

COMPUTER GRAPHICS

Volume 28 Number 2 May 1994

A publication of ACM SIGGRAPH



Focus: Interactive Entertainment
Design, Implementation and Adrenaline



Placeholder: Landscape and Narrative in Virtual Environments

Brenda Laurel, Rachel Strickland, Rob Tow, Interval Research Corp.

A note on authorship

Sections of this paper are authored by different individuals. The people who worked on Placeholder shared a core set of goals, but each individual contributor had distinct interests, values and ways of working. Each of us reflects on it differently, seeing different shortcomings and strengths in the work, and learning different lessons from the process. For these reasons it seems important to preserve the singularity of each author's voice.

Virtual Reality as Entertainment

Brenda Laurel, Interval Research Corp.

The idea of using virtual reality for entertainment purposes is actually quite recent in the history of the technology. Early VR entertainment applications, appearing in the late 1980s, were extensions of the existing "serious" application of simulation-based training. Flight-sim technology—motion platforms used in sync with motion video or animation—were much more readily adapted to theme park applications than immersive VR with head-mounted displays. Motion platform rides, of which *Star Tours* is probably the best known example, trade off individual viewpoint control and the sense of agency for thrilling, finely calibrated effects and the optimization of "throughput"—that is, getting the most people through the ride in the least time. Networked pods, as used in Virtual World Entertainment systems (previously Battletech), are also derived from flight-sim technology.

"Classic" virtual reality, employing head-mounted displays and various forms of body tracking, presents serious throughput problems in theme park settings. It takes time to get the gear onto the participants. Only a handful of people can experience the attraction simultaneously (although a much larger audience might watch what the people "inside" the VR are doing). A hard-driving plot with distinct beginning, middle and end is a great way to control how long an experience takes, but "classic" VR is inimical to this kind of authorial control—it works best when people can move about and do things in virtual environments in a relatively unconstrained way.

In fact, it may be that the nature of VR makes it inappropriate to think of it as an entertainment medium. Mass entertainment implies the consumption of some performance by a large audience. Generally speaking, the size of the audience is inversely proportional to the degree of influence over the course of events that can be afforded any one person.¹ Nor can we simply turn to human-to-human interaction as the source of engagement and still support a large number of simultaneous participants; virtual spaces seem not to differ from actual ones in terms of social and attentional constraints posed by crowds. There seems to be an upper limit (probably the famous seven plus or minus two) on the number of people who can interact meaningfully or pleasurably with one another (which is why little clusters form at cocktail parties). Even for a small audience, VR is inimical to the kind of passive receptivity that is appropriate in films and amusement park rides. The experience of VR hinges on human action and the environment's response. This is true in both perceptual and emotional terms. In VR, one is not *done unto*, but *doing*.

If, on the other hand, the goal is to create a technologically mediated environment where people can *play*—as opposed to being entertained—then VR is the best game in town. When children play, they typically use their imaginations quite actively and constructively to invent action and assign meaning to materials (or

make or find new ones) as the need arises. In VR as in children's play, there is no sharp distinction between "authoring" and "experiencing." With Placeholder, we learned that adults can play in the same way—when their imaginations are booted up by a rich virtual environment.

Placeholder was presented in a public venue with very low throughput (two participants approximately every 15 minutes). A handful of onlookers were also usually present, watching the participants in the studio and tapping into their experiences through monitors and earphones. It is not likely that people could have a satisfactory experience in a rich VR world in less time or with many more simultaneous participants, but one could imagine working out a way for a much larger audience to look in on the action—by utilizing the virtual camera techniques developed by David Zeltzer's group at the MIT Media Lab, for instance, to automatically synthesize a cinematic presentation of what is going on rather than forcing people to try to construct a coherent mental picture of the action by looking at two individuals' eye-views of the world.

We believe that the theme parks of the world—at least some of them—will eventually be virtual. In the end, we think that VR works like Placeholder will find their natural home on computer networks rather than in public venues. When our version of the Grand Canyon is up on the net, there won't be several thousand tourists staring over the virtual rim. Instead, there will be a thousand copies of the Canyon running where people can experience a magical place in solitude or in the company of a few good friends—each subtly distinguished by season, weather, historical epoch, critters, voices in the wind and rocks and the marks of those who have gone before.

About Placeholder

Experiences are said to take place. One comes to know a place with all one's senses and by virtue of the actions that one performs there, from an embodied and situated point of view. The mind, observes naturalist Barry Lopez, is a kind of projection within a person of the place which that person inhabits: "Each individual undertakes to order his interior landscape according to the exterior landscape." The environment proceeds to record our presence and actions and the marks that we place there—this is a reciprocal affair.

Placeholder is the name of a research project which explored a new paradigm for narrative action in virtual environments. The geography of Placeholder took inspiration from three actual locations in the vicinity of Banff National Park in Alberta, Canada—the Middle Spring (a sulfur hot spring in a natural cave), a waterfall in Johnston Canyon and hoodoos (rock formations created by erosion) overlooking the Bow River. Three-dimensional videographic scene elements, spatialized sounds and words and simple character animation were employed to construct a composite landscape that could be visited concurrently by two physically remote participants using head-mounted displays. People were able to walk about, speak and use both hands to touch and move virtual objects.

People's relationships with places and the creatures who inhabit them have formed the basis of many traditions and spiritual practices, as well as ancient stories and myths. The graphic elements in Placeholder were adapted from iconography that has

¹ An interesting exception is the Cinematrix™ technology developed by Loren Carpenter (and premiered at SIGGRAPH'91), where a very large group interacts on a large video screen via wands with red and green reflective material on them. Each individual controls the color of a logical pixel by turning the wand. Crowds have been observed to learn very quickly to cooperate well enough to play mass games of Pong, make intricate patterns, and even control a flight simulator. This is certainly not virtual reality, but it is a robust technologically enabled form of mass interaction.

been inscribed upon the landscape since paleolithic times. Narrative motifs that revealed the archetypal characters of landscape features and animals were selected from aboriginal tales. Four animated spirit critters—spider, snake, fish, and crow—inhabited this virtual world. A person visiting the world could assume the character of one of the spirit animals and thereby experience aspects of its unique visual perception, its way of moving about and its voice. Thus the critters functioned as “smart costumes” that changed more than the appearance of the person within.

People sometimes leave marks in natural places—pictograms, petroglyphs, graffiti or trail signs for example. In Placeholder, people were able to leave voicemarks—bits of spoken narrative—that could be listened to and rearranged by anyone who passed through. The virtual landscape accumulated definition through messages and storylines that participants left along the way. We hope that the ideas we explored in Placeholder will foster the emergence of new forms of narrative play.

Capturing the Sense of a Place

Rachel Strickland, Interval Research Corp.

Most computer graphic virtual environments—and video adventure games in particular—that had heretofore fallen in our paths consist of synthetically generated scenery that comes from nowhere on earth. Even the flight simulator examples with airport runways precisely dimensioned and positioned on geographically correct terrain models most starkly reflect the world of cartoons.

One of our objectives with Placeholder was to experiment with capturing actual places—in the attitude of landscape painting traditions or documentary cinema, for example—using video and audio recorded on location as the raw material for constructing the virtual environment. It must be emphasized that we were not concerned with achieving a high degree of sensory realism—something bristling with polygons and mips that might induce a perfect audiovisual delusion of sticking your head in the “real” waterfall. No, it gets more slippery than that. What we have really set out to capture or reproduce is just the simplest “sense of place.”

Genius loci

For one thing, there is the *genius loci*—a Latin phrase for the “guardian spirit of a place,” whose presence accounts for the life of the place and determines its character or essence. Something like this ancient Roman concept is common among indigenous cultures throughout the world. Architectural scholar Christian Norberg-Schulz, who wrote the book entitled *Genius Loci*, proposes two levels of analysis for articulating the structure of place:

- Spatial organization, or knowing where one is here. “Space” denotes the three-dimensional organization of the elements which define a place. This organization may correspond to landscape features. It may reflect early parcel subdivisions or transportation routes. It may be shaped by local building methods, materials and codes. It may derive from abstract geometric principles. The spatial organization of the place influences how people orient themselves in this environment, how they find their way around, how they avoid getting lost.
- Character, or knowing how one is here. Compared with space, which can be mapped, character is too comprehensive to permit succinct specification, and it is also more concrete. It is a condition of atmosphere, weather, lighting and the passage of time, as well as a function of space-defining form. “Any real presence is intimately linked with a character.”

Umwelt

Another sense of the sense of place that influenced our designs for Placeholder is suggested by the German word “*umwelt*.” The naturalist Jakob von Uexküll tried to imagine the physical world as lived and perceived by different animals. He used the word *umwelt* to express the organized experience—or point of view—

unique to any creature, which depends on that particular creature’s sensory and cognitive apparatus. Employing virtual environment technology to explore alternate *umwelten* has been one of our irrefutable motives. The scheme for Placeholder included:

- experimenting with approximations of the sensory-motor experiences of non-human creatures
- experimenting with narrative strategies for supporting transformation, locomotion, spatial and temporal discontinuities
- exploring concepts of space and time that underlie representation systems which have been developed by cultures whose views of the world differ dramatically from our own

Place recording and modeling techniques

The three locations selected for Placeholder—cave, waterfall and river valley—asserted strikingly differentiated characters. By matching capture and representation techniques to the unique qualities of the respective places and by amplifying their contrasts, we hoped to distill distinct environmental caricatures. For example, it was determined that the waterfall model should incorporate motion video to render the dynamic flow of the water. The sense of the cave should be auditory rather than visual—a dimly illuminated quick sketch surrounding a lively array of localized sound sources.

Our predilection for sampling and representing actual places rather than synthesizing environments from scratch was reinforced by the collaboration of Michael Naimark. Techniques for constructing 3D computer models out of camera originated imagery were based on Naimark’s previous experiments. One method involved panoramic tiling of multiple video images onto a spherical wireframe. Another used the video picture as a guide for deforming the surface of the wireframe model to approximate the contours of the original.

Extensive location scouting and a series of preliminary trials reminded us that nature affords few landscapes of sufficiently simple form to reveal themselves to a single point of view. Two questions that abided with us throughout the process deserve further pondering:

- How to capture a place simultaneously rather than sequentially with a time based medium like video—in order to provide for more than one way of experiencing the representation of the place during a stretch of time?
- How to capture a place from multiple camera positions all at once, and how to join the several vantage points into a spatially coherent representation?

Traditions such as Chinese landscape painting, impressionism and cubism experimented with a range of strategies for integrating time and multiple viewpoints into the depiction of places. In the entire history of painting, considering the many achievements that have been made by artists—in the pictorial representation of light and color and texture, for example—it’s curious, observed historian Ernst Gombrich, that the development of linear perspective by Brunelleschi and associates is the sole achievement that has been consistently regarded in the category of true scientific invention.

Placemarks

Several years ago Brenda Laurel and I found ourselves along with Michael Naimark working on a video production in Zion National Park. One of the ingredients that Zion offered for our videotaping was Anasazi petroglyphs. We reverently regarded these evocative figures inscribed on the red sandstone faces of Zion as evidence of their creators’ profound spiritual connection with the land. Once taping was finished we hiked to a particularly spectacular trail that climaxed in the Angels Flight. The steep, tortuous ascent that involved clinging to chains for dear life eventually terminated in a wide ledge that afforded a panoramic view of just about everything. What arrested my attention more than the view, was the graffiti. It had spread like a virus over every square inch of stone surface that humans could reach. All the inscriptions were alphabetical, of course—just the rude monograms that a population of Kilroys leave to mark their excursions far and wide.

The question that occurred to me just then was whether the impulse that motivates people to carve unsightly initials on places might have anything at all in common with the impulse that produced those ancient environmental art works we had been admiring.

Why and how do people make/erase their marks on places? Had we overlooked some innate human proclivity that deserved to be trained and cultivated rather than discouraged? The phenomenon of placemaking—a form of behavior that I've studied more carefully since that trip to Zion—yields promising insights about how people might be encouraged to take action in a virtual environment. For Placeholder, the initial idea was that self-representation of the participants would amount to the marks they leave on things. What would be the virtual equivalent of footprints, graffiti and shadows? Planting flags on the moon or peeing in corners?

Notes on staging: The magic circle

A concept found in fairy tales and traditions of theater, the magic circle is the primordial stage—that zone differentiated from darkness by the illumination of the campfire. Consider it this way: We spend most of our lives stumbling around in the dark. Occasionally anyone wanders inside a magic circle here or there where she finds her bearings and everything suddenly falls into place. The problem is that we can never stay in such places for long, for they are not the end of the journey.

For the virtual environment of Placeholder, magic circles yielded a solution to the limited tracking range of the electromagnetic position sensing devices. The ten foot diameter of these particular circles was determined by the maximum reliable distance from the Polhemus receivers worn by participants to the transmitters mounted overhead.

However, there is no need to draw a sharp boundary between the physical and the virtual world. Why not do ambient sound with speakers, for instance, or make wind with electric fans, or create physical definition? Rather than restrain people with tethers and railings, perhaps changing the surface underfoot would be a gentler way of letting them know that they were stepping out of tracking range.

Narrative and Interaction in Placeholder

Brenda Laurel, Interval Research Corp.

Notes on conventions and constraints

Film began to emerge as an art form distinct from the technology of cinema when conventions began to be established for representing time and space. One is tempted to add to the last sentence, "in other than directly mimetic ways"—but representations are by definition distinguishable from actualities, and film could not have achieved absolute verisimilitude by its very nature. Nevertheless, the active use of the camera to orchestrate gaze and define space through attention, for instance, or the use of transitions like cuts, fades and dissolves to represent spatial and temporal discontinuities are examples of *intentionally* non-realistic treatments of time and space. Such techniques arose in order to communicate subjective experience and to serve as syntactic elements in the artistic construction of meaning. They became conventions because they were successful in forming the basis of a language of cinema that enabled artists to create works of increasing complexity and power.

Space and time

In a similar vein, one of the goals of Placeholder was to experiment with various techniques for representing space, time and distance. The three environments were separated by several miles of Canadian mountains and forests. We needed a way for participants to move among them without simulating the actual traversal of the intervening landscape. We knew of a few previous experiments in the area, conducted at the NASA Ames Research Laboratory in

the mid-80s, that used various permutations of windowing to allow people to move among unconnected spaces, but we were unsatisfied with the window metaphor, finding it too close to the visual language of computers. Our wanderings through cultural anthropology, mythology and folklore eventually led us to adopt the idea of active portals that would transport people among the worlds. Our encounters with rock art and aboriginal visual symbols brought us to the spiral as an appropriate sign for the portal.

When a person approached, the portal emitted ambient sound from the next environment to which that person might transport. Another person in the same environment might hear the same portal sounds from a distance, but upon approach, might hear the sound of another environment coming through the portal, since the destination of each person was determined individually by random choice. Within the portal, time was compressed but not absent—the duration of a transit was about 10 seconds, in darkness, accompanied by the environmental sounds coming from "ahead." People were able to see two glowing points of light representing the "grip" of each hand (if they happened to raise their hands to within their field of view), and some people seemed to use these points of light to orient themselves and maintain their balance. Many questions remain about the duration of the interval (people might be too disoriented by instant teleportation, but was the transit too long?) and the visual effects of the transit (would a neutral color with the suggestion of a flowfield better represent this metaphorical movement through space?).

Conventions relating to the body

In Placeholder we reexamined some techniques that were already becoming conventions of the VR medium. For example, as of 1993 it was standard practice to infer both direction of gaze and desired direction of movement from a position sensor mounted on the head. As people would learn this constraint, they would stop moving their heads independent of their torsos, often increasing muscle tension across the neck, shoulders and upper back. In order to "give people back their necks" without resorting to an expensive and encumbering bodysuit, we resolved to infer direction of gaze from the head-mounted sensor and to make guesses about the desired direction of movement by looking at a sensor mounted on a belt worn just below the waist, on the theory that the pelvis generally gives much more accurate information about the direction in which one intends to move than the head.

Another set of VR traditions has to do with the treatment of the hands. The assumption is that one needs to gather information as detailed as possible about the movement and position of the hand, and that in general only the right (or dominant) hand need be tracked. These assumptions arise from typical VR applications involving manipulations of remote or virtual objects in teleoperations and simulation scenarios. We determined that, for people to be able to play in the world of Placeholder, the system need only know whether a hand was touching or grasping something, and so we invented a simple device called a "grippee." A grippee was a piece of flexible plastic, held in the semicircle defined by thumb and forefinger, which used a sliding variable resistor to measure the distance between the tips of those two fingers, and a Polhemus FastTrak position sensor to define the location and orientation of the hand in three-space. We also reasoned that a person might want to use both hands, alternately or together, and that people would try different things if they had both hands available to them—so we put grippees on both hands. Since there was no instrumented glove from which one could construct a virtual hand, and since participants in the world would spend most of their time in the virtual environments in the characters of critters rather than humans, we decided that the traditional visual representation of a (wireframe or shaded-polygon) human hand would be inappropriate. Instead, each person saw two points of light for each hand as described above. These points provided unambiguous but minimal feedback about the location of the hand in space and whether the hand was open or closed.

We also questioned certain VR interface conventions using gestural language. It has become customary to move about in virtual environments by "flying," a mode invoked by pointing two fingers in the desired direction of flight (and which must also be explicitly terminated by another gesture). These gestures are formal—that is, not mimetic of any activity one might undertake in order to move or fly—and they demand a good deal of accuracy for reliable recognition. We did not wish to require people to learn formal gestures in order to use our system, and we did want them to move around fairly naturally in the environment. These desires led us to make two key decisions about movement. First, we resolved to let people walk around by walking around. Since the tracker range was only reliable for a circle of about 10 feet in diameter, we fell upon the notion of the "magic circle"—an ancient theatrical and storytelling convention—as a way to contextualize the technical constraints imposed by the system. Second, we wanted a person to be able to fly like a bird when they assumed the character of Crow. We were aware of experiments by Mark Bolas at NASA Ames in the late 80s with alternative interfaces for flying, including flapping, gliding and following a virtual paper airplane, but none of these techniques had supplanted "finger-flying" as the convention for movement. We resolved to let Crow fly by flapping his wings.²

In summary, two issues were central to our work with conventions in Placeholder. One was the definition of the medium—thinking about what it was and could be, and how conventions could be used to shape its potential. If VR is to be used as a medium for narrative, dramatic or playful activity, we should question the appropriateness of conventions derived from computer displays, teleoperations or training simulators. The other issue was the question of the interface—thinking about how people were being sensed and how they were being constrained to behave. Our motto was "no interface," expressing our desire to maximize naturalness, to enable the body to act directly in the world and to minimize distraction and cognitive load.

Narrative elements

We thought of Placeholder as a set of environments imbued with narrative potential—places that could be experienced and marked through narrative activity. When a person visits a place, the stories that are told about it—by companions, by rock art or graffiti, or even by oneself through memories or fantasies—become part of the character of the place. Stories give us ideas about what can be done or imagined in a place; learning that a particular canyon was an outlaw's hiding place, for instance, or remembering a child saying that a particular rock resembled an old woman's face will certainly influence our experience of that place. It's hard to experience a natural place without remembering or constructing some stories about it.

Motifs

Lucinda deLorimer, a professional storyteller, worked with us to uncover narrative motifs from mythology and folklore that would influence the design and representation of virtual environments. We used many of these motifs as indirect inspirations to design—making decisions about the "feel" of a given environment, for example, and more concretely, determining what the environments said about themselves. Place motifs were embedded in the virtual environments as "placemarks"—fragments of narrative spoken in the "voice of the place" (performed by an actor), emanating from sound sources located within landscape features and triggered by the proximity of a person.³

A second use of motifs was in the selection of critters and the creation of their dialogue. We wanted to populate the environments with archetypal critters with which human participants could merge. The narrative goal here was to give people character materials to play with. A not-so-obvious goal was to make humans aware of being embodied by inducing them to intentionally enter the body of a critter. We did not want the body in the vir-

tual world to be taken for granted, and we did want to explore the idea of how places look and feel different to different kinds of beings. The experience began in the cave (selected as the starting place for obvious Jungian reasons), where people encountered the critters as large petroglyphs placed around a pool. Triggered by proximity, the critters would begin to speak about themselves—their powers, characteristics and opinions of other critters. These narratives were typically based on motifs that could be found in stories from widely diverse cultures, or on stories that were improvised by actors who had been exposed to those motifs. As a person moved closer to a critter, its narrative became more elaborate and persuasive, urging the person to "come closer." Once a person's head intersected the petroglyph, he or she would join that critter, taking on its appearance, voice, perceptual characteristics and means of locomotion (see the section on smart costumes, below).

Collaboration with the Precipice Theatre Society

We were fortunate to find and form an alliance with the Precipice Theatre Society, an environmentally-oriented improvisational theatre company based in Banff. Under the direction of Colin Funk, the Precipice troupe studies pending legislation or development plans that threaten the natural environment in Canada, then develops scripts on these topics through improvisation. The company gives rollicking, commedia-style performances throughout western Canada, often meeting—and overcoming—vociferous audience disagreement with their point of view. We sought the company's help in developing the characters and interface concepts for Placeholder. They were ideal collaborators, both because of their excellent improv and performance skills, and because they were generally naive about computers and free of preconceived notions about virtual reality.

By way of preparation we gave the actors lists of story motifs that Lucinda had found about the kinds of places and critters we were using. As soon as we had settled on the locations and characters we wanted to use in the piece, we asked members of the troupe to improvise action with the characters in the actual environments. Immediately we began to see new ways in which people could play in the environments and what kinds of affordances were needed in order to facilitate such play. The actors also improvised vocal and physical characterizations of the environments. We later recreated some of the environmental vocalizations in the studio and mixed them with natural sounds to create an auditory signature for each environment that was played through the portals. The improvs were crucial in shaping the individual critter characters as well.

After the improvs in the field, we did another series of improvs in the studio, recreating the spaces through mime and vocalization, and improvising interface features like voiceholders to see what would seem "natural" to a person who was not technologically inclined. When it came time to script the critters' voices and placemarks, I took much of the dialogue directly from videotapes of these improv sessions. I cast actors from the company in the roles of the various critters and places and rerecorded the dialogue digitally. The dialogue was post-processed in some cases to apply the same critter voice filter that would be applied in real-time to a participant embodied as that critter (Crow always sounded like Crow, whether he was speaking stored dialogue or with the live

² Designing crow flight was a wonderfully interesting problem. At first, we asked everyone we could find to tell us how they fly in their dreams. We were discouraged because so many different methods were reported. But in early tests with the system, we observed that whenever someone became embodied as Crow, they inevitably flapped their arms. Viola! It remained for Rob Tow and Graham Lundgren to figure out what constituted a "flap." Rob designed the strategy for returning the flyer to the position corresponding to the current location of her body—an elegant landing that made every Crow feel like an expert flyer—and Graham produced a superb implementation of flight.

³ Due to time constraints, only a few placemarks were actually implemented, although many were scripted and recorded.

voice of a participant.) At performance time the dialogue was spatialized using the Convolvotrons.

When Placeholder opened, the Precipice troupe members were our first participants. Most of the actors were fascinated with the system and with VR. Their physical fluidity and improvisational skills made their interactions in the environments a joy to watch.

Voiceholders and voicemarks

Rachel has described our interest in how people leave marks on places. We wanted to give people the ability to "mark" the virtual environments, and we arrived at voice as a convenient modality for doing so. Voice offered several advantages over writing or drawing. Through prosody, voice permits greater expressiveness and personalization than writing; it is also more immediate. Most people are less self-conscious about speaking than about drawing. While drawing would require that we build special virtual drawing tools, capturing voice was relatively easy to implement.

Where and how could voices be stored and replayed? We designed the voiceholders as virtual record/playback devices. A voiceholder would capture and store an utterance, called a voicemark. A "full" voiceholder (that is, one containing a voicemark) would play its contents when touched. In order to encourage people to play with relationships among voicemarks and between voicemarks and landscape features, we made the voiceholders moveable, exempt from gravity and able to be placed anywhere one could reach. Voiceholders could be moved by grabbing them (closing the grippes while the points of light were "inside" the rock), dragging them to the desired location (they would stick to your hand), and releasing the grip. We wanted people to think of the voiceholders more as tools or agents than as devices; machines (including tape recorders) were inconsistent with the fantasy context. We designed them as rocks with faces, using the facial expressions to indicate the state of the voiceholder [see figure 1]. This was as close as we came to an iconic or symbolic interface element.

The voiceholders turned out to be difficult for many people to use. A major problem was performance speed. The sampling rate for the grippes' location sensors was constrained by the frame rate of the display (sometimes as low as 5 Hz). That was not fast enough to feel natural; a person's hand could pass through a voiceholder in less than .2 seconds and therefore not be sensed as a touch. People had to become conscious of and careful with their hand positions in order to activate and grasp the voiceholders; that distraction made using them more problematic and less natural than it might have been. We hypothesize that these technical difficulties, and not intrinsic design flaws, limited the amount and complexity of narrative activity and other kinds of play with the voiceholders.



Figure 1: When a voiceholder was empty, its eyes and mouth were closed and it was dark. When a person touched an empty voiceholder, its eyes would open and it would light up from the inside (rather like a jack-o'-lantern) and a voice (emanating from the voiceholder) would say "I'm listening." If a person spoke when a voiceholder was in that state, their speech would be recorded. When the voiceholder became "full" (i.e., ran out of space in the sound file it was creating) the inner light flickered and then went out (rather like a guttering candle), the eyes would close and the mouth would open to indicate that it was ready to speak. If a person touched a voiceholder in this state, it would open its eyes, light up and play back its contents, then it would "go to sleep" again.

"Smart Costumes"

The sense and status of the body in virtual space has been problematic since VR was invented. As mentioned in the section on conventions and constraints, the body has typically been highly constrained by sensing technologies and strategies as well as by the emphasis on formal gesture. The absence of haptic affordances in VR interfaces has reinforced a sense of incorporeality. The now-iconic disembodied hand that floats before one's eyes in most virtual worlds struck us as emblematic of a fundamental difficulty.

We considered two rather oblique approaches in an attempt to create a different sense of the body in VR than we had heretofore experienced. The first was an extension of the idea of place-marks—that one knows where one is or has been through evidence in the environment; e.g., shadows and footprints. Our original specification included both but neither was implemented. The second approach was based on the idea of having to do something—to take some action—in order to have a body. Mere humans were invisible in the world, to themselves (save for the points of light on the hands) and to each other. They couldn't see the portals or see the voiceholders. All they could do was talk and explore the immediate environment of the cave. From the moment people entered the world the critters were talking to them, bragging about their qualities and enticing people to "come closer." When a person's head intersected one of the critters, he or she became "embodied" as that critter. The critter, now functioning as a "smart costume," changed how a person looked, sounded, moved and perceived the world. Thus the final point of our "body politics" was to draw attention to the sense of body by giving people novel bodies.

The critters—Crow, Spider, Fish and Snake—were chosen in large part on the basis of the narrative motifs associated with them. Universality of motifs was also a selection criterion. Complementarity was another, in terms of both pleasing contrasts and potential alliances. The narrative motifs formed the spines of the critter characters, giving rise to both their graphical representations and their voices and dialogue. We hoped that the critters' traits and the things they said about themselves would give people narrative material to play with after they had become embodied.

We tried to identify characteristics of perception and locomotion for the critters that were consistent with the narrative motifs. Crow, for instance, has a reputation for admiring and acquiring shiny things; Crow's vision might boost specular reflections in the environment. In myth and lore, Spider is often characterized as being able to "see into all the worlds"—multiple points of view or levels of reality—hence representation of Spider's eight eyes (some independently steerable) seemed apt. Snake, renowned for its ability to navigate the dark landscapes of sex and death, could

see in the dark, possessing infrared vision as pit vipers do. We also tried to give each critter physical characteristics that would create unique advantages—e.g., Fish could see clearly underwater while others' vision was blurred. In the end, "snake vision" was the only perceptual quirk that we had time to implement, and the results were equivocal—people certainly knew when they were Snake, but the implementation was poor in that it simply applied a red filter without increasing apparent luminance, thus effectively reducing rather than enhancing visibility.

Despite these shortcomings, we found that the smart costumes immediately and strongly influenced participants' behaviors. Their voices and body movements became more exaggerated and dramatic. Most people were "in character" the instant they realized that they had become embodied as a critter. I suspect that the "masquerade" aspects of the smart costumes—replacing or obscuring one's identity with an exotic persona, and also amplifying aspects of one's own identity that are obscured by one's ordinary persona—put people in a frame of mind that allowed them to play, often quite boldly and imaginatively.

The Goddess

Placeholder was a performance piece in that one of the characters was improvised live—the character of the Goddess. We originally conceived her as a playmate and trickster, with the goal of enriching dramatic interaction. When people first found themselves in the cave, the Goddess spoke to them about the world (this bit of the Goddess' dialogue was taped). Unlike the environmental sounds, her voice seemed to reside in the participant's head. We had planned for her to be able to cause many things to happen—change the weather, make rocks fall from the sky, send people through portals and send her minion, the Mosquito, to pester anyone who displeased her. As with many other narrative elements, the schedule did not permit us to implement these plans. In the end, she simply spoke.

The Goddess workstation consisted of two monitors, each showing the video for one eye of each participant, headphones with corresponding sound channels for each participant and sound board controls enabling her to speak to one or both participants or to the audio and computer control rooms. The workstation was located behind glass in a booth facing the two circles, so she could also see the participants' actual bodies. This came in handy when people got in physical trouble (for instance, one little girl's helmet had slipped down over her nose) so that the Goddess could provide realtime help; it also gave her additional cues about how people were actually feeling by watching their physical bodies as well as their virtual views.

Most of the time, the role of the Goddess was performed by me or Jennifer Lewis, a research associate at the Banff Centre. It was also occasionally performed by others, including men. The Goddess' character changed according to who was performing her and also in relation to the participants. With children, she tended to behave (and to be perceived) as a helper and friend. With adult couples, she was often a cupid and a tease. She answered questions about the worlds and about the interface and coached people who were having difficulties. She often made suggestions about things to do. Occasionally, as with a pair of young men who asked one another, "Can I eat you? Can I shoot you? Well, what can we do here?"—the Goddess became downright bitchy. Our interviews with participants after their experiences revealed that people had differing reactions to the Goddess, usually well correlated with the style of her performance in their session.

Some conclusions

Although our accomplishments with Placeholder fell far short of our hopes and plans, I believe that we achieved "proof of concept" with many elements of the piece, including the idea of smart costumes, the physical interface strategy and the various techniques of environmental capture and representation that we employed. In other areas, like the voiceholders, more work must

be done to determine how we can design environments with affordances that induce the kind of collaborative narrative construction we had in mind. A strong measure of our success is the number and quality of new questions that the piece enabled us to ask.

Working on this piece has demonstrated to me that the art of designing in VR is really the art of creating spaces with qualities that call forth active imagination. The VR artist does not bathe the participant in content; she invites the participant to produce content by constructing meanings, to experience the pleasure of embodied imagination.

Technology and the Senses in Placeholder

Rob Tow, Interval Research Corp.

Technology overview

Placeholder was a two person VR system, with helmets manufactured by Virtual Research that provided both visual and auditory stereo to the participants. We added a small microphone to the helmets to pick up the voices of the users. There were two physical spaces where the participants stood wearing display helmets and body sensors, and three virtual worlds through which they could independently move. Position sensors (Polhemus "FastTraks") tracked the three-space position and orientation of the users' heads, both hands and torsos within a circular stage of about 10 feet.

An additional sensor system was employed—the "grippees," designed by Steve Saunders of Interval. These were placed in each hand, and measured the distance between the thumb and forefinger (or middle finger) of the hand, allowing the development of a simple "grasping" interface for virtual objects.

A variety of computers were used in concert in the Placeholder project. The primary computer used in the project was an SGI Onyx Reality Engine, equipped with 64M of main RAM and 4M of texture memory. It was programmed in C and Unix, using the Minimal Reality Toolkit (authored by Chris Shaw of the University of Alberta), as the primary VR framework. John Harrison, the Banff Centre's chief programmer, modified the Minimal Reality Toolkit to provide support for two users and two hands per user from its original one person, one dataglove instantiation. Shaw and Lloyd White, also of the University of Alberta, visited to help with code coordinating support for two users. Glenn Fraser and Graham Lundgren of the Banff Centre and Rob Tow provided additional programming support within the framework of the MR Toolkit.

Tow wrote the C code on the SGI which managed the audio generation and spatialization. This was coordinated with the visual VR code running in the MR Toolkit and controlled sound generation by a NeXT workstation and a Macintosh II equipped with a SampleCell audio processing card, as well as the spatialization by two PC clones, both equipped with two four-source Crystal River Engineering Convolvotrons. The NeXT, the Macintosh II and two Yamaha sound processors were programmed by Dorota Blaszcak of the Banff Centre, who was also responsible for the general audio design and integration. Blaszcak also designed the realtime voice filters which altered participants' voices to match the critters' "smart costumes." Two SGI VGX computers were used with Alias architectural design tools to lay out the environments' geometries and to apply textures to the resulting wireframes; this effort was accomplished by Rachel Strickland working with Catherine McGinnis, Raonnul Conover, Douglas McLeod, Michael Naimark and Tow. Video capture of environments was accomplished by Strickland and Naimark; subsequent video digitization using Macintosh-based equipment was done by McGinnis and Strickland, with image enhancement in Photoshop done by McGinnis and Tow. Control of the grippees was done by C and TCP programming on a Macintosh Powerbook 180 by Sean White of Interval and C and TCP programming on an SGI by Fraser.

The effort was greatly hampered by the tools used. Due to budget constraints, an old C compiler for an earlier SGI model was used instead of the proper optimized compiler for the Reality engine. Debugging was largely accomplished by "print" statements, and required a minimum of three people in realtime to do one in the VR helmet, one running the Onyx Reality Engine and one running the sound processors. The design of the worlds proceeded in a non-immersive way, on workstation screens; the projective geometry of the Alias design software differed greatly from the immersive experience in the VR helmet, which led to tedious difficulties in the world construction. The MR Toolkit itself suffered from memory leaks, which would lead to a sudden slowing down of the worlds during performance from a frame rate of eight to 12 frames a second to two frames a second. This was due to the memory demands of the textures used exceeding the size of the 4M fast multipotted texture memory, requiring paging of textures from main RAM (SGI released a no-cost upgrade to 16M of texture memory in the fall, too late for our effort).

We suffered greatly from not designing while immersed in the medium itself. One notable exception to this occurred near the end of the process, when we did do a small piece of world layout from inside the virtual environment. This was the placement of the uninhabited critter icons; we placed a set of voiceholders randomly in the worlds, then Laurel donned a helmet and moved the voiceholders to where she wanted the various critters to be—and we replaced the voiceholders with the critters. This was a small presage of what it might be like to fluidly design from within an immersive environment, as opposed to painfully and explicitly calculating coordinates at a desk.

Another painful part of the construction was the process of capturing the environments and turning them into data structures. A tremendous amount of imprecise hand work was involved, from camera positioning to wireframe design. Automating this process is clearly possible, using computer controlled cameras and such techniques as deriving information about depth from stereo imagery.

Spatialized audio

Fully spatialized audio was produced by using Crystal River Engineering Convolotrons. These are DSP subsystems that are integrated into an IBM PC clone, which acts as an audio processing server. With these, sounds are input into a Convolotron, along with a virtual location and an orientation of the user, and processed to correspond to how a sound would be perceived at such a distance—producing inverse square attenuation, atmospheric coloration (differential frequency attenuation) and the effects of the reflections and absorptions of a modeled upper body and shoulders, head and the pinnae of the ears. The Convolotrons were indirectly controlled by the main computer used to produce the visual VR—an SGI Onyx "Reality Engine"—through a connection to a NeXT computer which directly controlled both their input sound sources and the location data.

There were several sound sources fed through the Convolotrons. The first kind were the voices of the two participants from small Sony microphones we added to the Virtual Research VR helmets. This was done so that although the participants occupied differing physical spaces, we could map them into the same virtual space in a way that was coherent with their body movements within the "magic circles." A third voice, the "voice of the Goddess" (VOG), was spatialized so as to always appear to emanate within the users own head.⁴

A second class of sounds were environmental sounds, such as waterfalls, water drips, river and wind. Some of these were recorded in the field, and some were drawn from standard sound effects libraries. These were digitized, and stored in a Macintosh equipped with a "SampleCell" audio card. These sounds were produced on command from the NeXT—which in turn was commanded by the Onyx—and fed into the Convolotrons.

The most challenging environmental sound was that of a

waterfall, which we recorded at four positions—two stereo pairs, one at the base of the falls and one near the top. This was done so as to emulate a field of sound, as opposed to the usual point source like a voice or a drip.

Several waterfalls were recorded using a Sony D-7 DAT recorder and a variety of microphones. We compared conventional directional microphones, carefully shielded against spray,⁵ and a pair of Sonic Studios head-mounted microphones. The latter proved to have superior frequency response⁶ and much greater convenience and utility than the conventional shotgun microphones.⁷

The four waterfall sound sources were positioned at the four corners of the virtual waterfall. This world was composed of 30 temporally successive video fields, warped in 2x space to correspond to the actual topography of place (with a 3x vertical exaggeration)—this world did not attempt a complete embodiment in space, but rather presented a dynamic loop through time of the flow field of the moving water.

Although the Convolotrons are able to spatialize sounds with a reflective model, with up to six walls of differing sound reflection qualities, we used only the anechoic model for all of the sounds—although we had initially planned to use the reflective model for the cave world, a fully enclosed space in which echoes would have been evocative of "caveness." There were two reasons for this choice: firstly, the reflective model used much more of the computing power of the Convolotrons, reducing the number of channels; secondly, it was not completely supported in the existing client on the SGI, and would have required work to complete its interface. We used Yamaha sound processors to add reverb to sounds within the cave world. This proved evocative to most participants, but severely annoyed two professional musicians who objected to the sounds presented in the cave world in an interview following their session.

There were onerous difficulties in integrating and debugging the sounds (and everything else!) within the overall system. There were entirely too many computers involved (the Onyx, a VGX to control the Onyx because its video output went to the helmets, the NeXT, the two Convolotron systems, two Macintoshes—one running gripper server code, the other equipped with the SampleCell audio card—and the Yamaha sound processors), with three major software subsystems held by three programmers and two minor secondary systems. Running the system took a minimum of three people in realtime to manage the various controls. Debugging the spatialized sound, especially achieving coherence with the visual VR, was extremely challenging—errors sometimes went unnoticed for days. It was striking that a visual cue would often apparently "pin" the apparent sound location—but that the sound would become startlingly more lifelike when it really came from the correct location. I remember the first time that the "enticement" critter sounds were properly spatialized—and I was severely startled while "in" the cave world by a Crow voice speaking just over and behind my left shoulder at a very close distance—I jumped, turned and looked up, all quite involuntarily.

Ultimately, good spatialization was achieved for all of the sound sources. For example, the hoodoo world had as one of its continuous environmental sounds a river, positioned in the far distance, corresponding to a real river in the actual place. When users were asked to close their eyes, turn around several times and point

⁴Actually, we positioned VOG at a point six virtual inches above the head, as an annoying clipping occurred when the position was within the head and the participant moved—this was apparently a limit of the Convolotrons' model, which does not account for sounds inside a head's volume.

⁵We used non-lubricated condoms to shield the microphones, accomplishing safe audio recording in a very slippery and wet environs.

⁶As measured by a frequency spectrum analyzer program running on one of the Convolotron systems, they were flat from 40 hertz to 18 khz, with a mild dip from 18 kilohertz to 22 khz.

⁷At Brenda's direction that I record the sound of putting one's head into a waterfall, I wore these into a 4°C twenty meter high waterfall under a rain poncho...twice, since the first time I did it, I fell and accidentally turned the recorder off.

at the river, they were able to point accurately at the apparent location.

In conclusion, I think that the use of multiple channels of spatialized sound in an environment that was integrated with the kinesthetics of participants' bodies in space was powerfully evocative and pleasant. It provided affordances both for navigation in space and for locating things that had action in the world beyond visual boundaries; it directed attention and it enhanced conversation.

Some notes on perceptual issues

The principle of action

The movement of the body in space creates changes in the information impinging on the ears and the eyes. These changes are very important in building the awareness of what is out "there." The psychologist James Gibson elucidated this idea for vision in his 1979 book *The Ecological Approach to Visual Perception*. We were informed by this principle in the construction of Placeholder, both in vision and in audition.

Sound is difficult to perceive as being located at a particular place in three dimensional space when the head and body are immobilized relative to the sound source, or when there is no correlation between body movements and changes in the sound. In the real world, when the head and the body are allowed to move and change position relative to an object that makes sound, quite accurate judgments as to the direction and distance of sound sources may be made. Ordinary stereo heard through headphones does not take into account the differences in sound caused by the movements of the body in space.

Remarkably, a white noise source, which consists of random sounds of all frequencies and all phases, may be localized in space by a person—but only if they are able to move their head. Body movements help, too. This is because the movement causes changes in the frequency mix; some frequencies are attenuated or diffracted by the head, the pinnae of the ear and the torso more than others; this relative change in the mix allows the perceiver to deduce where the sound source is located despite the pure randomness of its content.

In Placeholder we made use of technologies which track the position and orientation of the head and body, and which produce synthetic changes in vision and sound which correlate with the movement of the body. Visually, this was done with the SGI Onyx Reality Engine, which computed what each eye would see from their separate positions in space; auditorially, this was accomplished with CRE Convolvotrons, which process sounds to produce the effects of the reflections, attenuations and diffractions caused by the head, torso and pinnae of the ears, and the delays between sound arriving at the two ears separated in space. These disparate systems were tightly coupled in Placeholder to produce a high degree of coherence between vision and audition.

Achieving a sense of place requires a degree of sensitivity to sensory cues and combinatorics. For example, salient verticals and horizontals—such as trees, falling objects and horizons—must be coherent with each other. A lack of coordination of such cues rapidly produces disorientation, sometimes results in "simulator sickness," and may even make people fall down if they are in a walking VR environment. Considerable attention in Placeholder was devoted to these issues, and some world constructions were discarded because of such difficulties in achieving coherence. Each of the three final worlds had strong cues of this nature; the horizon line and visible trees in the hoodoo world; the falling water and the flat sharp edged canyon floor in the waterfall world; and the flat floor in the cave world. Additionally, the orientation of the petroglyphs provided information about verticality.

A strong sense of the body in a place was provided by the spatialized sound sources. When a person stood in the center of the circle, there were sound sources in all surrounding directions, which could be accurately localized. Sound sources varied in dis-

tance from zero in the case of someone putting their head in the waterfall near a corner, to two to 10 feet for a critter icon or a voiceholder or a placemark, to hundreds of feet for the wandering wind sound and the distant river; walking about quickly provided a robust feeling of spatial extent. The other participant's voice and its correlated movement with the movement of the embodied critter icon provided cues for distance and size of the space—which could be quite extensive in the case of watching and listening to Crow fly away above the waterfall to an apparent distance of hundreds of feet. In the waterfall world, the four-point sample grid which modeled the sound field of the falling water was powerfully felt to be coextensive with the moving visual flow field of the water—into which participants invariably tried to put their heads. The voiceholders were small objects that could be walked around, peered at from more than one direction, manipulated (albeit clumsily) and which emitted sounds which were coherent in apparent locality and temporality with visual location and appearance—in this regard they had the greatest degree of multisensory physicality as objects.

These results of the action of the body in space and correlations with changes in the sensorium—"the principle of action"—mark the major defining characteristic of immersive VR as a medium. It is fundamentally different from the older technologies of television or cinema or stereophonic music in exactly this regard. High resolution is less important than tightly coupled coherent action in the sensorium resulting from the participant's action. Adding low or medium resolution affordances in different senses or modalities that are coherent in their combinatorics and which follow the principle of action—like adding spatialized sound to stereo video—results in a greater sense of "immersion" than does ultra-high resolution high frame rate cinema passively viewed by mass audiences. These latter expressions of the same basic technologies of wide angle stereo visual display and of multichannel sound, which do not present to individual participants the results of their body acting in the synthetic world, are more similar to movies than to what we strove for in Placeholder—and have subtle political implications for the social constructions of body and self.

Reconstructing the body

An intent in the design of Placeholder is to cause participants to become more aware of what it is to be an embodied human.

We sought to problematize issues around body and gender in *the realm of the senses*—in studied contrast to the usual literary post-modern deconstructionism, which negates the visual sense and insists on the primacy of text, and results in a profound disembodiment of cognition and feeling.

Our approach was to remove direct visual evidence of the participant's primate body and substitute iconic non-mammalian representations that moved through space in accordance with their actions, and by a series of sensory transformations. Auditorially, we achieved this by distorting participants' voices in ways that were artistically inspired by the various critters. One result of this was to render them difficult to identify according to gender. Visually, only one of the planned sensory transformation was implemented (Snake's "infrared" vision), due to time constraints, although a series of visual mappings were designed, each loosely based on the actual psychophysics of each animal. For Crow, these would have included double foveation in each eye, increased specularly of reflections and the fading from vision of stationary elements of the visual field. For Spider, we intended to emulate the vision of the jumping spider, with its multiple eyes of differing resolutions and spatial extents, merged into the single extent of the VR helmet's presentation. For Fish, we wished to provide sharper vision under water contrasted with blurry vision out of water, combined with a gradual fading to black when out of water—countered by a "reviving" when returned to water. For Snake, we had hoped to emulate the low resolution infrared sense that pit vipers such as rattlesnakes enjoy by virtue of the IR sensing pits

located on their snouts; this would have provided a spatially low resolution but bright image in the cave world—we implemented a quick hack of this that was underwhelming.

This effort was incomplete, and highly tentative; many questions remain.

Sources

- Anderson, Sherry Ruth and Hopkins, Patricia. *The Feminine Face of God*. New York, Bantam Books, 1991.
- Baring, Anne and Cashford, Jules. *The Myth of the Goddess*. London, Viking Arkana, 1991.
- Blauert, Jens. *Spatial Hearing*. Cambridge, MA: MIT Press, 1983.
- Bogert, Charles Mitchell. *Sensory Cues used by Rattlesnakes in their Recognition of Ophidian Enemies*. New York Academy of Sciences, 1941: 329-344.
- Brown, Joseph Eps. *The Spiritual Legacy of the American Indian*. New York, Crossroad, 1982.
- Buser, Pierre and Imbert, Michel. *Audition*. Cambridge, MA: MIT Press, 1992.
- Buser, Pierre and Imbert, Michel. *Vision*. Cambridge, MA: MIT Press, 1992.
- Campbell, Joseph. *Historical Atlas of World Mythology, Volume I: The Way of the Animal Powers*. New York, Harper & Row, 1988.
- Chatwin, Bruce. *The Songlines*. London, Cape, 1987.
- Cowan, James. *Mysteries of the Dreaming: The Spiritual Life of Australian Aborigines*. Bridport, Dorset, UK, Prism Press, 1989.
- Darwin, Charles. *The Expression of the Emotions in Man and Animals*. St. Martin, 1980.
- de Angulo, Jaime. *Indian Tales*. New York, Hill & Wang, 1953.
- Fisher, S.S., Wenzel, E.M., Coler, C., McGreevy, M.W., "Virtual Interface Environment Workstations" presented at Human Factors Society 32nd Annual Meeting, October 24-28, 1993, Anaheim.
- Gibson, James. *The Ecological Approach to Vision*. Boston: The Houghton Mifflin Company, 1979.
- Gimbutas, Marija. *The Goddesses and Gods of Old Europe: Myths and Cult Images*. Berkeley, CA: University of California Press, 1974.
- Gimbutas, Marija. *The Language of the Goddess*. San Francisco, Harper & Row, 1989.
- Gombrich, E.H. *Art and Illusion: A Study in the Psychology of Pictorial Representation*. The A.W. Mellon Lectures in the Fine Arts, National Gallery of Art, 1956. Princeton University Press, Second Edition, 1961.
- Highwater, Jamake. *The Primal Mind: Vision and Reality in Indian America*. New York, Harper & Row, 1981.
- Higuchi, Tadahiko. *The Visual and Spatial Structure of Landscapes*. Translated by Charles Terry. Cambridge, MA, MIT Press, 1983.
- Klauber, Laurence Monroe. *Rattlesnakes, their Habits, Life Histories, and Influence on Mankind*. Berkeley: University of California Press, 1982.
- Land, Michael. "Vision in other Animals," *Images and Understanding*, edited by Horace Barlow, Colin Blakemore and Miranda Weston-Smith. Cambridge, England: Cambridge University Press, 1990.
- Laurel, Brenda. *Computers as Theatre*. Reading, MA: Addison-Wesley Publishing Company, 1991; paperback edition 1993.
- Lopez, Barry. *Crossing Open Ground*. New York, Charles Scribner's Sons, 1988.
- Lopez, Barry. *Giving Birth to Thunder, Sleeping With His Daughter: Coyote Builds North America*. Kansas City, MO, Andrews & McMeel, 1977.
- Marshak, Alexander. "An Ice Age Ancestor?," *National Geographic*, October 1988, pp. 478-481.
- McKenna, Terence. *The Archaic Revival*. New York, Harper Collins, 1991.
- McLuhan, Marshall and Parker, Harley. *Through the Vanishing Point*. New York, Harper & Row, 1968.
- Montessori, Maria. *The Discovery of the Child*. New York, Ballantine Books, 1972. ©1948.
- Naimark, Michael. "Elements of Realspace Imaging: a Proposed Taxonomy," *SPIE/SPSE Electronic Imaging Proceedings*, vol. 1457, San Jose, 1991.
- Naimark, Michael. "Presence at the Interface, or Sense of Place, Essence of Place," *Wide Angle*, vol.15, no. 4, Ohio University School of Film, 1994.
- National Geographic*. "The Peopling of the Earth," "In Search of Modern Humans," "An Ice Age Ancestor?" and "Treasures of Lescaux Cave," vol. 174, no. 4, October, 1988.
- Norberg-Schulz, Christian. *Genius Loci: Towards a Phenomenology of Architecture*. New York, Rizzoli, 1980.
- Putman, John. "The Search for Modern Humans," *National Geographic*, October 1988, pp. 438-476.
- Rigaud, Jean. "Treasures of Lescaux Cave," *National Geographic*, October 1988, pp. 482-497.
- Sams, Jamie and Carson, David. *Medicine Cards: The Discovery of Power Through the Ways of Animals*. Santa Fe, Bear & Company, 1988.
- Taylor, Rogan. *The Death and Resurrection Show*. London, Anthony Blond, 1985.
- von Uexküll, Jakob. "A Stroll Through the World of Animals and Men: A Picture Book of Invisible Worlds."
- Walls, G. L. *The Vertebrate Eye and its Adaptive Radiation*, New York: Hafner, 1963.
- Wise, David H. *Spiders in Ecological Webs*. Cambridge, New York: Cambridge University Press, 1993.

Brenda Laurel, Rachel Strickland, Rob Tow, Interval Research Corp., 1801 Page Mill Rd. Bldg. C, Palo Alto, CA 94304; (415) 354-0900

