, avatars) in the Caracol world account for most of its texture and polygonal complexity. However, modern graphics cards easily manage that complexity. In contrast, truly complex virtual worlds may contain lush scenes, through which users navigate, perhaps carrying out tasks in which they must observe minute details of their surroundings. This requires detailed models, typically including trees, other types of vegetation and many, diverse man-made structures. To manage this complexity, level-of-detail (simplified) models must be created and managed. LOD-based management is essentially a strategy to select coarse-level representations as an object becomes less and less detailed in the eyes of the viewer.

The essential LOD management problem is 'How can we get the best overall image quality while limiting the number of polygons and maintaining a guaranteed frame rate?' Consideration of this problem leads to several criteria for the effective selection of LOD. Two common ones are frame rate-oriented LOD selection as exemplified in Funkhouser and Sequin[25] and perceptual oriented LOD selection as exemplified in Reddy *et al*.[26] Additional schemes focus on visibility, discarding polygons that are off-screen, oriented away from the viewer or occluded.[27] Cohen-Or *et al*.[28] outline a visibility preprocessing method for an outdoor environment that partitions the view space into cells. Using this method, they formulate the probability for a given object to be visible from a given view cell as a function of distance from the view cell and density of the intervening scene. In comparison to Euclidean (pure distance-based) techniques, this visibility measure provides a better criterion of LOD selection for densely occluded scenes.

The worlds that we are now building contain lush navigable scenes. Several of these are intended to train people on 'search and rescue' missions. Others are designed to train students in research and field techniques. Yet others are intended to teach children about the ecology of rainforests. Our approach is to use occlusion-based LOD selection, creating these less detailed models based on psychological experiments. In general, test subjects more accurately recall typical symmetrical trees; accuracy is greatest for the gross aspects of trees, such as height and leaf size, and accuracy diminishes for features involving smaller branches or unusual curvature. We also learned that subjects are particularly sensitive to color saturation. These experiments are now motivating our development of algorithms to automatically produce lower-LOD models that can be generated from a set of canonical trees and plants.

Streaming Digital Media into 3D Environments

Traditionally, most VE systems lack support for delivering and rendering real-time digital media. In place of this, many systems offer support for canned content such as audio files via HTTP, key-framed animations, and preprocessed animated textures. The absence of live, interactive content has limited interest in Internet-based VE systems to a passing curiosity since such 3D worlds are largely static in nature.

Until recently, bandwidth limitations, the lack of formal network protocol definitions, and the lack of low-cost, high-performance rendering hardware have been at the root of the absence of real-time content. Fortunately, these issues have been addressed by streaming protocols such as RTP and chipsets that incorporate video decoders and are able to render millions of polygons per second.

Stream-based media can now be delivered and rendered in real-time to 3D worlds such as the Caracol Time Travel Project. We have recently developed a streaming server called StreamZ that is capable of capturing and delivering real-time video, audio, and preprocessed content such as AVI and MP3 files. Since the StreamZ server acts both as a point-to-point and a multicasting streamer, a virtual world participant benefits by being able to stream content directly to some selected single user (e.g., telephony conversation), or send and receive multicast streams arbitrated by a central server.

A Convergence of the Virtual and the Real

Worlds like Caracol have the characteristic that much of their content evolves as the world is used. Clearly, any virtual world themed on a real world is a living structure with the potential to become richer over time. Interestingly, even purely fictional worlds, as one usually encounters in MUDs and MOOs, are designed to evolve. However, their content is based solely on the imagination of their inhabitants, not on evolving knowledge of the real world.

Some of our most recent work has focused on the convergence of the virtual and the real. Here we are interested in virtual worlds that are overlays for real places. These augmented realities allow virtual objects to be perceived in the real world. Using location-aware devices with wireless broadband communication, a user can then traverse the real world, interacting with both virtual and real artifacts.

We are currently using this notion of convergence to build virtual worlds in which stories, representing history and culture, are placed for the enjoyment of real and virtual travelers. This project, called Earth Echoes, is being used to capture, preserve and deliver the oral traditions and histories of several unique towns in Florida. One is Winter Park, a city originally developed as a winter resort for wealthy people from the northern United Stares. The other is Eatonville, the oldest incorporated African-American township in the United States.

Conclusions

Experimentation with virtual drama is very challenging. To have much of a chance of success, one needs to have a reliable and well-documented technical substrate, including modeling systems, in place at the outset. The student teams should include subject matter experts, skilled 3D modelers and programmers, as well as humanists with a flair for storytelling. Drama students could add ideas and know-how about organizing a theatrical production.

All of the above elements were missing at the beginning of the Caracol Time Travel Project, and yet many useful lessons were learned, and a compelling virtual performance was achieved. The final product was perhaps 50% of the way toward being sufficiently mature to undergo field trials in middle schools.

ACKNOWLEDGEMENTS

Research was supported by a grant from the National Science Foundation (EIA 9986051) to the first author, and by a contract from the US Army Simulation and Training Command to the first two authors.

References

1. Hayes-Roth B, Brownston L, Gent RV. Multiagent Collaboration in Directed Improvisation. Knowledge Systems Laboratory: San Francisco, CA, 1995. Available: http://ksl.stanford.edu/KSL_Abstracts/KSL-94-69.html [11 May 2001].
2. Maes P. Artificial life meets entertainment: interacting with lifelike autonomous agents. *Communications of the ACM* 1995; **38**(11): 108–114.
3. Kelso M, Weyhrauch P, Bates J. Dramatic presence. *Presence* 1993; **2**(1): 1–15.
4. Moshell JM, Winn W (eds). A special issue on virtual reality and education. *Presence* 1999; **8**(3): 247–365
5. Youngblut C. *Educational Uses of Virtual Reality Technology*. Technical Report D-2128, Institute for Defense Analysis, Alexandria, VA, 1998.
6. Johnson A, Moher T, Ohlsson S, Gillingham M. The Round Earth Project: collaborative VR for conceptual learning. *IEEE Computer Graphics and Applications* 1999; **16**(6): 60–69.
7. Jesup G. Available: http://www-rohan.sdsu.edu/faculty/geosoc/sociology/ [10 May 2001].
8. Morningstar C, Farmer FR. The lessons of Lucasfilm's Habitat. In *Cyberspace: First Steps*, Benedikt M (ed.). MIT Press: Cambridge, MA, 1991.
9. Hughes CE, Moshell JM. Shared virtual worlds for education: the exploreNet experiment. *ACM Multimedia* 1997; **5**(2): 145–154.
10. Hughes CE, Moshell JM, Pullen M. Two dimensional shared virtual worlds in middle and elementary schools: lessons learned. In *Proceedings of the Virtual Worlds and Simulation Conference*, San Diego, CA, January 1998; 139–144.
11. Moshell JM, Hughes CE. The virtual academy: a simulated environment for constructionist learning. *International Journal of Human–Computer Interaction* 1996; **8**(1): 95–110.
12. CTG (Cognition and Technology Group at Vanderbilt). Designing learning environments that support thinking: the Jasper series as a case study. In *Designing Environments for Constructivist Learning*, Duffy TM, Lowyck J, Jonassen DH (eds). Springer: New York, 1993; 9–36.
13. Caracol. Caracol Archaeological Project. Available: http://www.caracol.org/ [26 April 2001].
14. Chase AF, Chase DZ. *Investigations at the Classic Maya City of Caracol, Belize: 1985–1987*. Monograph 3. Pre-Columbian Art Research Institute: San Francisco, 1987.
15. Chase DZ, Chase AF (eds). *Studies in the Archaeology of Caracol, Belize*. Monograph 7. Pre-Columbian Art Research Institute: San Francisco, 1994.
16. Chase AF, Chase DZ. A mighty maya nation: how caracol built an empire by cultivating its 'middle class'. *Archaeology* 1996; **49**(5): 66–72.
17. Chase AF, Chase DZ. The Royal Court of Caracol, Belize: its palaces and people. In *Royal Courts of the Ancient Maya: Data and Case Studies*, Inomata T, Houston SD (eds). Westview Press: Boulder, CO, 2001; 102–137.
18. Schele L, Freidel DA. *A Forest of Kings: The Untold Story of the Ancient Maya*. William Morrow: New York, 1990.
19. Schele L, Matthews P. *The Code of Kings: The Language of Seven Sacred Maya Temples and Tombs*. Scribner: New York, 1998.
20. Martin S, Grube N. *Chronicle of the Maya Kings and Queens: Deciphering the Dynasties of the Ancient Maya*. Thames & Hudson: New York, 2000.
21. Chase AF, Chase DZ. Sixth century change and variation in the southern maya lowlands: integration and disbursement at Caracol, Belize. In *The Years Without Summer: Tracing A.D. 536 and its Aftermath*, Gunn J (ed.). BAR International Series 872, Archaeopress: Oxford, 2000; 55–56.
22. Feld WA. The caves of Caracol: initial impressions. In *Studies in the Archaeology of Caracol, Belize*. Chase DZ, Chase AF (eds). Pre-Columbian Art Research Institute: San Francisco, 1994; 76–82.
23. Chase DZ. Human osteology, pathology, and demography as represented in the burials of Caracol, Belize. In *Studies in the Archaeology of Caracol, Belize*, Chase DZ, Chase AF (eds). Pre-Columbian Art Research Institute: San Francisco, 1994; 123–138.
24. *Caracol: The Lost Maya City*, Pyramid Films, 1988. *Searching for the Maya*, Pyamid Films, 1995.
25. Funkhouser TA, Sequin CH. Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments. In *Computer Graphics Proceedings*, Annual Conference Series, August 1993; 247–254.
26. Reddy M, Walker N, Watson B, Hodges L. Managing level of detail in virtual environments: a perceptual framework. *Presence* 1997; **6**(6): 658–666.
27. Law F, Tan T. Preprocessing occlusion for real-time selective refinement. In *Symposium on Interactive 3D Graphics*, Atlanta, GA, April 1999; 47–53.
28. Cohen-Or D, Fibich G, Halperin D, Zadicario E. Conservative visibility and strong occlusion for viewspace partitioning of densely occluded scenes. *Computer Graphics Forum* 1998; **17**(3): 243–253.

Authors' biographies:

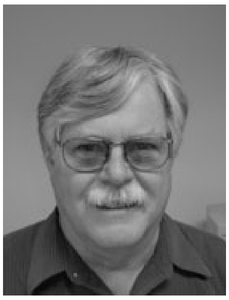

Charles E. Hughes is Professor of Computer Science in the School of Electrical Engineering and Computer Science at the University of Central Florida (UCF), Orlando, Florida. He received the B.A. degree in Mathematics from Northeastern University in 1966, and the M.S. and Ph.D. degrees in Computer Science from Penn State University in 1968 and 1970, respectively. He served on the faculties of Computer Science at Penn State University and the University of Tennessee prior to joining UCF in 1980. His research interests are in distributed interactive simulation, distributed object-oriented programming, and models of distributed computation. He is a member of the ACM, the IEEE and the IEEE Computer Society.

J. Michael Moshell is Professor of Computer Science and Director of the CREAT Digital Media Program at the University of Central Florida (UCF), Orlando, Florida. He received the B.S. degree in Physics from Georgia Tech in 1968 and the Ph.D. degree in Computer Science from Ohio State University in 1975. He served on the faculty of Computer Science at the University of Tennessee prior to joining UCF in 1984. He founded and directed the Visual Systems Laboratory of UCF's Institute for Simulation and Training from 1990 through 1994. His research interests are in the use of simulation and virtual environments for learning and problem-solving. He is a member of the ACM, the IEEE and the IEEE Computer Society.

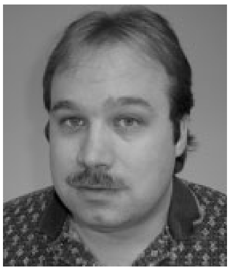

Dean Reed is a Systems Programmer in the Institute for Simulation and Training at the University of Central Florida (UCF), Orlando, Florida. He received the B.S. degree in Computer Science from UCF in 1999. His interests are in distributed interactive simulation and streaming media.

Diane Z. Chase is Professor of Anthropology at the University of Central Florida, Orlando, Florida. She received B.A. and Ph.D. degrees from the University of Pennsylvania in 1975 and 1982, respectively. Her research interests focus on archaeological method and theory in the Maya area, with a particular emphasis on complex societies and hermeneutics, ethnohistory, and ceramic and mortuary analysis. For the last 17 years she has co-directed excavations at Caracol, Belize; before that she directed a seven-year project at Santa Rita Corozal in the same country.

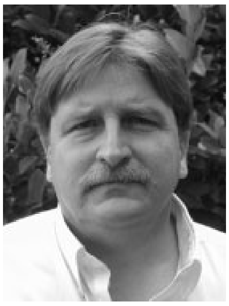

Arlen F. Chase is Professor of Anthropology at the University of Central Florida, Orlando, Florida. He received his B.A. and Ph.D. from the University of Pennsylvania in 1975 and 1983, respectively. His research interests focus on archaeological method and theory in the Maya area, with a particular emphasis on urbanism and ethnicity, hieroglyphic interpretation, settlement patterns, and ceramic and contextual analysis. After working in Mexico and Guatemala, he began archaeological research in Belize in 1978; since 1985 he has co-directed research at the Classic period site of Caracol in the country on an annual basis.