

Disney's Aladdin: First Steps Toward Storytelling in Virtual Reality

Randy Pausch¹, Jon Snoddy², Robert Taylor², Scott Watson², Eric Haseltine²

¹University of Virginia

²Walt Disney Imagineering



Figure 1: A Guest's View of the Virtual Environment

ABSTRACT

Disney Imagineering has developed a high-fidelity virtual reality (VR) attraction where guests fly a magic carpet through a virtual world based on the animated film "Aladdin." Unlike most existing work on VR, which has focused on hardware and systems software, we assumed high fidelity and focused on using VR as a new medium to tell stories. We fielded our system at EPCOT Center for a period of fourteen months and conducted controlled experiments, observing the reactions of over 45,000 guests.

contact author: Randy Pausch, Computer Science Department, Thornton Hall, University of Virginia, Charlottesville, VA 22903. Pausch@virginia.edu, 804/982-2211

Permission to make digital or hard copies of part or all of this work or personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1996 ACM-0-89791-746-4/96/008...\$3.50

Riders filled out an exit survey after the experience, and with select groups we used a number of other data-gathering techniques, including interviews and mechanically logging where guests looked and flew.

Our major finding is that in a high fidelity VR experience, men and women of all ages suspend disbelief and accept the illusion that they are in a different place. We have found that in VR, as in all media, content matters. Novices are unimpressed with the technology for its own sake; they care about what there is to do in the virtual world. We can improve the experience by telling a pre-immersion "background story" and by giving the guest a concrete goal to perform in the virtual environment. Our eventual goal is to develop the lexicon for this new storytelling medium: the set of communication techniques shared between directors and the audience. We conclude with a discussion of our second version of the Aladdin project, which contains a large number of synthetic characters and a narrative story line.

INTRODUCTION

Most existing work on virtual reality (VR) has focused on hardware and system software [1, 3, 5, 6, 7, 10, 12, 23, 24]. The price of a high quality system has placed it out of reach for most people interested in content. Building high quality, low cost VR systems is important, but we believe the exciting challenge in VR is learning what to do with the medium.

We believe that the content questions are the really hard ones. The goal of this project has been to allow the content producers, or authors, to assume the existence of satisfactory technology and to focus directly on authoring in the new medium of VR. We produced high-quality content based on flying a magic carpet in the animated film “Aladdin” [2]. Figure 1 shows a screen shot from the system.

We field-tested the system on over 45,000 guests at EPCOT Center. In this paper we report our detailed observations, the guests’ exit surveys, and data we recorded during guest experiences. This is not a systems implementation paper; we describe the hardware and software only as context for describing the guest experience. In addition to guest experiences, we also describe industrial design solutions to the problems of high volume usage. In early 1996, we will deploy a second version with a narrative story line and a large number of reactive characters. We conclude with lessons learned from creating virtual environments and characters for our second version, especially controlling the narrative in an interactive medium.

Our underlying premise is that VR is a new medium, as film, radio, and television once were. As motion pictures matured, directors and audiences developed a lexicon including close ups, cross cuts, flash backs, etc. Over time a common language, or lexicon, will evolve for VR; this project is our first step towards that goal.

SYSTEM DESCRIPTION

In each of our field trials, four guests donned head-mounted displays and piloted a flying carpet. Because they were running on separate systems, the pilots could neither see nor interact with each other.

We designed the system for robustness, high volume usage, and high accessibility. Unlike research setups, theme park equipment is used extensively, continuously, and abusively. Failures with a one-in-a-million chance of happening can occur once a week in a typical theme park attraction.

The Head Mounted Display

The system used an internally developed head-mounted display (HMD), shown in Figure 2. The two main design constraints were to provide high image quality and to make it easy to put the HMD on quickly, to support the high throughput of guests in a theme park attraction. In early trials we learned that having adjustments such as a focus knob on the HMD confused guests, since they had no baseline to distinguish between high and low image quality. Therefore, we designed a system that would accommodate a large variation in where a guest’s eyes sit with respect to the optics.

Figure 2: The Head Mounted Display

Image quality considerations drove us to use CRTs instead of LCDs, a tradeoff that increased the HMD’s weight and extended its center of mass. We partially compensated for this by providing spring-suspension of the HMD from the ceiling. Major design challenges in the HMD included avoiding visible pixel boundaries, obtaining high contrast, minimizing inter-ocular rivalry, and addressing the weight balance and packaging issues. For head-tracking, we used a magnetic position/orientation tracker.

Unlike many other VR systems, our HMD display was bi-ocular, not stereo. We rendered a single, horizontally wide graphics window and fed partially overlapping view windows to each of the CRTs in the HMD. For applications such as ours, stereoscopy is surprisingly unimportant as a depth cue [8].

We addressed hygiene issues by having the HMD snap onto a per-guest inner “cap” that can be cleaned separately. The inner liner also allowed us to adjust tightness to each guest’s head before monopolizing the more expensive HMD and image generator. The HMD fit comfortably over eyeglasses; the only notable issue was guests with hair tied in buns.

Sound

The HMD contained two speakers that rested close to, but not in physical contact with, the guest’s ears. We used a combination of stereo ambient sound, binaural recorded sound, and eight channels of localized sound. We recorded the binaural sound track via a high quality binaural head (essentially, microphones placed in a mannequin head). The binaural soundtrack included background voices, animals, and other “clutter” sounds. We recorded multiple binaural tracks, and then mixed those layers to form a composite recording. When the binaural recording was played during the VR experience, even though those sounds “moved with the head,” they established a believable background sound field. It is in this context that the eight special channels were convolved to localize in real-time based on head tracking [26]. The localized channels provided main characters and large sound

effects. The stereo sound (primarily music) established emotional context, the binaural sound established the believable three-dimensional space, and the localized sounds gave strong, specific cues for orientation. The three levels increasingly traded recording quality for localization, and the binaural and localized sounds worked well together because they employed the same head-related transfer functions [4].

Seating, Controls, and Motion Base

Guests were seated straddling a motorcycle-style seat as shown in Figure 3. A benefit of this design is that the guests were firmly grounded, with weight on their buttocks, knees, and feet. Additionally, this design accommodated a wide range of heights. Guests gripped a steering mechanism representing the front of a magic carpet. Turning left or right controlled yaw of the carpet, and tilting controlled the pitch of the carpet. Imagine a car's steering wheel; pulling the top of the wheel toward the driver pitched the carpet up, pushing it pitched the carpet down. Pushing the entire mechanism forward controlled velocity. Figure 4 shows a schematic diagram of the carpet controls.



Figure 3: The Physical Setup




Figure 4: Schematic Diagram of Carpet Controls

We mounted the seat on a movable base that pitched up and down in response to the steering control. Originally, the motion base also tilted side-to-side, but this caused discomfort during early testing so we removed the side-to-side tilt. Surprisingly, the presence or absence of a motion base had no substantial effect on guest satisfaction, or anything else we measured with exit surveys.

An early version of the system simulated wind with a rate-controlled fan blowing air over the guests. Much to our disappointment, most guests wearing the HMD did not notice it.

Image Generation

For each guest, we used a custom Silicon Graphics computer with 512 megabytes of RAM, 16 megabytes of texture memory, eight 150 MHz MIPS R4400 CPUs and three Reality Engine pipelines with four RM5 raster managers each. We rendered 20 frames per second on each pipe, interleaving the frames to achieve a 60 Hz frame rate. Although the frame rates could vary between 15 and 60 during a flight, the overwhelming majority of the time the system rendered at 60 Hz.

Because hardware lighting can draw attention to edges in models with low polygon count, our artists decided to render all polygons with hand-painted textures, with no hardware lighting. This also improves rendering time slightly, but we did it for image quality, not speed.

Model Management And Show Programming

A custom software system, called the *player*, provided scene management and character animation playback. The player provided a Scheme interface on top of a C/C++ layer, all on top of SGI Performer [19].

The player used Performer's support for multiple levels of detail. Unlike a flight simulator, most of our scene was close, so we used only two levels of detail per object. Artists created both models for each object because degrading a model "by hand" still produces better results than automatic means [20]. We sometimes used large texture "flats" for distant objects, and switched to three dimensional models as the guest approached.

Programming of various show elements, such as an object's reaction when hit by the carpet, was performed in the topmost layer of the player, a locally developed "Story Animation Language." This SAL layer implemented cooperative lightweight threads on top of Chez Scheme [9], an incrementally compiled Scheme.

In our second version, the database is much larger, and is partitioned into distinct scenes. The player software pre-fetches geometry and texture maps as guests fly from one scene to another [11]. Between scenes, we include explicit "transition areas," such as hallways and caverns. Transition areas have a smaller number of polygons, which buys us time to pre-fetch textures. Transitions bend and twist, thus ensuring that at no point can the guest quickly look back to the previous scene.

GUEST SELECTION

We deployed the system at EPCOT Center in Orlando, Florida, from July 1, 1994 until September 8, 1995. Every twenty minutes a group of up to 120 guests was given a brief technical lecture about VR followed by a demonstration where four guests were selected to "fly a magic carpet."

The attraction was intentionally hidden in a remote area of the park. Most guests entered not because they had a strong interest in VR, but because our attraction was "the next thing" to do. Guests could not volunteer to fly; they were selected by the ride operators. The operators maintained a strict policy of avoiding guests who showed an active interest in VR. Therefore, rather than pertaining to a small subset of VR enthusiasts, we believe that our results are essentially a fair cross section of the theme park population. Some guests did decline the invitation to fly. Interviews revealed this was primarily due to stage fright, not an aversion to trying VR.

The selected pilots did not hear the technical lecture about VR. We gave them a background story that they would be stepping into the feature film "Aladdin." We instructed them that the main goal was to have fun, but that they were also searching for a character from the film. The marketplace scene was chosen because it 1) contains familiar objects such as doors which establish scale, 2) is a brightly lit daytime scene, and 3) contains wide variety, encouraging exploration. There was typically time for a one-to-three minute practice flight followed by a few minutes of rest before the audience entered and the four minute flight began.

NOVICES' EXPERIENCES

We exposed a large, non-self selected, population of guests to a high-fidelity virtual experience in a controlled environment. At least one other system has exposed large numbers of novices to VR [25]. However, Virtuality's users were self-

selected. Their users wanted to try VR, and paid for the experience. Our sample is much more diverse.

Our findings are drawn from a variety of sources, including written post-flight guest surveys, logged flight data, extensive conversations with the day-to-day attraction operators, observations of guests' flights, and interviews of guests before, during and after their flights.

Technologists should be aware that most guests were not impressed by the technology itself; guests assumed VR was possible and had an expectation of extremely high quality. Many had seen the "holodeck" on Star Trek, and expected that level of quality. Once in a virtual environment, guests focused on what there was to do in the virtual world – content matters!

General Observations

We were able to sustain the illusion that the guests were in another place. Men and women of all ages suspended disbelief and a large number reported the sensation that they were "in" the scene. This is hard to conclude from exit surveys, but guests also provided unsolicited cues, such as panicking or ducking their heads as they approached obstacles.

Guests cared about the experience, not the technology. Most guests had no concept of how VR works, nor did they care. They focused on the sensation, which was exhilarating for most guests. Many guests shouted "Wow!" or "Whee!" in their first thirty seconds.

The experience was overwhelming. Between sensory overload and the task of trying to control the carpet's flight, many guests were so cognitively taxed that they had trouble answering questions early in their flights.

Guests needed a goal. If not given a specific goal, guests would ask "What should I be doing?"

Guests needed a background story. We found that giving as much context as possible about the scene helped reduce the severity of the transition from the real to the virtual environment. *Background story* is the set of expectations, goals, major characters, and set of rules that apply to the virtual world. Ironically, in lower fidelity, less believable VR systems, this need for background story may not be as evident. We believe it is the abrupt transition into a *believable* virtual world that is problematic. Performing a good transition from the real to the virtual world is an open challenge.

Guests liked exploring, and seeing new spaces. Most guests did not spend much time studying detail in a given place; they tended to move on quickly to new vistas.

Guests did not turn their heads very much. This could be because they were piloting a vehicle, or because they were not accustomed to being able to turn their heads when looking at displayed images. For many, we believe the latter. Guests often watched characters "walk out of frame," as would happen with television or movies. Our strongest indication came from many pilots where we waited 90 seconds into their flight, then explicitly told them to turn their heads. At that point, they clearly had the "aha" of the head-tracking

experience. While we suspect that different content would be more conducive to head turning, head tracking is far enough from most guests' experiences with film and television that we suspect this will be a problem for many systems.

Controlling the carpet was a problem for many guests. This prompted the addition of test flights before the show began. Many guests flew out into the desert or up above the city to find a space where there were fewer obstacles, making flight easier. Although we could have had the magic carpet fly itself, our surveys indicated that the control and freedom are important parts of the experience. Six-axis control is a very difficult problem and an important design challenge is finding appropriate control constraints.

VR must be personally experienced. In addition to the 45,000 guests who piloted carpets, we had over one million audience members who observed the pilots' progress on display monitors. The audience members enjoyed the show and understood that *something* fascinating was going on with the pilots, but it was clear that VR is foreign enough that most people can not fully comprehend it without direct personal experience. Audience members often asked if the pilots could see or interact with each other.

Presence and Immersion

Although it is difficult to formally measure, we believe that most guests suspended disbelief and had the experience of being in a new place. Our choice of an animated world underscored that believability is different from photo-realism. In fact, we reject the term "simulation," as we provide an experience not possible in reality. Our virtual environment was not realistic, but it was consistent with the large number of animated worlds that guests had seen before. Guests flew, but had no fear of heights; guests reacted to the characters, but were not afraid of a guard who brandished a sword. In many ways, this environment was compelling without being disturbing.

A common sight in a 3-D theater is to see large numbers of guests reaching out to grab the projected image. We speculate that they are compelled to conduct this test because their perceptual and cognitive systems are in conflict; their eyes tell them the object is within arm's length, but their brain tells them it is just a projection. In our system, we saw no evidence of the need to test. Guests did not intentionally run into objects to see if the objects really existed. In fact, guests did the opposite, often involuntarily ducking when they felt they could not avoid a collision.

In general, we believe that the need for high fidelity can be reduced by engaging the user in a complex, real-time task. For example, the desktop DOOM game [14] and the SIMNET tank simulator [18] both get users to the point where the interface becomes transparent and the user focuses on task performance, which requires a sense of presence. Our system did so with the mildest of tasks, that of searching for a character. At first, we suspected that the difficult task of piloting the carpet might lower our fidelity requirements. Therefore, we ran experiments where the carpet flew itself. During those tests guests achieved the same suspension of disbelief, with the only task being to look around. Our metric for suspension of disbelief was their reactions to the environment, such as ducking when flying near objects.

What produced the effect of immersion is difficult to know. Even for guests who did not turn their heads much, the HMD physically blocked out the real world. Sound was also very important, as many guests remarked that the sound of wind when they flew high, or the crashing noises when they ran into walls strongly reinforced their sense of being there. In post flight interviews, guests told us that their illusion of presence was destroyed when characters did not react.

Reaction to Virtual Characters

It is more difficult to build a believable character than a believable scene. Although our major focus was on building the environments, we were pleased that a few of our guests did respond to characters. The show began with instructions from a parrot who told the pilots to nod their heads. Some guests actually heeded his command. Another character covered his head and shouted "Don't you have a horn on that thing?!" when guests flew near him. Many guests shouted back at this character. One young girl finished the attraction in tears because she had spent several minutes attempting to apologize to him, but instead continually triggered hostile responses whenever she approached him. (All the characters had a small set of dialog sequences that could be triggered).

The key to a successful character is the suspension of disbelief; one must talk to the puppet, not the puppeteer. Most guests flew at high speed, zooming past the characters. When guests did slow down, they expected the characters to respond and were very disappointed when the characters did not. At the very least, characters should orient their heads and eyes and look at the guest. Our next system is incorporating this feature.

We suspect that the limited believability of our first system's characters is due to low fidelity. All characters in the first show, such as those shown in Figure 5, were animated with motion capture, where sensors recorded an actor's body motions in real time, and those values were used to drive the animation. Our second version uses higher quality key frame animation. While testing of the second version is not yet complete, early indications are that we will cross a fidelity threshold in character animation much as this project crossed one in environment fidelity.

Figure 5: Animated Characters

Men vs. Women

One of our original objectives was to discover whether VR appealed only to the narrow (primarily young male) video game market, or was more like feature films, appealing to males and females of all ages. While *content* will still matter, the technology itself did not turn away any guests. On post flight surveys, the reaction of both genders and all age groups was almost identical on all questions. One major difference was that many women are afraid that they would not be able to operate the equipment properly. This surfaced both as a pre-flight concern and as a post-flight comparison. They often asked how they performed relative to the other pilots. Also, during in-flight interviewing men were more likely to talk about the technology, whereas women were more likely to talk about the experience and emotional impact. Neither men nor women complained about having to wear the HMD.

VR for the Disabled

Everyone involved with the project noted the impact on both the pilots and the audience when motion-impaired guests flew. Accessibility is a fundamental design constraint at Disney parks, and we have a substantial wheelchair population. One of our four stations could be converted for wheelchair access in about ten seconds, and we had several wheelchair fliers per day. The sense of mobility and the joy it brought them was overwhelming.

Motion Sickness

We did not find motion sickness to be as significant an issue as we had feared. During selection, we asked guests if they were prone to motion sickness, and warned that they might feel motion sick during the experience. We also told them they could stop at any time and remove the HMD. Post flight surveys indicated that, as with many theme park attractions, some guests reported discomfort or dizziness, but they mostly described it as a mild sensation. We do not know if guests who felt discomfort or dizziness self-limited their head motion; our logged data showed no such correlation. Reports of discomfort went up when the room was warmer, which is consistent with discomfort reports from platform-based simulator rides. We were careful to limit the duration of the experience. As with any “thrill” experience, discomfort increases with ride length.

GUEST POST-RIDE SURVEYS

After their flights, we asked guests to complete a one page survey with about five multiple choice questions. Guests were identified on the survey only by first name, and over 95 percent of the guests completed a survey. Most who declined did so because of low English skills. We asked many questions and report here a relevant subset. Our sample was 48.5 percent female, and included all ages.

We tried to ask questions that would yield different responses by gender and age. However, we were unable to design questions where the responses were not reasonably consistent across all groups. Thus, we conclude that VR experiences have broad appeal. Responses are presented here by gender (M = male, F = female); breakdown by age is equally similar. The possible responses are listed in the same order as they

appeared on the printed survey form. Because we made ongoing changes to the surveys, the number of total responses to any question is variable -- after each question is the total number of responses.

What did you LIKE the most? (N=25,038)

	all	M	F
characters	11%	10%	12%
helmet fit	4%	4%	3%
motion	32%	32%	32%
picture quality	17%	19%	15%
sound	8%	7%	9%
steering control	21%	21%	21%
town	7%	7%	7%

What did you DISLIKE the most? (N=22,479)

	all	M	F
characters	5%	6%	4%
helmet fit	20%	20%	20%
motion	13%	14%	13%
picture quality	13%	13%	13%
sound	6%	6%	6%
steering control	34%	33%	36%
town	8%	8%	7%

Guest rating of the Experience (N=1,903)

	all	M	F
terrible	1%	1%	1%
okay	4%	4%	5%
good	11%	9%	13%
great	54%	49%	57%
best thing at Disney	23%	28%	20%
best thing in my life	7%	9%	5%

As an absolute answer, we take this with a grain of salt. It is unlikely that our system is really the best thing in seven percent of our guests' lives. However, the scale is useful for comparing males and females; again, we found an overwhelming similarity.

Would You Recommend it To a Friend? (N=273)

	all	M	F
yes	99%	100%	98%
no	1%	0%	2%

It Made Me Feel Like I Was... (N=1,336)

	all	M	F
visiting a town	14%	15%	14%
playing a video game	23%	19%	25%
being inside a movie	45%	49%	43%
in the middle of a dream	17%	16%	17%
invisible	1%	1%	2%

Had You Heard About Virtual Reality Before Today? (N=307)

	all	M	F
no	16%	12%	18%
I had read about it	36%	37%	34%
seen on TV or movies	49%	50%	47%

On My Next Ride, I Would Most Like To... (N=324)

	all	M	F
--	-----	---	---

see more characters	35%	32%	40%
see more towns/places	38%	37%	38%
see the other pilots	27%	31%	22%

The Best Thing About it Was... (N=439)

	all	M	F
the characters	5%	3%	6%
flying	42%	41%	43%
exploring/seeing new things	23%	23%	23%
being able to go where I wanted	30%	33%	28%

I would most like to... (N=426)

	all	M	F
have the carpet fly itself	9%	5%	12%
fly the carpet myself	84%	90%	80%
ride while a friend is flying	6%	4%	8%

LOGGED DATA

For over two thousand guests we recorded the position and orientation of the pilot's head and the carpet twenty times each second. Our original hope was that we could see patterns of where guests flew and what they found interesting. In fact, we discovered that guests flew almost indiscriminately; no obvious patterns of travel emerged from the data.

The analysis of head turning data was more interesting. Our first question was "How much do guests turn their heads?" The data confirmed what many researchers describe as the dirty secret of VR. In many scenarios, people in HMDs do not turn their heads very much. Figure 6 shows a top-view polar histogram of head yaw; for a guest facing right, the length of each line shows the proportional amount of time spent at each angle. Figure 7 shows a conventional histogram of guest head yaw; the height of each bar is the portion of total time spent at that angle.

Figure 6: Polar Histogram of Head Yaw




Figure 7: Conventional Histogram of Head Yaw

Figure 8 shows that the difference between male and female head yaw is negligible. In fact, every category that we examined (gender, age, which lab technician instructed them, whether or not they experienced motion discomfort, how much they enjoyed the ride) yielded essentially the same profile.

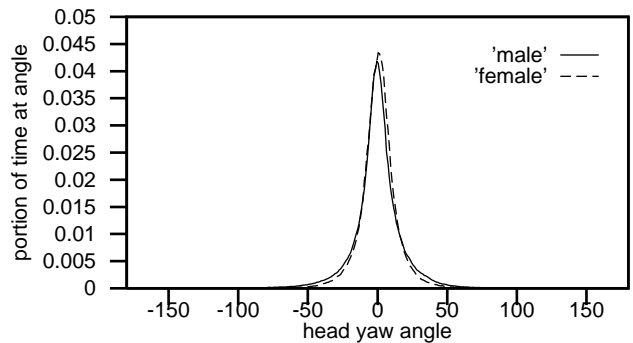


Figure 8: Male vs. Female Head Yaw

Figure 9 shows that head pitch, or up/down tilt, is even more confined than head yaw.

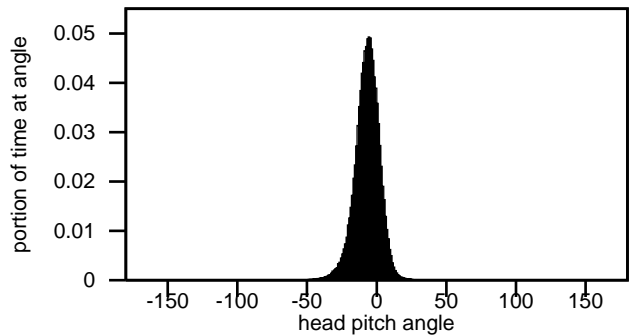


Figure 9: Head Pitch

We were not surprised that guests looked straight ahead most of the time. However, we were surprised by the following: instead of portion of total time, Figure 10 graphs the widest yaw angle guests *ever* experienced. One way to read this graph is that 90% of the guests never looked more than 75 degrees to either side. One could infer that building a 150 degree wide, screen-based display would be as good as an HMD for 90% of the guests. That conclusion would ignore that the HMD field of view must be added to the head yaw, and that the HMD also prevents visual intrusion from the real world.

TELLING STORIES IN VR

Given that VR can present a compelling illusion, researchers can and should pursue its uses for education, training, medical applications, games, and many other purposes. As a storytelling company, we are focusing on using VR as a storytelling medium.

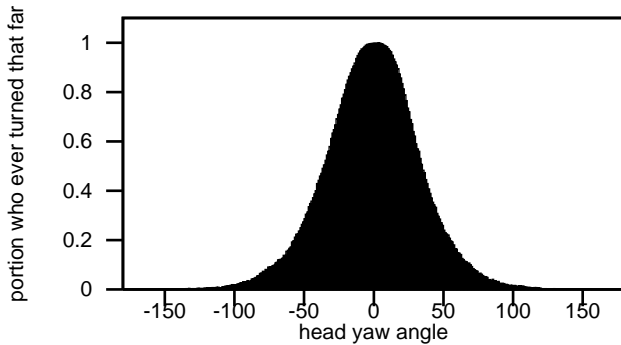


Figure 10: Widest Yaw Angles Seen by Various Guest Percentages

Script vs. Guest Controlled Cameras

Our first system was the first of many steps towards telling stories in VR. Our next show contains over twenty scenes, approximately fifteen high-fidelity characters, and a narrative story line that includes the ability to alter the sequence of events. Our guest assumes a role in an immersive feature film. The major challenge of allowing the guest to become a character is that the director gives up control of the narrative. While this is true of many interactive, non-HMD based games, the problem becomes acute with an HMD.

Because we let the guest control the viewpoint we must build characters and scenes that look good from all vantage points. By establishing entrances to scenes we control the initial view of each scene, a technique used in well-designed theme parks. The inability to cut from scene to scene or view to view is very frustrating for content authors. We have experimented with having characters that are attached to the guest's head, and appear to be hanging off the front of the HMD. This allows us to interject a brief "scene" including that character.

There is an intrinsic conflict between a pre-constructed narrative and a guest-controlled exploration. An interactive system can dynamically re-configure the story to avoid

omission of critical portions. As our director said, "It's as if you decide to leave a movie early, and the projectionist edits the film to make sure you see the important ending before you leave."

In other perceptually intensive theme park attractions such as effects-laden stereoscope films or platform simulators, we have learned to keep the story line simple and clear. We must do the same for VR. Our initial experience indicates that VR is good at placing guests in an environment, and we look forward to seeing its storytelling capacities evolve.

All our experiences to date have been with a novice audience. Filmmakers once used devices such as tearing away calendar pages to show flashbacks or passage of time. As the audience became more experienced, these devices became unnecessary.

Controlling the Narrative

We are fairly successful at composing a scene that draws the guest's attention to a desired spot. We have also experimented with using characters to direct attention. In some scenes characters point where we want the guest to look, and in others, we have a character move to be in line with another object we want in view. All these techniques can be quickly tried; the key is to test them on novices.

We have experimented with explicit techniques for controlling the guest's position, such as having a character grab the carpet and drag the guest to a desired location. Another coarse grain technique is to close doors behind guests to keep them from back-tracking. We have also experimented with implicit techniques such as a "water skiing tow rope" metaphor [13], where an invisible boat is controlling the eventual position and the guest is free to fly within a moving envelope.

Sound

In films, the sound track, particularly the musical score, tends to carry the emotional tone for most scenes. Because we no longer control timing we must choose sound tracks that work with wide variation in duration, and we must be able to make the transition smoothly from one ambient sound to the next based on guest actions.

Many VR system architects are concerned with the underlying technology for localizing sound. In our experience, the careful selection/creation of ambient sounds and music, i.e. the content, is much more important than the specific details, or even the use of, sound localization.

AUTHORING

In the process of building our first show we have learned a number of lessons regarding the process of authoring in VR.

Rapid prototyping is essential for authoring. Flight simulator technologies often guarantee rendering frame rates, but require long (many hour) periods to change the show's content. Our SAL/Scheme layer allows code interpretation at run time, similar to MR/OML [22], Alice [16], and World Toolkit [21].

We could not have developed the show without this interpreted layer.

Character animation in our second system approaches the look of traditional animation. This is not surprising, since the principles of animation apply regardless of the medium [15]. The key to achieving this was involving the artists in the development of the underlying technology, rather treating it as a given. We can now generate a new scene or character animation in under a week.

The fidelity trap for VR is that unlike many other media, a low-quality “quick and dirty” mockup is often misleading. Since a partial or low-quality mockup may not yield accurate results when we test guests on it, we often must build systems to completion before we know what works well. This is partially because there are not yet good tools for sketching three dimensional scenes and animations.

Motion capture vs. key frame animation: Our first system used motion capture to animate characters. This allowed us to produce a large amount of animation quickly, but the quality was not as high as the key frame-based animation we are using for the second version. Motion capture is troublesome for non-human characters, often seems too realistic, and requires laborious post-processing of the data.

Branching story: In a linear narrative, a character’s behavior is completely pre-planned. When the guest’s actions can cause a character to perform different pre-animated sequences, we refer to that as a *branch*. In the original show most characters performed a repose animation until the guest approached, and then branched to perform a reaction animation. While this makes an interesting and active scene, in most cases it does not provide enough different branches to allow the guest to easily suspend disbelief.

Autonomous characters vs. authored branching: Artificially intelligent characters are an interesting concept, but it will be a long time before they are believable in any but the simplest background role. For the next few years, we feel that believable character performances will be made up of branches of pre-animated material rather than computer programs for several reasons: 1) “thinking” characters are far enough into the future to be off the planning horizon, 2) characters who can construct decent sounding sentences are not much closer, 3) a good animator can achieve a much more believable dramatic performance than a computer program, and 4) even simple branching works when properly done.

In our second version characters have multiple possible behaviors that are triggered by context. Even our simplest characters have a default behavior, a reaction to the guest’s presence, and a reaction to the guest’s departure. A major technical challenge is to smoothly blend between various pre-defined animations [17].

Rotation of characters to face guests is important to present the illusion that characters are real. We find that first turning a character’s head, then his or her body, works best. The technical challenge is to avoid bad interactions between the automatic rotation and the character’s key frame animation.

RESEARCH CHALLENGES

Based on our experiences, we present the following as open challenges to the research community:

- 1) Finding mechanisms that allow guests to self-calibrate the intensity of the experience. Currently we must keep the experience tame enough to be enjoyable for our more sensitive guests.
- 2) Developing constraints to solve the six-degree-of-freedom problems in controlling flight; i.e. navigation and motion through virtual spaces.
- 3) Development of software to better support animators, especially in the sketching phase. Animators use onion skin paper to superimpose views from multiple frames; this ability is lacking in most software tools.
- 4) The automatic generation of mouth animation from sound source. This is currently labor intensive and not particularly creative work.

CONCLUSIONS

This project gave over 45,000 people a first exposure to virtual reality (VR). While we have made what we consider to be substantial advances in HMD and rendering technology, our major advances have been in learning how to create and present compelling virtual environments. We stress that this is an exercise that requires both artistic and engineering talent and creativity.

Our guests completed written surveys, and with subsets we logged head and carpet motions. Based on that data and interviews before, during and after guest flights, we conclude that:

Guests suspend disbelief. The illusion is compelling enough that most guests accept being in a synthetically generated environment and focus on the environment, not the technology.

VR appeals to everyone. Both genders and all ages had similar responses to our attraction. This leads us to conclude that VR is like feature films in that different content may segment the market, but the basic technology does not. We also note that wheelchair guests find mobility within VR extremely exciting.

VR must be personally experienced. VR is foreign enough that most people can not comprehend it without direct personal experience.

Fidelity matters. To get most guests to suspend disbelief requires extremely high fidelity. We provide 60 frames per second (at the expense of stereo), for polygonal models with hand-painted texture maps, and we do not use hardware lighting. Texture quality matters much more than polygon count.

Content matters. People love the experience of VR, but even at high fidelity VR by itself is not enough. The public, unlike

the developers, is not impressed with the technology. In fact, the public *assumes* that high fidelity VR exists and immediately focuses on what there is to do in the synthetic environment.

The illusion of presence is fragile. Although guests suspend disbelief, inconsistencies can *instantly* shatter the illusion. For example, objects inter-penetrating, or characters not responding to the guest's presence completely shatter the sense of presence.

Guests need a background story. VR is an overwhelming experience of being thrust into a new environment. A good way to soften this transition is to provide a background story that familiarizes the guest with the new environment before the immersion. This is a standard technique in theme park attractions, typically provided in a pre-show.

Guests need a goal. Guests need to know why they are in the virtual world and what they are supposed to do.

Guests do not turn their heads much. We were surprised at how little people turned their heads in this flight-based experience. We attribute this to the mass of the HMD, the need to look where one is flying, and guests' inexperience with a head-tracked medium.

Input controls are hard. We developed a novel input mechanism for controlling flight. Since no one flies magic carpets in the real world we could not transfer everyday skills. After many design iterations we believe that six axis control is a phenomenally difficult problem and conclude that designers must limit degrees of freedom.

Tell a straightforward story. As we have learned with other intensive media, such as effects laden stereoscopic films and motion-base simulators, when the guest is perceptually overwhelmed it helps to keep the story short and clear.

Aladdin is a beginning, not an end. Our original goal was to move past the technology. Our first system produces a compelling illusion and our next efforts are to examine whether we can tell stories in this new medium. Our second version of the project, scheduled for release in early 1996, contains a large number of characters and a narrative story line.

ACKNOWLEDGMENTS

This work is the effort of many talented people over several years; we mention here only a subset, but express our gratitude to all involved.

Special thanks, in alphabetical order, to: Daniele Colajacomo, for managing the character modeling in the EPCOT show; Dave Fink, for helping start the project; Phillip Freer and Gary Daines, for their art direction and world design; Andy Ogden, for his industrial design on the steering and HMD; George Scribner, for his work on story and character in the EPCOT show; and Dave Spencer, for his management of the EPCOT show installation.

We thank all the other artists and engineers who worked on this project, and we would especially like to express our deepest gratitude to the families and significant others who supported these individuals in their efforts.

We would also like to thank Evans & Sutherland, Silicon Graphics, NASA Ames Research Center, the staff who ran the attraction at EPCOT Center, and Matt Conway.

REFERENCES

- [1] Jon Airey, John Rohlf, Frederick Brooks, Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments, ACM SIGGRAPH Special Issue on 1990 Symposium on Interactive 3D Graphics 24:2, 1990, pages 41-50.
- [2] Walt Disney Home Video. Distributed by Buena Vista Home Video, Dept. CS, Burbank, CA, 91521. ISBN 1-55890-663-0. Originally released in 1992 as a motion picture.
- [3] Chuck Blanchard, Scott Burgess, Young Harvill, Jaron Lanier, Ann Lasko, Reality Built for Two: A Virtual Reality Tool, ACM SIGGRAPH 1990 Symposium on Interactive 3D Graphics, March 1990.
- [4] J. P. Blauert, Spatial Hearing, MIT Press, Cambridge, MA, 1983.
- [5] Frederick Brooks, Walkthrough -- A Dynamic Graphics System for Simulating Virtual Buildings, Proceedings of the 1986 Workshop on 3D Graphics, Chapel Hill, NC, October 23-24, 1986, ACM, pages 9-21.
- [6] Steve Bryson, Creon Levit, The Virtual Wind Tunnel, IEEE Computer Graphics and Applications, July 1992, pages 25-34.
- [7] Christopher Codella et al., Interactive Simulation in a Multi-Person Virtual World, Proceedings of the ACM SIGCHI Human Factors in Computer Systems Conference, May 1992, pages 329-334.
- [8] James Cutting and Peter Vishton, Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Contextual Use of Different Information About Depth, Handbook of Perception and Cognition: Perception of Space and Motion, Vol. 5, Academic Press (to appear).
- [9] R. Kent Dybvig. The Scheme Programming Language. Prentice-Hall, 1987.
- [10] S. S. Fisher, A. M. McGreevy, J. Humphries, W. Robinett, Virtual Environment Display System, Proceedings on the 1986 Workshop on Interactive 3D Graphics, pages 77-87, October 23-24, 1986.
- [11] Thomas Funkhouser, Carlo Sequin, Seth Teller, Management of Large Amounts of Data in Interactive Building Walkthroughs, Proceedings of the 1992 ACM

Symposium on Interactive Three-Dimensional Graphics, April, 1992, pages 11-20.

7801-7805 Mesquite End Drive, Suite 105, Irving, Texas 75063, USA 001-214-556-1800, 1-800-ILLUSION, enquiries@tx.virtuality.com.

- [12] Thomas Furness, The Super Cockpit and Human Factors Challenges, Human Interface Technology (HIT) Laboratory of the Washington Technology Center, Tech Report HITL-M-886-1, October, 1986.
- [13] Tinsley A. Galyean, Guided Navigation of Virtual Environments, 1995 ACM Symposium on Interactive 3D Graphics, April 1995, pages 103-104, Monterey, CA.
- [14] Id Software, Inc.; information available via <http://www.idsoftware.com/>.
- [15] John Lasseter, Principles of Traditional Animation Applied to 3D Computer Animation, Computer Graphics (SIGGRAPH '87 Proceedings) Volume 21, number 4, pages 35-44, July 1987.
- [16] Randy Pausch, et al, A Brief Architectural Overview of Alice, a Rapid Prototyping System for Virtual Reality, IEEE Computer Graphics and Applications, May 1995.
- [17] Ken Perlin, Real Time Responsive Animation with Personality, IEEE Transactions on Visualization and Computer Graphics, Vol. 1, No. 1.
- [18] A. Pope, BBN Report No 7102. The SIMNET Network and Protocols. BBN Systems and Technologies, Cambridge, Massachusetts, 1989.
- [19] John Rohlf and James Helman, IRIS Performer: A High Performance Multiprocessing Toolkit for Real-Time 3D Graphics, SIGGRAPH '94 Conference Proceedings, Computer Graphics, July, 1994.
- [20] William J. Schroeder, Jonathan A. Zarge, William E. Lorensen, Decimation of Triangle Meshes, Computer Graphics (SIGGRAPH '92 Proceedings) Volume 26, pages 65-70, July 1992.
- [21] Sense8 Corporation, 100 Shoreline Highway, Suite 282, Mill Valley, CA 94941, 415/331-6318, <http://www.sense8.com/>, info@sense8.com.
- [22] Chris Shaw, Mark Green, Jiandong Liang, Yunqi Sun, Decoupled Simulation in Virtual Reality with the MR Toolkit, ACM Transactions on Information Systems, 11:3, pages 287-317, July 1993.
- [23] Ivan Sutherland, The Ultimate Display, Proceedings of IFIP (International Federation of Information Processing) '65, Vol. 2, pages 506-508.
- [24] Ivan Sutherland, A Head-mounted Three-dimensional Display, Proceedings of the Fall Joint Computer Conference, AFIPS, Vol. 33, pages 757-764.
- [25] Virtuality, <http://www.virtuality.com>, UK office: Virtuality House, 3 Oswin Road, Brailsford Industrial Park, Leicester LE3 1HR, United Kingdom, Tel: +44(0) 116 233 7000, enquiries@virtuality.com. USA office: [26] E. Wenzel, F. Wightman, S. Fisher, A Virtual Display System for Conveying Three-dimensional Acoustic Information, Proceedings of the Human Factors Society, 32nd Annual Meeting, 1988, pages 86-90.



Figure 1: A Guest's View of the Virtual Environment



Figure 2: The Head Mounted Display

High-resolution TIFF versions of these images can be found on the CD-ROM in
S96PR/papers/pausch



Figure 3: The Physical Setup



Figure 5: Animated Characters

High-resolution TIFF versions of these images can be found on the CD-ROM in
[S96PR/papers/pausch](#)