

Virtual Great Barrier Reef: A theoretical approach towards an evolving, interactive VR environment using a distributed DOME and CAVE System

Scot Thrane Refsland¹, Takeo Ojika¹, Tom Defanti², Andy Johnson², Jason Leigh², Carl Loeffler³, and Xiaoyuan Tu⁴

1:Virtual System Laboratory, Gifu University, 1-1 Yanagido, Gifu, 501 JAPAN Thrane@vsl.gifu-u.ac.jp, ojika@oik.info.gifu-u.ac.jp

2:Electronic Visualization Laboratory, University of Illinois at Chicago, Tom@uic.edu, ajohnson@eecs.uic.edu, jleigh@eecs.uic.edu

3: Simation, P.O. Box 1984, Cranberry PA 16066, cel@well.com

4: Intel Corporation, 2200 Mission College Blvd., RN6-35, Santa Clara, CA 95052 xiaoyuan_tu@ccm.sc.intel.com

Abstract. The Australian Great Barrier Reef is a natural wonder of our world and a registered UNESCO World Heritage site hosting 1.5 million visitor-days in 1994/95. Tourism is currently the main commercial use and is estimated to generate over \$1 billion annually.[1] With the coming 2000 Olympics in Australia, tourism increases will substantially present a major conservation and preservation problem to the reef. This paper proposes a solution to this problem through establishing a virtual reality installation that is interactive and evolving, enabling many visitors to discover the reef through high quality immersive entertainment. This paper considers the technical implications required for a system based in Complexity: a distributed DOME and CAVE architectural system; a mixed reality environment; artificial life; multi-user interactivity; and hardware interfaces.

1 Introduction

"Well, life emerged in the oceans," he adds, "so there you are at the edge [diving off the continental shelf], alive and appreciating that enormous fluid nursery. And that's why the 'edge of chaos' carries for me a very similar feeling: because I believe life also originated at the edge of chaos. So here we are at the edge, alive and appreciating the fact that physics itself should yield up such a nursery..."

- Chris Langton, artificial life pioneer, commenting on Complexity.

1.1 A Cinematic Narrative Virtual Walkthrough

Jack, Jill and Tanaka have entered the Virtual Barrier Reef Exhibition, where "Nintendo" meets classroom. Before going into the environment, they have to choose from several different experiential options: Jack picks the "Marine Biologist," Jill picks the "Dolphin" and Tanaka the "Snorkeler."




Fig. 1. Overview of the Virtual Great Barrier Reef installation consisting of a 15 meter diameter DOME, Interactive WALLs, 3 CAVE systems and a motion simulator.

When they put on their equipment mainly consisting of 3D shutter glasses, audio communication and other interactive devices, they enter the Dive Vehicle (motion simulator) that takes them on a small "journey" to the reef. The Immersion hatch opens and they exit into the main DOME [Fig. 1] finding themselves immersed in 20 meters (60 feet) of water. The main DOME [2] (15 m diameter x 15m high, 360° x 180°), enables them to see fish and other species swim about in 3D. They see a mix of autonomous fish, avatars and avatar guides that assist, guide and role-play. They hear soundscapes of the reef. They can also see 3 CAVEs [3]. They were told that the installation was directly linked to the real reef to affect the climatic conditions within the installation. The installation's environment was persistent, evolving and adapting just like the real reef.

Jack finds a school of fish that holds his interest, and points to it using the "laser wand" pointer on top of his Dive Panel. The panel displays the information about the fish, life habits, predators, eating, mating habits, etc. Then, a big grouper swims to him (guide avatar) and engages him in a conversation about the reef. A dolphin joins in and he finds out its Jill, and they are able to communicate together through the group communication channel with their headsets.

Tanaka finds an autonomous fish guide and several descriptor icons floating nearby. He selects the Japanese version. He inserts his smart card in the nearby stand, then experiments with the inter-related lives of the species inhabiting that community by moving the various avatars around with a Fish data Glove. At one point a mythical aboriginal character swims by and heads for one of the CAVEs. At the warning beep on each of their timers, Jill and Tanaka to return to the visitor's center. Jack extends his visit by buying more air directly through his smart card and Dive Panel. Once in the visitors' center, Jill and Tanaka insert their smart cards to get a color print out of their activities during their visit.

1.2 The Complex Architecture

The system is a large distributed network based upon proven, stable DOME and CAVE technologies. Similar reference models in which this project is based upon can be found in the past projects of NICE [4] and ROBOTIX: MARS MISSION [5].

The NICE project is a collaborative CAVE-based environment where 6 to 8 year old children, represented as avatars, collaboratively tend a persistent virtual garden. This highly graphical, immersive virtual space has primarily been designed for use in the CAVE, a multi-person, room-sized virtual reality system. As the CAVE supports multiple simultaneous physical users, a number of children can participate in the learning activities at the same time.

Interactive DOME projects including the Carnegie Science Center's "ROBOTIX: MARS MISSION, was the first example of Interactive Virtual Reality Cinema. Each audience member had a three button mouse device that allowed them to make group decisions on story and piloting spacecraft. The graphics were rendered in real-time, called from an extensive database of Mars's terrain from NASA. The partial dome sat 35 people, and was 200 degrees wide and 60 degrees high. The system used a SGI ONYX, with 4 R10000 processors, Infinite Reality Engine, Audio Serial Option board, and display boards. The polling system, custom built, fed into a PC which talked to the SGI. Majority vote ruled. The dome display used 3 state of the art Electrohome RGB projectors. Most important, only one graphic pipeline was used, and split into three views using custom built software and edge blending. The overlap on edges was 25%, and it appeared seamless and no distortion. Performance, 30-90+ fps, no lag. Audio was 3-D spatialized, custom built software, with floor shakers. It ran for 9 months, no glitches. Effective, there were some people who *thought* they went to Mars.

1.3 Methodology of Application

Whilst the true objective of this installation is to install a strong sense of conservation and preservation to the future visitors of the Great Barrier Reef, it seeks to apply it through a Constructionist [6] and "Immersive role-play" educative model steeped heavy in "Nintendo" style interaction of total consciousness submersion. Thus, our methodology objectives are simple – give the visitor an experience that they constructed for themselves, immerse them in wonder and inspiration, give them a reason to care, in their minds really take them there, and finally, afterwards give them something material they can use to prove and remember their experience.

2 Distributed DOME and CAVE Architectural System

2.1 The DOME zone

The DOME zone is constructed of CG and real-time/pre-captured video which is projected onto the scrim of the DOME together in layers to assist with overall processing quality of CG and simulations. The background is a mix of live or pre-recorded scenes in the reef, and the CG fish are projected bright enough over the top of it to create the mixed illusion. Of course there will have to be certain areas, such as coral and other big blocks of CG that have to be animated, but these aren't that computationally intensive as they update on a slightly slower level.

2.2 Multiplexing and Mixed topologies

Since the DOME is the central "Hub" [Fig. 2] of the environment physically, and holds the most connectivity to all other zones, we will use it as the focus to consider the many facets of multiplexing within the environment and the various other network topologies and protocols that connect to it.

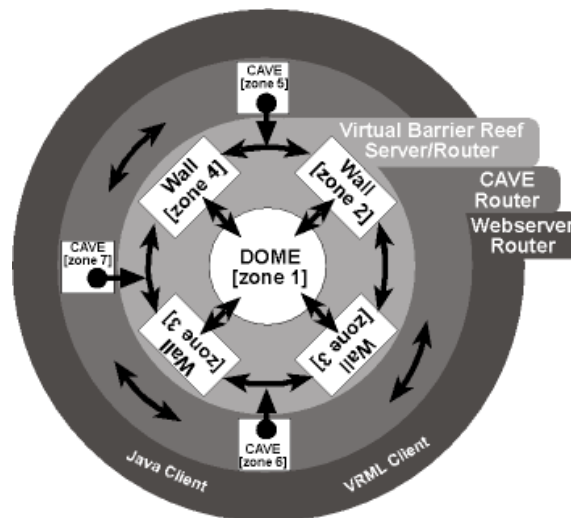


Fig. 2. Distributed Relationship of the DOME, CAVES, WALLs and the Internet.

2.3 The WALL zones (4 WALLs)

The walls will be the most computational expensive of all the zones in the environment as most of the direct interactivity will happen around them. Walls are interactive in that they might have species that will react when visitors walk by them. Example would be an alife moray eel might come out to feed, and when a visitor walks up to it quickly, a sensing device gives it spatial reaction information. Spot projectors and smaller independent computer systems are displaying various small interactive displays, while the meta environment was separate and being processed by the SGI system.

2.4 CAVE zones: (3 CAVEs)

CAVEs are separate zones or environments that are story related. They can host indigenous or mythical stories related to the reef and its inhabitants. Here the visitors can have a more immersive, intimate experience with the denizens of the reef, as the sea creatures swim and interact with the visitors within the 3 meter cube of each CAVE.

The CAVEs use the standard network topologies and protocols being developed as part of the CAVERN network [7] to deal with the multiple information flows maintaining the virtual environments.

Each of the CAVEs, WALLs, and the DOME making up the environment will be connected via a high-speed network, allowing information such as avatar data, virtual world state / meta data and persistent evolutionary data to pass between them. This will allow avatars in the CAVEs to swim over onto the WALLS or the DOME, and allow museum guides in the guise of reef inhabitants to interact with the visitors. This will also allow remotely located CAVE users in the United States, Europe, and Asia to join into the collaboration. Similar connectivity has been used in the NICE environment [8] to link multiple CAVE sites in the United States to Europe and Japan.

3 Mixed Reality Environment

The environment of this project uses video, lighting, and sound techniques to establish an environment that convinces the user that they are indeed immersed in the Great Barrier Reef.

Video: Video will be used in strategic places includes some parts of the walls and the floor. The main imaging of the floor is video, with spot CG projections to provide a better illusion of immersion and non-repetition.

Sound: immersive and 3D in its localization, the sound system provides authentic sounds from autonomous life and "extra spatial" depth sounds.

3.1 Real-time Connectivity to the Reef

Climatic Interference: The virtual model is connected real-time to a section of the real reef so that it can automatically adjust the virtual model. Temperature, current, air pressure and other climatic conditions on the reef will cause several actions in the virtual environment, such as temperature of the room, flow of the virtual debris, virtual water currents, etc.

3.2 Complex and evolving environment

This environment is considered to be a "living entity", possessing an authentic as possible simulation to the real reef.

For simulating simple fluid flow, we consider techniques similar to the ones used by Wejchert and Haumann Wejchert 91 [9] [Fig. 3] for animating aerodynamics.




Fig. 3. The flow primitives (reproduced from the paper by Wejchert and Haumann)

Assuming inviscid and irrotational fluid, we can construct a model of a non-turbulent flow field with low computational cost using a set of *flow primitives*. These include *uniform flow* where the fluid velocity follow straight lines; *source flow*--a point from which fluid moves out from all directions; *sink flow* which is opposite to the source flow; and *vortex flow* where fluid moves around in concentric circles. This process is also directly linked with the real reef current flow that is constantly supplying real-time information. Both of the processes integrated together enable lower computation requirements and increase the authenticity of the virtual environment.

Much of the seaweed, plankton, coral and other movable plant/marine life are simulated and respond dynamically to water currents and movements of other life in a realistic manner. In this case it is proposed to use a system which builds the plants in a mass-spring chain assembly. In order to obtain computational efficiency, we do not calculate the forces acting on the geometric surface of each plant leaf, rather, we approximate the hydrodynamic force at each mass point.

4 Artificial Life

The artificial life treatment in this project is an exciting and innovative one. As mentioned earlier, one of the educative goals in this project is to enable people to "think like a fish" so they can experience first hand through the "eyes of a fish" what it's like to live on the reef.

To achieve this, we have developed two types of models for Alife. The first model is a totally autonomous animal that behaves according to the events in the environment. It has a life span, is integrated into the ecological system, and possesses authentic behaviors of its real life counterpart.

The second model is a similar model, only that the "brain and motor skills" are detached and placed into the control of a visitor using a dataglove to direct the fish. The innovation here is that even though a visitor is controlling the "will and motor" of the artificial fish, it still has behavior and characteristic traits that will not allow you to act outside of the character set for that fish.

4.1 Artificial Animals for Life-like Interactive Virtual Reef Creatures

What's the key to bringing you a captivating experience as a tourist of the virtual Great Barrier Reef? Whether you are merely observing the schools of sardines, playful dolphins and vicious sharks, or you are living the aquatic life by playing the role of one of the virtual reef dwellers (want to be a dolphin, parrot fish or maybe a reef shark? Sure you can), realism is the password to the ultimate immersive experience. Here we are not just talking about visual authenticity in the appearance of the digital sea creatures, but also in their physical environment, and most importantly, in the way they move, perceive and behave. In an interactive environment like ours, behavioral realism, especially, autonomy of the virtual creatures is indispensable.

How are we to achieve such integrated realism? Our answer is to build Artificial Life (or Alife) models of the aquatic animals. The properties and internal control mechanisms of these artificial animals should be qualitatively similar to those of their natural counterparts. Specifically, we construct animal-like autonomy into traditional graphical models by integrating control structures for locomotion, perception and action. As a result, the artificial fishes and mammals have 'eyes' and other sensors to actively perceive their dynamic environments; they have 'brains' to interpret their perception and govern their actions. Just like real animals, they autonomously make decisions about what to do and how to do in everyday aquatic life, be it dramatic or mundane. This approach has been successfully demonstrated by the lifelike virtual undersea world developed in [10, 11].

We tackle the complexity of the artificial life model by decomposing it into three sub-models:

- A *graphical display model* that uses geometry and texture to capture the form and appearance of any specific species of sea animal.
- A *bio-mechanical model* that captures the physical and anatomical structure of the animal's body, including its muscle actuators, and simulates its deformation and dynamics.
- A *brain model* that is responsible for motor, perception and behavior control of the animal.

As an example of such an artificial life model, [Fig.4] shows a functional overview of the artificial fish that was implemented in [10, 11]. As the figure illustrates, the body of the fish harbors its brain. The brain itself consists of three control centers: the motor center, the perception center, and the behavior center. These centers are part of the motor, perception, and behavior control systems of the artificial fish. The function of each of these

systems will be previewed next.

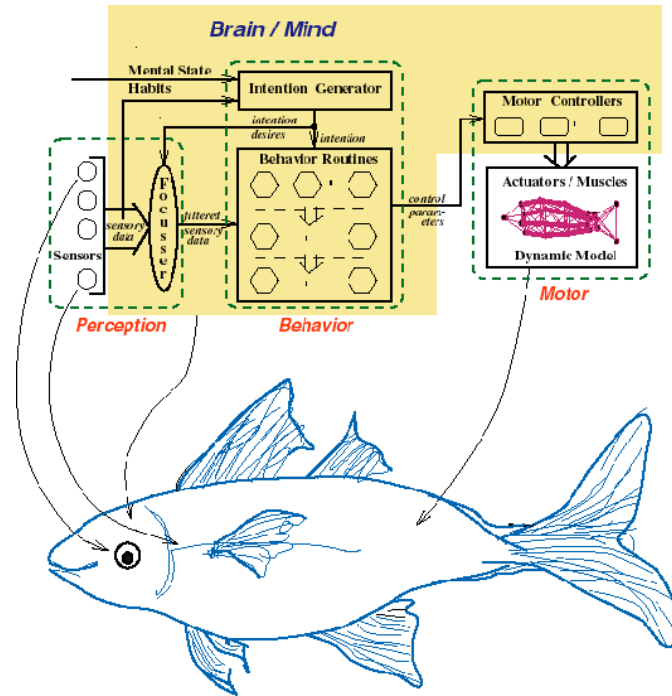


Fig. 4. Diagram of an autonomous and avatar model fish

4.2 Motor System

The motor system comprises the dynamic model of the sea creature, the actuators, and a set of motor controllers (MCs) which constitutes the motor control center in the artificial animal's brain. Since our goal is to animate an animal realistically and at reasonable computational cost, we seek to design a mechanical model that represents a good compromise between anatomical consistency, hence realism, and computational efficiency. The dynamic fish model [12, 13] represents a good example. It is important to realize that adequate model fidelity allows us to build motor controllers by gleaned information from bio-mechanical literature of the animal [14, 15]. Motor controllers are parameterized procedures, each of which is dedicated to carrying out a specific motor function, such as "swim forward", "turn left" or "ascend". They translate natural control parameters such as the forward speed, angle of the turn or angle of ascent into detailed muscle or fin/leg actions. Abstracting locomotion control into parameterized procedures enables behavior control to operate on the level of motor skills, rather than that of tedious individual muscle/limb movements. The repertoire of motor skills forms the foundation of the artificial animal's functionality.

Given that our application is in interactive VR systems, locomotion control and simulation must be executed at interactive speed. Physics-based dynamic models of animals offer physical realism but require numerical integration. This can be expensive when the models are complex.

There are various ways to speed things up. One way is to have multiple models of increasing complexity for each species of animal. For example, the simplest model will just be a particle with mass, velocity and acceleration (which still exhibits basic physical properties and hence can automatically react to external forces, such as water current, in a realistic manner). The speed-up can then be achieved by way of 'motion culling', where when the animal is not visible or is far away, only the simplest model is used and the full mechanical model is used only when the animal is near by. This method, however, is view-dependent and can be tricky when the number of users is large. Another way is to use bio-mechanical models whenever possible (when they are simple enough to run the simulation at real time) and with more complex animals, we can build pre-processed, parameterized motion libraries (which can be played back in real time), like what is used by most of today's games. These (canned) motion libraries can be built from off-line simulations of elaborate bio-mechanical models of the animal or from motion-captured data. One thing to keep in mind is that, no matter what the underlying locomotion model is, its interface to the behavior system should be kept comparable as parameterized procedures of motor skills.

4.3 Perception System

Perception modeling is concerned with:

1. Simulating the physical and optical abilities and limitations of the animal's perception.
2. Interpreting sensory data by simulating the results of perceptual information processing within the brain of the animal.

When modeling perception for the purposes of interactive entertainment, our first task is to model the perceptual capabilities of the animal. Many animals employ eyes as their primary sense organ and perceptual information is extracted from retinal images. In a VR system, such "retinal" images correspond to the 2D projection of the 3D virtual world rendered from the point of view of the artificial animal's "eyes". However, many animals do not rely on vision as their primary perceptual mode, in which case vision models alone may not be able to appropriately capture the animal's perceptual abilities.

It is equally important to model the limitations of natural perception. Animal sensory organs cannot provide unlimited information about their habitats. Most animals cannot detect objects that are beyond a certain distance away and they usually can detect moving objects much better than static objects [14]. If these properties are not adequately modeled, unrealistic behaviors may result.

Moreover, at any moment in time, an animal receives a relatively large amount of sensory information to which its brain cannot attend all at once. Hence there must be some mechanism for deciding what particular information to attend to at any particular time. This process is often referred to as attention. The focus of attention is determined based upon the animal's behavioral needs and is a crucial part of perception that directly connects perception to behavior.

Unfortunately, it is not at all well understood how to model animal sensory organs, let alone the information processing in the brain that mediate an animal's perception of its world. Fortunately, for our purposes, an artificial animal in its virtual world can readily glean whatever sensory information is necessary to support life-like behavior by directly interrogating the world model and/or exploiting the graphics-rendering pipeline. In this way, our perception model synthesizes the results of perception in as simple, direct and efficient a manner as possible.

The perception system relies on a set of on-board virtual sensors to provide sensory information about the dynamic environment, including eyes that can produce time-varying retinal images of the environment. The brain's perception control center includes a perceptual attention mechanism which allows the artificial animal to train its sensors at the world in a task-specific way, hence filtering out sensory information superfluous to its current behavioral needs. For example, the artificial animal attends to sensory information about nearby food sources when foraging.

4.4 Behavior System

The behavior system of the artificial animal mediates between its perception system and its motor system. An *intention generator*, the animal's cognitive faculty, harnesses the dynamics of the perception-action cycle and controls action selection in the artificial animal. The animator establishes the innate character of the animal through a set of *habit parameters* that determine whether or not it likes darkness or whether it is a male or female, etc. Unlike the static habits of the animal, its mental state is dynamic and is modeled in the behavior system by several *mental state* variables. Each mental state variable represents a distinct desire. For example, the desire to drink or the desire to eat. In order to model an artificial animal's mental state, it is important to make certain that the modeled desires resemble the three fundamental properties of natural desires: (a) they should be time varying; (b) they should depend on either internal urge or external stimuli or both; (c) they should be satisfiable. The intention generator combines the habits and mental state with the incoming stream of sensory information to generate dynamic goals for the animal, such as to chase and feed on prey. It ensures that goals have some persistence by exploiting a single-item memory. The intention generator also controls the perceptual attention mechanism to filter out sensory information unnecessary to accomplishing the goal in hand. For example, if the intention is to eat food, then the artificial animal attends to sensory information related to nearby food sources. Moreover, at any given moment in time, there is only one intention or one active behavior in the artificial animal's behavior system. This hypothesis is commonly made by ethologists when analyzing the behavior of fishes, birds and four-legged animals of or below intermediate complexity (e.g. dogs, cats) [15, 16]. At every simulation time step, the intention generator activates behavior routines that input the filtered sensory

information and compute the appropriate motor control parameters to carry the animal one step closer to fulfilling the current intention. The intention generator Primitive behavior routines, such as obstacle avoidance, and more sophisticated motivational behavior routines, such as mating, are the building blocks of the behavioral repertoire of the artificial animal.

5. Interfaces

Because it is our past experience that all technology must be bulletproof and withstand the rigorous demands of an unrelenting public, this equipment is designed along the lines of actual DIVE equipment which is very durable and rugged, big and easy to use.

5.1 Tracking/Sensing

A number of tracking options are available, his environment will consider the use of mix from magnetic [17], to un-encumbered technology such as camera based systems [18]. Sensing the bulk of visitors is done with infrared and video. Specific occupant tracking is done either with localized gloves cabled to the specific site location, or through wireless. Shutter glasses should be independent and able to sense any WALL or CAVE. Artificial Life use sensing to establish the location of visitors within the environment and can respond to them as a normal fish would respond.

5.2 Weight Belt with Dive Panel

This equipment enables visitors to interact with the environment and gain statistical information about the life species in the installation. The weight belt holds the battery pack and wireless communications, tracking and wand/pointer hardware, and a smart/memory card slot for recording individual data and experiences. Attached to it is a "Dive Panel" which is a pointing wand, and small LCD monitor to display information about the species that are under investigation, and other climatic information. Through this, the visitor is able to access the knowledge base stored on the memory card and other data being generated real-time.

5.3 Enhanced Shutterglasses

Shutter glasses are connected to a timer so that when the visitor's use all their "air", they shut down and are inoperable. This makes the environment very difficult to see and forces the visitor to exit the installation. Some of the more sophisticated models for avatars and "Marine Biologists" have audio communications built in.

5.4 Taking home a bit of the Virtual Reef

Something else that we think is really important from NICE is giving people some kind of artifact from the experience. Since VR is such a rare experience, and you really can't take anything physical away from the virtual space, it helps people when they try to describe what they did to other people, and enhance the experience and memory of the event.

Each participant in the virtual reef environment gets a smart/memory card to record a "narrative" of everything they did, characters they talked to, actions they performed. The narrative structure captures these interactions in the form of simple sentences such as: "Mary (disguised as a dolphin) stops by to see Murray the Moray eel to find out what he eats. Mary and Murray have dinner together."

The story sequence goes through a simple parser, which replaces some of the words with their iconic representations and stores the transcript onto their Memory card. This gives the story a "picturebook" style that the visitor can print out and keep to remember the various aquatic life they met that day.

5.5 Smart/Memory Cards

We are currently investigating the use of Smart/Memory Cards to perform functions like: tracking and recording the user's actions, supplying a knowledge database, and providing financial calculations for visitation fees.

6 Interactivity

6.1 Visitor Interactivity

Avatar: There are many species in which a visitor can explore, from fish and other animals. Avatars can be as simple as the ones found at the localized exhibits for tourists to try out, to the sophisticated ones like dolphins and sharks. Avatars ultimately come under the domain of the simulated eco-cycle, so anything could happen: the visitor's avatar could be eaten at any time, and eventually will go through the entire life cycle of the reef.

Snorkeler: This level is simple roaming and exploring the environment with minimal hardware, such as shutter-glasses to immerse them in 3D. The smart card enables them to utilize the local interactive stands available throughout the environment.

Marine Biologist: This level of interaction is highly independent and is able to program their own interactivity levels for simulation, planning, studying, etc. The equipment this visitor has is the most advanced versions of the standard models.

6.2 Artificial Life Interactivity

The artificial life in the environment is just as interactive towards visitors as visitors are to them. They can and will interact with visitors in many different behaviors, from being friendly, territorial, viscous, funny, etc. Fast movement will scatter fish. In this way, true interactivity is two-way, creating a complex environment.

6.3 Staff Interactivity

Guides: These are sophisticated avatars that can interact with anyone in the environment. These people can be in the environment or hidden from view. They have audio and visual communication devices to drive avatars such as Mermaids, large fish, human animations or indigenous characters. This level of interactivity can be used for guided tours, assistants, observers and other special requirements. At this level, the guide is able to leave discovery objects which visitors can touch and find out more information.

7 Conclusion and future works

This system has been designed using existing and known solutions of virtual reality, networking, evolutionary theory and artificial life, and yet, the innovation lies within the organization of each application into complexity. It has been designed with the forethought of being able to use the physical architecture as a generic shell to enable other stories besides the reef to be told, and the application able to be used in other systems, including retrofitted Planetariums.

Acknowledgements

The Gifu Research and Development Foundation and The International Society on Virtual Systems and MultiMedia, Gifu Japan

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible through major funding from the National Science Foundation (NSF), the Defense Advanced Research Projects Agency, and the US Department of Energy; specifically NSF awards CDA-9303433, CDA-9512272, NCR-9712283, CDA-9720351, and the NSF ASC Partnerships for Advanced Computational Infrastructure (PACI) program. EVL also wishes to acknowledge Silicon Graphics, Inc. and Advanced Network and Services for their support. The CAVE and ImmersaDesk are trademarks of the Board of Trustees of the University of Illinois.

References

1. McPhail, I. The Changing Environment of Managing Use in the Great Barrier Reef World Heritage Area Presented at the ABARE Outlook '96 Conference, Melbourne, February 1996. Great Barrier Reef Marine Authority <http://www.gbrmpa.gov.au/chair/#Top>
2. Goto Virtuarium DOME system. www.bekkoame.or.jp/~goto-co/
3. Cruz-Neira, C., Sandin, D.J., and DeFanti, T.A. "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE." In Proceedings of SIGGRAPH '93 Computer Graphics Conference, ACM SIGGRAPH, August 1993, pp. 135-142.
4. Johnson, A., Roussos, M., Leigh, J., Barnes, C., Vasilakis, C., Moher, T., "The NICE Project: Learning Together in a Virtual World,." In the proceedings of VRAIS '98, Atlanta, Georgia, Mar 14-18, 1998, pp.

176-183.

5. Carnegie Science Center, "ROBOTIX: MARS MISSION",
<http://www.csc.clpgh.org/Exhibits/Mars/index.htm>
6. Piaget, J. To Understand is to Invent: The Future of Education. Grossman, New York, 1973.
7. Leigh, J., Johnson, A., DeFanti, T., "Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments,." In the proceedings of Supercomputing '97 San Jose, California, Nov 15-21, 1997.
8. Johnson, A., Leigh, J., Costigan, C., "Multiway Tele-Immersion at Supercomputing '97." To appear in IEEE Computer Graphics and Applications, July 1998.
9. Wejchert, J. and Haumann, D. Animation aerodynamics. ACM Computer Graphics, SIGGRAPH'91, 25(4): pp.19-22. 1991.
10. Tu, X. 1996. Ph.D Thesis, Department of Computer Science, University of Toronto.
<http://www.dgp.toronto.edu/people/tu>
11. Tu, X. and Terzopoulos, D. Artificial fishes: Physics, locomotion, perception, behavior. In ACM Computer Graphics, Annual Conference Series, Proceedings SIGGRAPH'94, pages 43-50, Orlando, FL. ACM Computer Graphics. 1994
12. Blake, R. Fish Locomotion. Cambridge University Press, Cambridge, England. 1983.
13. Alexander, R.. Exploring Bio-mechanics. Scientific American Library, New York. 1992.
14. Tansley, K. Vision in Vertebrates. Chapman and Hall, London, England. 1965.
15. Tinbergen, N. The Study of Instinct. Clarendon Press, Oxford, England. 1951.
16. Manning, A. An Introduction to Animal Behavior. Addison-Wesley Publications, Massachusetts, 3rd edition. 1979.
17. Semwal, S., Hightower, R. Stansfield, S., Closed form and Geometric Algorithms for Real-Time Control of an Avatar, Proceedings of IEEE, VRAIS96, pp. 177-184 (1996).
18. Krueger, M., Artificial Reality II. Addison Wesley Publishing Company, Reading, MA pp. 1-277 (1991).