


~ MAKING VIRTUAL ENVIRONMENTS COMPELLING

DELIVERING A COMPELLING USER EXPERIENCE AND ENSURING
APPLICATION SUCCESS BOTH DEPEND ON THE FIDELITY OF
THE USER'S SENSORY IMMERSION.

[By Mary C. Whitton]




THE VERY IDEA OF A VIRTUAL ENVIRONMENT (VE) is compelling—being able to go places and do and experience things you couldn't or wouldn't in the real world. The wow factor certainly makes people's initial experience with the technology exciting, along with the fact that most people who have experienced VEs have done so in entertainment venues carefully designed to be engaging and fun. It is the user's immersion in the sights and sounds of the virtual world that sets VE applications apart from their conventional counterparts.

The reality of today's VE systems isn't what Hollywood films like *The Matrix* and *Disclosure* have depicted over the past several years. We don't yet have the technologies to build a *Star Trek*-like Holodeck, where virtual space is unlimited and objects have all the affordances they have in the real world, and where you can feel and manipulate things and sit on virtual furniture. However, despite VE-system limitations, compelling and successful

VE applications do exist. Some are gut-wrenchingly compelling because of their realism; Figure 1 (right) shows a VE that makes user heart rates increase by about eight beats/minute. Other VEs, while not compelling in the sense of personally gripping, are impressive simply because the application wouldn't exist without VE. Figure 2 shows an engineer evaluating a manufacturing process at full scale before it is built.

A compelling VE application depends on a VE system being able to provide high-quality sensory immersion, a well-designed application, and a motivated user. A minimal VE system today includes a mathematical model of the VE (the virtual world), a head-mounted display presenting images of the virtual world to the user, and a tracker on the user's head reporting which way the user is looking. The user moves through and interacts with the environment through a handheld controller. Using the model and data from the tracker and controller, the computer draws, as quickly as possible, the



REALISTIC CLOUD SIMULATION AND RENDERING PRODUCE MORE REALISTIC REAL-TIME FLIGHT SIMULATIONS. THE IMAGE WAS CAPTURED FROM SKYWORKS, A 3D ENGINE WRITTEN TO DEMONSTRATE REAL-TIME CLOUD RENDERING RESEARCH; SEE www.cs.unc.edu/~harrism/SkyWorks. (ENGINE DEVELOPED AND IMAGE CREATED BY MARK J. HARRIS OF THE UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL, 2002.)



Figure 1. (left) A user stands on a wooden ledge (passive haptics) corresponding to the ledge in the VE; wires go to the video source, tracker, wand, and physiological measuring devices. (right) The user's view into the pit; the avatar of the user's hand is visible as he prepares to drop the red-and-white ball on the target below (University of North Carolina at Chapel Hill, inspired by Mel Slater, University College London).

head motion causes appropriate changes in the visuals, and the visuals appear life-size. This definition establishes a minimum level of user immersion and includes projection-based systems, like the one in Figure 2.

Sensory Immersion

While visual immersion remains the defining quality of VE systems, modern computing and VE technologies can immerse users not only in low-latency, high-quality visual stimuli but also in full spatial audio. Some VE systems include motion platforms, scent-disbursal systems, and active and passive haptic devices that allow users to feel objects in the VE. Some systems enhance user input with gesture recognition and voice input. As a general rule (application factors being equal), the more a VE system immerses its users in stimuli corresponding to their expectations of the virtual

world as seen from the user's point-of-view and sends the images to the head-mounted display. In the VE laboratory at the University of North Carolina at Chapel Hill, my colleagues and I define a VE system as having two characteristics: the user's

world, the more compelling their VE experience.

We find it useful to have language that distinguishes the technologies in the VE system and the sensory stimuli they deliver to the user from the effect of those stimuli on the user; for example, we follow [9, 10], reserving the word immersion to mean what the VE system technology delivers to the user, or the stimuli that collectively represent the virtual world. Presence is further defined in [10] as the effect of immersion on users, that is, their mental state (see the sidebar "Presence"). This usage, however, is controversial; [12] offers an alternative perspective.

Degree of immersion depends on how many senses, including vision, hearing, and touch, are simulated and stimulated by the VE system and how well users are isolated from the real-world environment. Quality of immersion varies with the fidelity of physical simulations, rendering (for all senses), and presentation/display of the data. Factors contributing to immersion are often measured and compared objectively. Examples of such factors include: geometric resolution of models; time resolution of a particle-system simulation of falling water; vehicle simulation physics; how the graphics simulate the physics of light transport; detection of and realistic response to collisions between objects; display field-of-view, resolution, brightness, and refresh rate; frequency response of speakers or headphones; processor speed; and latency of response to user input.



Consistency across senses. Sensory stimuli must be consistent and synchronized for users to perceive the world they represent as coherent and predictable. For instance, for visual consistency, if window curtains are moving in a breeze, the window should be open; for visual and haptic consistency, users should feel a breeze (a fan) as they approach the open window. The sounds coming through a city window should be street noises, not lowing cattle. Sensory conflicts resulting from fundamental limitations of the immersion system are more important and less easily addressed. For example, while passive haptics are sometimes used in VEs (such as the real wooden ledge in Figure 1 left), there is no general solution to the problem of including solid objects in VEs. In VEs, users may see an object that should be solid but be unable to feel it. Similarly, unless the system includes a motion platform, when

users push a button to “run” through the environment, only visual cues tell them they are moving. Because they are actually standing still, the vestibular system in the inner ear detects no acceleration, resulting in a conflict between visual and vestibular cues.

Even if a VE system is capable of providing realistic and compelling immersion, it is just technology until it is used in applications. Like traditional developers, VE developers must first understand the application’s goals, then identify target users, decide what they need to do to accomplish those goals, and decide what user interface tools will be available to help them. Unlike most application developers, VE developers must define where and under what conditions users might perform a task.

Users bring skills, experiences, and motivation to VE applications, and a good application designer recognizes and utilizes these personal characteristics. If an application is too difficult, the user is likely to be frustrated; if it is too easy, the user is likely to be bored. Even within the target user population, individual differences can affect the quality of an individual user’s VE experience. Some people find it easy to suspend their disbelief and “go” to the virtual place; others remain highly aware they are in the laboratory. Some users are susceptible to cybersickness and can spend only minutes in a VE; others can work immersed for extended periods. User interface devices for VEs are too often nonintuitive and unnatural, and

are arguably the least satisfactory component of VE systems today. The affordances of the tracked gloves, as in Figure 2, and tracked wands with buttons, as in Figure 1, aren’t a good match for many tasks. Imagine trying to open a virtual jar of peanut butter or tie a virtual suture with a wand and push button.

The user performs the application task while immersed in the virtual world. One of the reasons VE applications are so costly is the developer must define everything about that world, including the objects and entities in it, the behaviors of the objects and entities (whether autonomous or in response to user input), and even the physics of the world. Any feature or behavior that isn’t explicitly defined can’t and shouldn’t be there. If the developer doesn’t design a feature or behavior, it won’t be there. For example, if the application requires objects to fall to the floor when they are dropped and accelerate when falling,

SOME PEOPLE find it easy to suspend their disbelief and ‘go’ to the virtual place; others remain highly aware they are in the laboratory.

the VE design must include a gravity model. The amount of detail in the models making up the VE—visual, aural, or haptic—must be determined by the application requirements. Good design principles apply in virtual worlds just as they do in the real one; “just because you can” isn’t a good enough reason to add embellishments (such as detailed wallpaper patterns) that may, in fact, distract the user.

The application’s goals are the major determinant of the design of the virtual world. Consider how different two applications based on the same medical procedure can be. The goal of application A is teaching a medical procedure under clinical conditions; the goal of application B is teaching the user to perform the procedure in the midst of chaos in a disaster-relief unit. Meeting the goals of these applications requires enormously different virtual worlds. When the goal is training, developers exploit the fact that the virtual world and the conditions in it can be changed rapidly. Flight simulators, the original and best-known train-




Figure 2. A user immersed in a virtual factory observes a discrete event simulation of a manufacturing process (Carolina Cruz-Neira, Virtual Reality Applications Center, Iowa State University, Ames, IA).

changes of condition. Developing effective, emotion-evoking scenarios and environments is part of the art of designing training applications (see Swartout's and van Lent's article in this section).

Pitfalls and Promise

The ultimate test for VEs is their effectiveness in supporting application goals. Reports of successful VE applications appearing outside the computer graphics literature are one measure of that success (see the sidebar "Immersion Requirements"). However, despite demonstrated application success, the use of VEs is not likely to expand quickly for technological, market, and social reasons. The cost of VE systems and developing VE applications is relatively high. Although the computer game industry has driven down the cost of computer graphics hardware, other fundamental VE technologies, including head-trackers and displays (both head-mounted and multi-projector stereo) remain costly. There is as yet no high-volume market or its incentive for cost reduction.

Public reaction to real and projected dangers of VEs may also slow VE use in applications, even where

ing application of VE, have convincingly demonstrated the value of being able to safely practice a range of probable and improbable scenarios and to repeat a scenario many times with random

it offers significant advantage. The possibility of long-term personality effects from participation in violent VE-based entertainment is a hot issue, especially when the individuals exposed to the violence aren't simultaneously being trained in morals and ethics. Cybersickness, or the adverse physical effects of VE use, is a concern for all responsible researchers and application developers. Symptoms include unsteady-

User interface devices for VEs are too often nonintuitive and unnatural, and are arguably the least satisfactory component of VE systems today.

ness, mild nausea, and eye fatigue. Though infrequent, more subtle effects (such as disorientation and flashbacks) can occur, with potentially serious consequences. The frequencies of various adverse effects are outlined in [11], which also describes protocols and system performance characteristics that minimize the risk of cybersickness.

Realizing the promise of VEs won't, in the short



IMMERSION REQUIREMENTS

VES CAN be effective without providing perfect immersion, though several VE system elements are critical for success in a variety of applications. For example, when working with realistic virtual prototypes in VEs in design applications, individual designers and cross-specialty product teams are able to evaluate styling, usability, manufacturability, and maintainability early in the design process. Different design activities make different demands on the VE system; they don't all require the highest-quality immersion on the same things. Model accuracy, lighting, and rendering matter most for styling; a full-scale display matters most for evaluating human-scale spaces. Upon completion of the manufacturing line layout being simulated in Figure 2, and before any equipment is moved or purchased, plant managers and operators are able to evaluate the line's efficiency (at full scale) by following individual pieces as they move through the manufacturing cycle.


Phobia treatment is another area of VE success. Mel Slater and his colleagues in the computer science and psychology departments at University College London devised and evaluated an application to help users reduce their fear of public speaking [6]. Instead of speaking to an audience of real people, study participants ($n=40$) presented to an audience of computer-generated people, or avatars. The avatars were programmed to exhibit a range of responses to the speaker, varying from high levels of interest to outright hostility (see figure above).

The application-defined behavior of the avatars is critical to evoking an emotional response in the users—and to their learning to overcome these emotions.

Working with Barbara O. Rothbaum of the Department of Psychiatry at Emory University, Larry Hodges of the University of North Carolina at Charlotte, developed phobia-treatment VR applications that are in routine use worldwide. Hodges and his colleagues found that when treating Vietnam veterans with post-traumatic stress disorder (see figure right), the quality of the sound model was more important than the quality of the visuals in helping participants re-experience situations and learn to manage the emotions they evoked. **C**

(TOP) BEHAVIOR OF THE AVATAR AUDIENCE IN THE FEAR OF PUBLIC SPEAKING SYSTEM CAN BE PROGRAMMED TO EXHIBIT VARIOUS LEVELS OF INTEREST IN THE SPEAKER, INCLUDING (LEFT) MILD DISINTEREST AND (RIGHT) HOSTILITY (DAVID-PAUL PERTAUB AND MEL SLATER, UNIVERSITY COLLEGE LONDON).

A USER (INSET) AND THE VIRTUAL VIETNAM SYSTEM (VIRTUALLY BETTER, INC., DECATUR, GA).



PRESENCE

IN ADDITION to quality of the immersion and application success, VEs are often evaluated on how well they induce a mental state in their users making them feel, act, and react as they would in a corresponding real-world setting. Using the definitions in [10], that mental state, or “presence,” is the user’s personal response to immersion. Presence is a difficult (and controversial) thing to measure directly.

Researchers collect indirect evidence of presence with behavioral and physiological measures. Behavioral observations are used to determine how “present” a user feels in an environment. Physiological responses appropriate to the VE, particularly stress and relaxation responses, are another indirect indicator of presence. Differences in physiological measures (such as heart rate and skin conductance) taken in two or more VEs can be used as indicators of relative levels of presence [4].

For good overviews of research on presence in VEs, including how different qualities of immersion affect presence and how presence and performance are related, see [1, 2, 5, 8]. Research by the Effective Virtual Environments Group in the Graphics and Image Laboratory at the University of North Carolina at Chapel Hill has shown the following factors increase presence: really walking in the environment instead of push-button flying; increased field-of-view in head-mounted displays; real things in the environment to touch and feel (passive haptics); high graphics frame update rate; and low end-to-end system latency (see www.cs.unc.edu/Research/eve). **C**

term, mean installations in grade-school classrooms or in homes. The VE promise is that a combination of immersing technologies and well-designed applications will let users experience real, recreated, abstract, or imaginary places that are too big, too small, too far, too costly, or too dangerous to visit in person and let users do things they can’t or wouldn’t do in the real world; for example, they might let medical personnel train, but not on human patients, and let emergency personnel train in dangerous situations, but out of harm’s way. The VE promise is also in as yet unthought-of applications in medicine, design, training, education, data visualization, entertainment, and the fine arts. Today, even without systems as intriguing as a Holodeck, VEs are proving their value through effective and compelling applications. The future promises much more. **C**

REFERENCES

1. Draper, J., Kaber, D., and Usher, J. Telepresence. *Human Factors* 40, 3 (fall 1998), 354–375.
2. IJsselstein, W., Freeman, J., and de Ridder, H. Presence: Where are we? (editorial). *Cyberpsych. Behav.* 4, 2 (Apr. 2001), 179–182.
3. Kelsick, J. and Vance, J. The VR Factory: Discrete event simulation implemented in a virtual environment. In *Proceedings of the 1998 ASME Design for Manufacture Conference* (Atlanta, Sept. 13–16). American Society of Mechanical Engineers, New York, 1998.
4. Meehan, M., Insko, B., Whitton, M., and Brooks, Jr., F. Physiological measures of presence in stressful virtual environments. *ACM Transact.*

Graph. 21, 3 (July 2002), 645–652.

5. Nash, E., Edwards, G., Thompson, J., and Barfield, W. A Review of presence and performance in virtual environments. *Int. J. Hum.-Comput. Interact.* 12, 1 (2000), 1–41.
6. Pertaub, D.-P., Slater, M., and Barker, C. An experiment on fear of public speaking in virtual reality. In *Medicine Meets Virtual Reality 2001: Outer Space, Inner Space, Virtual Space, Vol. 81 Studies in Health Technology and Informatics*, D. Stredney, J. Westwood, G. Mogel, and H. Hoffman, Eds. IOS Press, Amsterdam, The Netherlands, 2001, 372–378.
7. Rothbaum, B., Hodges, L., Ready, D., Graap, K., and Alarcon, R. Virtual reality exposure therapy for Vietnam veterans with post-traumatic stress disorder. *J. Clin. Psychiatry* 62, 8 (Aug. 2001), 617–622.
8. Schuemie, M., van der Straaten, P., Krijn, M., and van der Mast, C. Research on presence in virtual reality: A survey. *CyberPsych. Behav.* 4, 2 (Apr. 2001), 183–201.
9. Slater, M. A note on presence terminology. *Presence-Connect* 3, 1 (Jan. 2003).
10. Slater, M. Measuring presence: A response to the Witmer and Singer questionnaire. *Presence: Teleop. Virtual Environ.* 8, 5 (Oct. 1999), 560–566.
11. Stanney, K., Kennedy, R., and Kinson, K. Virtual environment usage protocols. In *Handbook of Virtual Environments: Design, Implementation, and Applications*, K. Stanney, Ed. Lawrence Erlbaum Associates, Mahwah, NJ, 2002.
12. Witmer, B. and Singer, M. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleop. Virtual Environ.* 7, 3 (June 1998), 225–240.

MARY WHITTON (whitton@cs.unc.edu) is a research assistant professor in the Department of Computer Science at the University of North Carolina at Chapel Hill.