Panels

The Year of the Human

1998 was an interesting year to serve as Panels chair for SIGGRAPH, for it is both the 25th Anniversary of the SIGGRAPH conference itself and the eve of the eve of the new millennium. Not to belabor a cliché, but we sit here teetering between the past and the future, in the precarious present of an "industry," such as it is, that, since its inception, has experienced nothing short of constant growth and erratic perturbations. Companies rise and fall, go public, go under, expand, and contract. While our tenured research colleagues become fixtures and establish institutions within their institutions, our commercial friends seem to be on an endless rollercoaster ride of enterpreneurial rags-to-riches-to-rags-again (though the sum total appears to be mostly riches; otherwise, the cycle would not continue).

This year, through the 25th Conference Celebration program, we add to the visionary mission of SIGGRAPH the gift of hindsight, which transforms through insight into foresight. We look back, we take stock of the present. Then, armed with the sense of perspective that can only be gained through hard-won experience, we look to a future in which we can only vaguely imagine the impact of what we are creating today. And this is where the Panels program comes in.

If anything is said of the SIGGRAPH 98 panels program, let it be that it served as a kind of prism, a bending point of light, a juncture at which everything changed direction entirely, if only slightly. For 25 years, we have concerned ourselves primarily with tools, technologies, and techniques. Over time, we have occasionally entered into the realm of the sublime, but mostly, we were just trying to get everything to work right. Fortunately, we've gotten better, so much so that we can now begin to take a hard look at what it all means. If the Panels program is any indication, 1998 appears to be "The Year of the Human" for computer graphics.

This year's Panels explore the art and science of image and interface, and address the technical, practical, aesthetic, and social challenges we face as we build the future into the next millennium. But more than anything else, they reflect an industry-wide trend in what could best be called "human-centered computing," as manifested through a concern for better and more "transparent" interfaces, more meaningful forms of interactivity, more immersive forms of display, more ubiquity in computing, more authenticity in synthetic characters and worlds, and a greater concern than ever for the psychological, social, and even physical impact of computer graphics on *people*.

Virtual reality is making a comeback, but it has transmuted from its two extremes of high-end rarification and pop-culture hype. We are now seeing real-world applications that bring virtual reality to realms as diverse as LEGO and Parkinson's disease. We are teaching characters how to think like us and teaching computers how to care about what we think. We are finally making computers smart enough to understand us. A certain amount of wisdom and maturity is reflected in the depth of such topics as virtual healing, out-of-the-box toys, and the psychological impact of presence in virtual worlds. Even in film effects, we are seeing trends such as behavior modeling and a greater concern for more lifelike characters. Social engineering is as important as real-time rendering. Storytelling is as important as image processing. Human processing is at the center of it all.

Celia Pearce
SIGGRAPH 98 Panels Chair

Visualization: The Hard Problems

Organizer
David Zeltzer
Sarnoff Corporation

Panelists
Ann M. Bisantz
State University of New York at Buffalo

Jock D. Mackinlay
Xerox PARC

Krzysztof Lenk
Dynamic Diagrams

Randall W. Simons
Sandia National Laboratories

It has been 10 years since publication in the SIGGRAPH Newsletter of the National Science Foundation report, Special Issue on "Visualization in Scientific Computing." In this report, the authors proposed a definition of the visualization problem and suggested important research directions. The "firehose of data," as described in the NSF report, consisted primarily of the numeric output of simulations running on supercomputers.

Today, however, information sources and categories to be visualized have expanded enormously, including such applications as visualizing Web search results, depiction of complex communication network topologies, medical imaging, and battlefield situation awareness. If a problem has a relatively straightforward mapping to 3D geometry, there are numerous visualization packages that do quite well at interactively portraying the problem and the solution space. While many problems require display of more than three dimensions, a few extra dimensions can readily be mapped to effects such as color, texture, or animation.

But many visualization designers are now confronted with visualization problems that require access to diverse and massive sources of information, often located in distributed databases. An aviation-related application, for example, may require distributed and varied information such as maps, position and location of many aircraft, predictive weather simulations, satellite communication envelopes, aircraft status and capabilities, flight crew status and availability, and various alerts and warnings. The world of finance, moreover, offers examples of information spaces that are, for the most part, abstract and highly-multidimensional, with no obvious mapping to three-space.

So the hard problems remain:

- How can many, varied kinds of information be accessed, retrieved, and coherently displayed and manipulated?
- How can information qualities such as timeliness, accuracy, and uncertainty be portrayed?
- What does it mean to "understand" data in the first place?
- How can the "information environment" of a visualization problem, and the concomitant "information operations," be defined and described?
- How can knowledge of human perception and cognition be incorporated in design of visualization tools and techniques?
- How can human perceptual and cognitive talents be enhanced and amplified through visualization?
- How can the long and rich history of visualization in the arts be exploited in the information age?

These and other visualization questions are addressed by a multi-disciplinary panel from a variety of pertinent disciplines: computer graphics, human factors, cognitive science, and the graphic arts. Panelists also present visualization solutions designed to solve these kinds of visualization problems.

David Zeltzer
Sarnoff Corporation

Maintaining a coherent tactical understanding in modern warfare is extraordinarily difficult due to the proliferation of high-performance weaponry, increases in the numbers and types of sensors, the mixing of combatants and non-combatants, and enhanced communications technologies. JOVE (the Joint Operations Visualization Environment) provides a significant enhancement to achieving and maintaining situational awareness. JOVE presents a common operational picture of a joint operation to the Joint Task Force Commander and staff officers, and enables these users to:

- Maintain an accurate and coherent understanding of the battlespace.
- Assimilate information from different echelons, modes, and data sources.
- Decrease reaction time by direct, effective, and timely presentation of data.
- Reduce error through direct user-system interaction.

JOVE presents, on an immersive display system, a computer-generated, three-dimensional visualization of the air, surface, land, and undersea battlespace. Tracks of land, air, surface, and undersea entities are represented symbolically, and the 3D stereo presentation of track history enables the user to judge intent. JOVE has been deployed at seven military exercises on three continents in the past two years, and is now operational in the combined U.S./Korean command center in Seoul, South Korea; the Joint Training, Analysis and Simulation Center (JTASC) at Norfolk, Virginia; and at a facility maintained by the U.S. Air Force Communications Agency at Scott Air Force Base.

The current JOVE system depends strongly on mapping of information to readily understood 4D presentations (3D + time). However, continued development of JOVE will require finding effective visualization tools and techniques for understanding and interacting with mission-spaces of high dimensionality, abstract quantities or relationships, and diverse kinds of information.

Ann M. Bisantz
State University of New York at Buffalo

From a cognitive engineering standpoint, the computer interface (including the data display) is a window between users and the goals they are trying to accomplish. The interface is not an end in itself, but in some sense should be as transparent as possible, allowing people to perform their tasks in a direct way. Any technique for displaying complex data relationships should be based on users' task-relevant goals, which generally are dependent on the task, time, and situation.

For example, visualization of weather systems may be very useful to meteorologists who must make sense of complex weather patterns and make predictions. But for airline pilots and air traffic controllers who are trying to find safe pathways for aircraft, the meteorologist's information display may be less useful, and potentially even hazardous. That is, successful visualization answers are not generic solutions, but instead are tailored to particular tasks and contexts.

Given this perspective, one can address the question of how to visualize complex information by noting first that the complexity lies not only in the structure of the data itself, but also in the complexity of the often-dynamic goals and activities of the users. However, analyses of these same changing goals and tasks to identify the kinds of actions the data need to support may provide some constraint on, and thus insight into, the best methods for displaying even multi-dimensional and abstract data.

Krzysztof Lenk
Dynamic Diagrams

Modern visualization of complicated data originated in the Renaissance. The idea of studying and describing human beings and their environment as a way to understand the intentions of the Lord is visible in analytical visual studies of Luca Paccioli, Leonardo da Vinci, and Albrecht Dürer. But the real progress in the development of visualization methods is related to the scientific revolution of the late 16th and 17th centuries that also witnessed development of book production on a large scale.

Sophisticated methods of presenting complicated, abstract data in the form of multilayered and multi-windows diagrams were developed, and works of Simon Stevin and Johann Comenius still surprise us by their genuine inventiveness. The best examples of their works show an understanding of the process of communication, where a part of an incoming message has to activate contexts already existing in the reader's mind, in order to decode the meaning.

This presentation shows some of the most interesting historical diagrams, as well as contemporary examples of diagrams visualizing the computer space from Dynamic Diagrams' portfolio.

Information Visualization: Using Vision to Think

Jock D. Mackinlay
Xerox PARC

The goal of information visualization is to use interactive visual representations of abstract data to amplify cognitive activities such as sense-making, decision-making, and large-scale monitoring. This amplification of cognition can arise in at least two ways. The first way is through transforming information into visual forms such that special powers of human perceptual operations can be brought to bear. For example, a display might suggest to a trader within seconds, which stock out of 2,000 actively trading issues should be immediately investigated. The second way is through indirect effects of perception, such as the ability to keep track of more items of work with an advanced workspace. My focus is on "using vision to think" rather than using graphics to present information to another person. I describe the following reference model for information visualization, which is based on developing mappings of data to visual forms:

(Raw Data -> Data Tables -> Visual Structures -> View)

Raw data are transformed into Data Tables that can be mapped fairly directly to Visual Structures, which are then combined with a View transformation such as zooming, lensing, and distortion. In the other direction, human interaction involves control of these mappings to create an environment for working with the information. Data come in a number of types and so do their visual encodings, from which arise a number of constraints on the mappings. Finally, I use this reference model as a framework to describe recent research on information visualization.

Randall W. Simons
Sandia National Laboratories

The Comprehensive Test Ban Treaty (CTBT) Research & Development Data Visualization Project prototyped and evaluated new approaches to presentation of data for CTBT monitoring applications. The great amount of data expected to be available, and the complex interrelationships in that data, made this a promising area for scientific data visualization techniques.

Project members gained experience with various data visualization and user-interface design tools, and prototyped some new tools. We found that while a good tool set is useful, there is no substitute for understanding the data, the science behind it, and the customers who will use it. That understanding is required to find appropriate metaphors to represent spatial and non-spatial data, and give users the interactivity to explore and focus on features of interest.

Look Ma! Four Hands! New Models for Interacting with 3D Environments

Organizer/Moderator
Julian E Gómez
LEGO A/S, SPU-Darwin

Panelists
Dan Mapes
LEGO A/S, SPU-Darwin

Andries van Dam
Brown University

Henry Sowizral
Sun Microsystems

Dan Venolia
Cosmo Software

Julian E Gómez
LEGO A/S, SPU-Darwin
jeg@digi.lego.com
www.LEGO.com

What is so hard about 3D interaction? What exactly is being done to improve the bandwidth of the man-machine interface in 3D environments?

The most important shift over the last few years is that the World Wide Web has changed the basis on which ideas can be disseminated and communicated. The presence of the Web means that code now can be practically developed that will run everywhere. In terms of 3D interaction, VRML 2.0 provides mechanisms for rudimentary interaction, and Java 3D increases the common base of what's possible. Thus, instead of interface paradigm development being localized to particular laboratories, technologies can easily be distributed unilaterally, and in a networked fashion.

However, there is a very serious question of how performance affects interaction. In real life, 3D manipulation is immediate and, in fact, generally involves a real-time feedback loop ("real-time" is used here in its technical rather than its marketing meaning). Most computers can't yet provide this kind of throughput, leading to the issue of how and if interaction techniques should be modified in the presence of slower update rates. This is especially a problem over the Web, because there is no guarantee as to the performance level of the target platform.

Beyond performance, there's the question of the complexity or non-complexity of the input and output device(s). Does adding one level of complexity (e.g., a three-axis mouse instead of the normal two-axis mouse) significantly increase the ability to work in a 3D environment? Does adding six more degrees of freedom do the trick? How exactly should degrees of freedom be mapped to input or output? A simplistic mapping of one scalar to another (e.g., mouse X controls X translation, mouse Y controls Y translation, etc.) is better than nothing, but would functional networks provide better relationships?

Finally, does interaction in a 3D environment even need to be fundamentally different from interaction in a 2D environment? Are there any paradigms from the well-understood 2D environment that transfer to the 3D environment, especially in light of today's performance and input-output hardware?

The panelists represent a cross-section of ideas on HCI in 3D environments. We discuss interface mechanisms, outgrowths of ideas originally based in VR research, and Java 3D.

Dan Mapes
LEGO A/S, SPU-Darwin
dmapes@digi.lego.com
www.LEGO.com

If "play is the work of children," then as HCI designers we should be asking ourselves how children's work might be better enabled through our technology. Technology in itself is almost certainly not the key to play. Naïvely pasting a trendy high-bandwidth interface onto a game concept has proven to be a sure recipe for disaster. When designing interfaces using different types of simultaneously controllable input freedoms and output displays, we need to better understand the tradeoffs we are making among performance, precision, creativity, and intuitivity. We also need to maintain a clear focus on the actual problem we are trying to solve. Without this understanding, our interfaces become more complicated and the work of children ultimately confounded.

It's not valid to assume that by buying into higher-bandwidth interfaces you ensure success.

Henry Sowizral
Sun Microsystems
henry.sowizral@eng.sun.com
www.sun.com

It doesn't make tremendous sense for researchers to spend valuable time finding ways to extend existing 2D input technology so people can manipulate 3D content. Yes, we need to provide ways for manipulating 3D objects and for navigating within 3D environments using existing input devices, but not at the cost of finding the new paradigms equivalent to those found with the widespread use of the mouse. Much as we can use a keyboard to mimic some of the functionality of a true 2D input device (mouse), I firmly believe we will be able to use a keyboard and mouse to mimic some of the functionality of a true 3D input device. However, until we develop effective 3D – true six-degree-of-freedom (6DOF) – input devices and the corresponding manipulation and navigation metaphors, we don't know what features we need to mimic.

The general marketplace is feeling the effects of high-volume availability of 3D technology. Not only are million triangle-per-second output devices flooding the marketplace, but also higher degree-of-freedom input devices such as joysticks, low cost spaceballs, and lower-cost head trackers have become widely available. New higher-dimensional output devices such as force-feedback joysticks and haptic displays are also entering the marketplace. This has generated a wealth of opportunity for human-computer input-and-output interaction research.

Because we do not know what devices will work best, we need an environment for enabling the research that will result in the breakthrough technologies that truly enable tomorrow's 3D interaction metaphors. The Java 3D API includes two features that make it ideal for those who want to use or experiment with 3D or higher-degree-of-freedom I/O devices, specifically a generalized view model and a new flexible input model. We have already implemented a variety of 6DOF input devices and anticipate demonstrating Java 3D in action using these facilities.

Andries van Dam
Brown University
avd@cs.brown.edu
www.cs.brown.edu/research/graphics/home.html

Three-dimensional environments depend even more heavily on the quality of HCI than 2D environments, and yet the best way to handle interaction in immersive, augmented, and even desktop 3D environments is far from clear. I will provide an update on our latest research in using gesture-based techniques, and then outline some of the key challenges of the current computing landscape. For example, a milestone will occur within a year after SIGGRAPH 98, when 3D hardware will become universally available, even on commodity platforms. But this opportunity is not supported by any convergence on a software standard for 3D and other media on the Web.

The development community is forced to choose among a number of options, such as VRML 2.0 and several APIs which are not yet in general circulation (Java 3D) or even specified (Fahrenheit, being defined by Microsoft, SGI, and HP). Even if there were agreement on the API, there still is not enough collective experience on 3D interaction techniques, metaphors, and widgets to allow anything like the "standard" WIMP GUI to appear. I will discuss some possible outcomes of this debilitating uncertainty.

Dan Venolia
Cosmo Software
dvenolia@sgi.com
www.cosmosoftware.com

I have prototyped several forms of 3D direct manipulation: using 2D and 3D pointing devices, 2D and 3D cursors, with and without Brown-style widgets, etc. My longstanding goal has been to make 3D accessible to "mere mortals" – those with traditional graphic arts background or with only basic computer graphic skills.

My current project, Cosmo PageFX, is a tool that allows Web graphics professionals to create interactive, animated 3D graphics. It does so in two ways. First, we have the liberty of simplifying the 3D arrangement problem by removing the possibility of camera motion. Second, we use tools that are similar to those used in traditional 2D graphic editors. The result combines Brown-style widgets with a page metaphor, gleaning the best of both worlds.

I compare our design with that of tools that are aimed at the 3D-savvy. I talk about the special problems introduced by eliminating camera motion and introducing a page metaphor. I suggest how this approach can be extended to a broader set of problems.

Listen Up! Realtime Auditory Interfaces for the Real World

Moderators

Maribeth Back
Xerox PARC

Elizabeth Mynatt
Xerox PARC

Panelists

Perry R. Cook
Princeton University

Robin Bargar
Virtual Environments Group, NCSA

Peter B.L. Meijer
Philips Research Laboratories

Designing Interactive Auditory Displays

Maribeth Back
Xerox PARC
back@parc.xerox.com
www.parc.xerox.com/back/

Auditory display pushes the boundaries of audio at the interface, exploiting the auditory modality to deliver multi-dimensional information in an efficient and intuitive way. This panel covers some of the most interesting current tools, methodologies, and applications from leading research institutions and industrial labs, including audio-augmented reality, sound design, physical modeling systems, cross-modal sensory affordances, and audio in immersive environments.

- From Xerox PARC, an audio-augmented reality system designed to provide peripheral awareness information in the workplace, as well as a novel approach to sound design for interactive systems.
- From Princeton, a tool that provides access to the parameters of a physical modeling system for sound, which allows users to map data to any of a number of auditory inputs.
- From Philips Research, a close look at cross-modal sensory affordances, which allow transformation from one sense, such as sight, to another, such as sound or touch.
- NCSA's sonification system, was developed for the CAVE immersive environment, which allows people to experience datasets as multimodal constructs of graphics and sound.

At some level, all these systems and practices address one of the hardest questions: what to put in? The discipline of sound design consists of creation and manipulation of physical and conceptual structures including speech, music, sound effects, and ambiences. In any effort to derive the mechanisms that support these structures, we must consider context, human perceptual and cognitive capabilities, constraints and affordances according to media type, and symbolic and semantic systems.

A design methodology uniquely suited to sound in interactive systems can be created by combining what psychoacoustic and perceptual researchers tell us about human perceptual mechanisms with what we know about the cultural mechanisms surrounding sound. If we think of sound as a type of narrative, we can construct sounds and soundscapes as though they were elements in a story. Narrative context and content provide the user with ready-made schema for interpreting sonic events as the designer intended, thus providing a set of shortcuts for communication in the real-time interactive environment.

Audio Augmented Reality

Elizabeth Mynatt
Xerox PARC
mynatt@parc.xerox.com
www.parc.xerox.com/mynatt/

The perception of sound and the perception of place are powerfully intertwined. How we perceive a space, its dimensions, its textures is strongly influenced by how the space sounds. Additionally, its sound tells us how the space is used, its "placeness," whether it is formal, informal, crowded, chaotic, or calm. For these reasons, we can use audio to change a person's perception of a physical place. By augmenting the natural sounds of a physical place, we can add to its richness and utility.

At PARC, we are exploring a system that focuses on connecting people, places, and their computers by providing auditory awareness cues. We leverage the physical environment to trigger the delivery of information. As a user moves through the office place, entering the coffee room or pausing at a colleague's office, information that is collected computationally is summarized and sent to the receiver. The second part of our strategy is to present information via metaphorical auditory cues that mimic the peripheral auditory cues people constantly process in their normal environment. By using physical-world triggers and auditory cues, we are creating a light-weight interaction that does not require active participation by the user.

Our system, Audio Aura, is based on three known technologies: active badges, distributed systems, and digital audio delivered via portable wireless headphones. An active badge is a small electronic tag designed to be worn by a person. It repeatedly emits a unique infrared signal and is detected by a low-cost network of IR sensors placed around a building. A location server combines all the information culled from the IR sensors and augments it with other information such as online calendars and email systems. The delivery of audio cues is triggered by changes in the location database. Digitized sound is converted to analog and then sent to the user's wireless headphones. We are exploring using a variety of sounds (natural, musical, and voice) that complement daily activities and blend into the existing aural backdrop.

Auditory Display Using Real-Time Sound Synthesis and Processing

Perry R. Cook
Princeton University
prc@cs.princeton.edu

The nature of many sounds in the real world is one of generation by non-linear and/or random dynamical processes. As a result, we as human listeners expect certain behaviors in sounds in response to certain types of changes in the parameters of the sound-producing system. Simply attaching static beeps, blops, and recordings of sounds to interactions in a virtual environment or auditory display, even if those recordings are of real-world events such as glass breaking or horns honking, will in many cases be insufficient to match the richness of our normal auditory experiences.

Flexible, parametric sound synthesis algorithms and tools should be available as the basis of most auditory display systems. Pulse Code Modulation (PCM) waveform playback should be treated as only one single algorithmic member of a much larger palette of synthesis and processing algorithms. Physically based synthesis algorithms are now possible for a variety of sound producing objects, including a new class of random-particle-based models, which can closely approximate many common interaction sounds. With models such as these, it is possible to perform direct mapping of parameters such as effort, hardness, etc., yielding more natural virtual auditory experiences and mappings for abstract auditory displays.

Sound Authoring for Real-Time Synchronous Display in Immersive Environments

Robin Bargar
Virtual Environments Group, NCSA
rbargar@ncsa.uiuc.edu

The meaning of sounds or images can undergo radical changes when they are placed together in a display. We are accustomed to cinematic combinations of picture and sound that are designed to support the fictional power and objectivity of the camera required in movies. Emerging as an alternative medium are interactive displays, which provide a subjective experience of computational spaces such as simulations, databases, or immersive environments. In multi-modal systems sound contributes a temporal refinement that exceeds the frame rates of graphical displays. When a display is real-time, interactive, and data-driven, the sounds and images can be modified by observers who are able to adjust their actions to accommodate the dynamics of events as they are generated, thereby optimizing the display.

The combination of sound, image, and action in an interactive cycle of observation requires a system to support asynchronous parallel processes such as differences in the cycle rates for simulations, control loops, and image and sound rendering rates. At the same time, the synchronous linkage of display events must be maintained when actions are applied. We can think of this as a system to support human-computer performance, where the actions of an observer are understood as a form of time-critical performance, similar to the actions of musicians who make constant adjustments to their instruments while listening to the consequences at each moment. Sound authoring is a process of creating conditions for sound production in a real-time performance system. A sound authoring system is demonstrated and made available for hands-on exploration in the Creative Applications Lab.

Cross-Modal Sensory Streams

Peter B.L. Meijer
Philips Research Laboratories
meijer@natlab.research.philips.com
ourworld.compuserve.com/homepages/Peter_Meijer/

The rapidly increasing computational power of multimedia PCs is beginning to enable real-time transformations of complex (real-life) sensory streams in software. An exciting challenge is to find out if real-time cross-modal sensory mappings could help in dealing with sensory disabilities (for example, to feel or hear images if one is blind, to feel or see sounds if one is deaf, or to feel sounds and/or images if one is deaf-blind). This presentation focuses on developing general and affordable auditory displays for the blind, using portable equipment (wearable computing). A demonstration of real-time video sonification will be included, based on an approach named "The vOICe," now running on a regular PC with a Webcam.

In using cross-modal sensory mappings, there are inevitable trade-offs between information preservation (both space and time resolution), aesthetic acceptability, and limitations in human perceptual capabilities. The technology for affordable cross-modal mappings is (almost) there, but little understanding currently exists about what mappings and mapping parameters would be best under what circumstances, or what the actual added long-term value of any given mapping would be to the disabled user. Cooperation among engineers, neuroscientists and psychologists to further evaluate the options would be a logical next step. A key issue is that learning to exploit new information-rich auditory displays may require a major training effort, while one does not know in advance if the resulting human performance level would indeed justify that effort.

Out of the Box: Toys Break The Screen Barrier

Moderator
Steve Schklair
Quantum Arts

Panelists
Christian Greuel
LEGO A/S, SPU-Darwin

Andy Rifkin
Mattel Media

Erick Strommen
Microsoft Corporation

Michael Patrick Johnson
Massachusetts Institute of Technology

Steve Sutyak
Hasbro Interactive

Steve Schklair
Quantum Arts
steve@quantumarts.com

Transmedia is a new genre of software-based product that only recently has emerged into the marketplace. These applications blur the line between physical play in the real world and virtual play in the digital world. Are these applications precursors to the eventual integration of the computer into daily life? Just as ATM's broke through a major sociological barrier to widespread acceptance of the computer as an appliance, will these products become the next vanguard of consumer acceptance?

Interestingly enough, these applications are primarily children's entertainment products. As they distribute the play experience over different media, are they still considered "applications," or have they become "toys" in which the computer is now only part of the total experience? This is not just a matter of producing branded spin-off products, but a new form of entertainment in which part of the play experience occurs onscreen and part occurs with physical objects, or within social or out-of-the-home scenarios. Examples of released products that began and continue to define this new genre are Barbie Fashion Designer and Talk with Me Barbie from Mattel, and Microsoft's Interactive Barney.

This panel features the people behind these releases and attempts to focus on issues such as:

- Are toys the trendsetters in this new market?
- Do these products integrate the computer more into the lifestyle of today's families?
- Do these products expand what the computer can do and begin to change our perception of the computer to more of a household appliance?
- By enabling activities that are also tactile, creative, and social, do these products defuse the criticism that we are turning our children into vegetables by fostering a generation of computer addicts?
- Will this new genre be limited by the peripheral market of input and output devices?
- Is this market limited to children's toys, or are there adult applications in the works as well?
- Does the potential for alternative distribution channels make these products more appealing to publishers and developers?
- What are the educational, social, and entertainment benefits of these new products? Is this just another marketing and merchandising ploy to get kids to buy more stuff?

Christian Greuel
LEGO A/S, SPU-Darwin
chris@digi.lego.com

When Ole Kirk Christiansen began making wooden toys by hand back in 1932, he might never have dreamed of the day when his humble company would move into the revolutionary new age of plastic. But in 1947, the LEGO Group purchased its first plastic-injection mold and soon began producing "Automatic Binding Bricks", the predecessors of today's classic LEGO Bricks.

Today, we face a new opportunity. SPU-Darwin was established in 1996 as a special project to explore the realm of possibilities for the company as it enters the era of digital toys. Our work includes research and development of digital technologies based on and related to LEGO products. This includes everything from Web pages and CD-ROMs to visual simulation and artificial intelligence.

We have a strong interest in the interplay between the real world and computers, as is visible in the Mindstorms intelligent brick product. We are also looking into new approaches to the digital playspace, allowing kids to enter the computer in a manner more analogous to traditional play.

Children today are growing up with computers, but why are we forcing them to adapt to tools designed for adults? Toys are a natural place to evolve fundamentally solid human-computer interfaces. We have much to learn from the way children interact with their toys, their world, and each other. They do not need a desktop, but rather a space in which to play.

Kids can play fine with sticks and mud, but computers still promote solitary game-play. Our focus should not be on products, so much as on the tools that allow children to imagine and construct their own toys. New paradigms need to be developed that allow free-form creativity with and within the computer, in social collaboration with friends.

In this panel presentation we are looking at LEGO's Magic Window, a test-bed for many of our ideas. We hope to answer many questions regarding how children play and interact with technology, as well as raise new questions along the way. It is our hope to be able to turn tools into toys and toys into tools, effectively blurring the line between computer technology and children's playthings.

Andy Rifkin
Mattel Media
rifkinan@mattel.com

The continual evolution of computer technology has resulted in the development of increasingly more affordable and more powerful home computers. Computers have become a tool for the masses, for entertainment even more than education; for fun even more than for functionality.

My goal is to bring a new kind of fun to this virtual playground – the kind of fun that centers around friends and family. This is what we accomplished with Barbie Fashion Designer. We created an opportunity for parents and grandparents to play creatively with their children.

At Mattel, we are harnessing technology so that we can create tools for collaboration as well as entertainment and learning. Our magic is that we are using technology to embrace the deepest virtues of play. The fun of turning a child's dream into reality. The fun of sharing in the creation and realization of that dream.

Erick Strommen
Microsoft Corporation
erikstr@microsoft.com

There are a variety of ways to think about transmedia integration. The most common way is to conceive of it as a form of synthesis, a merging of the properties of formerly distinct media together into a new form. A different viewpoint is to think of media integration as a way of using interactivity to complement existing media use, instead of replacing it with something new. In this model, individual media maintain their distinct features but interactivity is deployed to deepen and strengthen children's understanding in each mediated experience.

This is the model used in the development of the ActiMates interface and content. ActiMates is an interface that uses the social dynamics of pretend play to integrate technology and learning. Because they are animated plush dolls who resemble and behave like familiar media characters, ActiMates tap into powerful pretend play and toy experiences common to early childhood. Using speech and movement, ActiMates utilize social responses as an interface strategy in order to enhance children's engagement with electronic media like television and the PC.

The goal of ActiMate design is to use expectations of social behavior, combined with the differential responsiveness of interactive technologies, to provide scaffolded learning experiences for children in different media. The ability of ActiMates to interact with the content of television programs and PC interactions allows them to augment the media experience in educationally valuable ways, acting as an intelligent peer or adult would. As a co-viewer during TV viewing, the ActiMate models "active viewing" by reflecting on specific onscreen events, asking questions, and directing attention. As a learning partner at the PC, the ActiMate is able to offer praise, suggest hints, model appropriate performance, and more. ActiMates rely on established principles of social learning and peer learning with media to enhance existing media experiences.

ActiMates are a form of cross-media convergence that combines the properties of physical and virtual interfaces to achieve specific educational goals: improving learning through play experiences and bolstering media literacy. This model has the virtue that it can be applied to new interactive forms, such as the Internet and interactive television, as they become part of children's expanding suite of media experiences.

Michael Patrick Johnson
Massachusetts Institute of Technology
aries@media.mit.edu

An avatar is the virtual representation of a human in a virtual world. What is the analog in the physical world for a virtual creature? Does it make sense for a virtual character to have a real-world avatar? Are the multiple representations of the same entity confusing? We have been exploring these issues in the synthetic characters group at the MIT Media Lab. Having struggled with trying to control a complex virtual character with a mouse and keyboard, we embedded sensors in a stuffed animal. The augmented toy could then be used to control a user's avatar in a direct and obvious manner, rather than relying on complex mapping of buttons and keys.

As we used this interface device, we began exploring the idea of using the stuffed animal as the physical avatar for the virtual character. Several members of the synthetic characters group were involved in building the ALIVE system, in which a participant could interact with a virtual dog using computer vision. One of the important lessons we learned was that people wanted to "feel" the character. They wanted a physical instantiation of the character so they could pet it. The stuffed animal gives us the opportunity for this type of physical interaction. For example, if the virtual character were frightened, a child could pet it to calm it down.

Finally, the toy retains its original functionality. It can be used in traditional play without the computer involved at all. It's not exactly comfortable to bring a computer to bed at night or to hug a keyboard. By mixing the physical toy media with the virtual, we can hopefully leverage the strengths of each. The computer lets us make expressive and interactive virtual characters without needing to wonder about issues of robotics control, cost, and safety if we were to fully embed the character in the physical world. The physical toy grounds the interaction with a virtual character, in some sense letting us really "touch" the virtual character, rather than relying on the relative coldness of a mouse click or keypress.

Feature FX: Money Pit or Gold Mine?

Moderator
Patricia Rose Duignan
VFX Marketing
Management Consultant

Panelists
Scott Ross
Digital Domain

Jim Morris
Industrial Light & Magic

Carl Rosendahl
Pacific Data Images

Richard Hollander
Blue Sky | VIFX

Phil Tippett
Tippett Studio

Ray Feeney
RFX, INC.

Scott Ross
Digital Domain

Like Sisyphus, effects houses continue to push the stone up the mountain every year. On a marquee visual effects blockbuster, their clients have the opportunity to reap great rewards. But with little to no profit, huge risks, unbelievable capital needs, increasingly higher salaries, we need to own the content we create. Visual effects studios must become producers to survive.

Carl Rosendahl
Pacific Data Images

As recently as 1991, the only production company with any formal ties to the filmmaking community was ILM, and no one had significant outside financing. Contrast this to today where 80-90 percent of the larger companies have ties to the studios, filmmakers, or substantial outside investors. Why is this?

I believe the current market supports two kinds of production companies: small boutiques and large facilities. Small boutiques have the advantage of carrying limited overhead and the ability to move quickly. Large facilities have the ability to take on large complex projects. The mid-sized company is at a disadvantage from both sides, and therefore it is not a good place to be. Maintaining a large company involves substantial capital investments and an ability to take risks. In a highly competitive, low-margin, creatively driven business, this necessitates bringing in outside partners, and, in fact, that is what has happened in our industry.

But the game isn't over. Technology continues on its steep improvement curve, and the entertainment industry is still learning how to take advantage of it. The tumultuous ride we have had in the past two decades will continue.

Ray Feeney
RFX, INC.

The Renaissance is over, and a Dark Ages is beginning. The ability of computer graphics service companies to survive under today's profit picture is in doubt. The fixed costs involved do not match the traditional motion picture model of "for the run of the show." Only the very largest facilities and the smallest boutiques can hope for continuous activity, and neither can completely fulfill the visual needs (and price points) of the current industry demands. Couple that with the stress of transitioning from Unix to NT (without clear leadership and investment by manufacturers) and the inability to scale solutions to put the dollars up on the motion picture screen, and all indicators point towards a period of retrenchment.

Phil Tippett
Tippett Studio

A visual effects approach to creating and solving cinematic issues has become a critical part of the filmmaking process over the last two decades. As a result, visual effects companies have grown from garage shops to empires. The normal 50-shot shows of the 70s and 80s have given way to the 500- or 1000-shot shows of today. The organization, personnel, and facilities required to accomplish shows of this scope have revolutionized our business and invited deep-pocketed corporations into our little garages. For better or worse? That remains to be seen.

Jim Morris
Industrial Light & Magic

Over the last five years, we have seen dramatic transformations in the visual effects industry. We have experienced explosive growth, and gone through tumultuous shakeouts.

But though the visual effects industry now books hundreds of millions of dollars per year in revenues, the heart of it remains the same: artists and craftsmen making great images to help tell stories. And the digital tools available for creating effects now let us make those images better than ever before. The palette available to filmmakers for creating characters, settings, and events is nothing short of astonishing.

This palette, along with the continued box office success of effects films, ensures a lively future for the visual effects industry. Since so much of it is done for love instead of profit, the industry will likely remain marginal as a business. But it will remain and grow as a vibrant, essential step in the making of motion pictures.

Richard Hollander
Blue Sky | VIFX

The bad part of the visual effects industry:

It is extremely competitive. Small houses are able to do quality work with low overhead. Work definitions are usually vague. “Build me something I have never seen before, and if I like it I will approve it.” Time schedules usually shrink while working on the project. The talent pool is small and very expensive. Since the industry is art-based, profit margins are extremely small at best. It is capital intensive. It requires lots of research and development. Technological advances make the result more impressive but it still takes just as long, if not longer, to produce the result. Hardware companies design their equipment for other markets, and we have to adapt that hardware to our industry. Software is no longer built by the industry practitioners. Software is developed by large companies whose markets are larger than our specific industry.

The good part of the visual effects industry:

It is fun.

This enthusiasm has spun the effects industry into a buyers’ market, which has led to the creation of “small” and “large” houses, leaving no middle ground. The industry has not been idle in response to these outrageous odds. The current trend has been the alignment and purchase of visual effects houses by larger companies. Effects houses are beginning to originate and retain ownership of content, thus leveraging their in-house talent. I believe this will lead to an increase of funds targeted for substantial research, promising imagery that we have never seen before.

Dis-Illusion of Life: Becoming a Digital Character Animator

Organizer
Barbara Mones-Hattal
Industrial Light & Magic

Moderator
Jacquelyn Ford Morie
Blue Sky | VIFX

Panelists
Endla Burrows
Daniel Jeannette
Industrial Light & Magic
Pete Docter
Pixar Animation Studios

Ken Perlin
New York University

James Sayers
Sheridan College

While talented digital character animators are constantly in demand, they are very hard to find. This shortage has been felt most keenly in the past few years, as more and more studios have started creative projects requiring animation talent. Why is there a shortage when more and more schools have implemented programs to train new animators, and the technology has matured to provide very user-friendly and sophisticated tools for animation?

Is there a secret formula that creates a successful digital character animator? Analysis of the backgrounds and training of established animators reveals that some came through trade schools, some were trained on the job, some taught themselves, and a small proportion were educated in a streamlined and structured program that merged classical animation training with newer digital tools and processes.

Educational institutions can and do provide graduates with creative approaches to problem solving and the potential to mature into employees who provide invaluable contributions. At many schools, however, there is still a general lack of understanding about what skills the industry requires and how to train for these skills. The result: hiring companies must devote significant resources to training recent graduates.

In this panel, representatives from educational institutions, animators in major production studios, and industry trainers present widely divergent views on these topics. They also provide insight and suggestions for animators, those who train animators, and animators looking for employment in the field of digital character animation.

Endla Burrows
Industrial Light & Magic
endla@lucasdigital.com

Where are the animators? High-end visual effects houses are continually faced with the need to put new employees, especially animators, through extensive training before they can be placed on any projects. Even then, there are often gaping holes in the animator's understanding of the craft.

Many institutions of higher learning have the facilities to provide an excellent educational foundation for animators. Unfortunately, most fall short by choosing to focus narrowly on specific animation programs or packages and eschew expanding on the fundamental concepts that underlie the 12 principles of animation.

To be able to create believable characters for film, an animator needs a basic understanding of anatomy, physics, and film and computer usage. The necessary courses exist at many schools. It seems that the schools simply need to find an advisor who can direct the budding animator to an interdisciplinary course of study that would better serve both the student and our industry.

Daniel Jeannette
Industrial Light & Magic
dj@lucasdigital.com

Digital animation and character animation share a common requirement: the ability for an animator to understand the kinematic and timing principles of a moving object to efficiently convey its physicality on screen. However, character animation goes beyond this. It uses the same foundations but requires much more than just applying the right amount of kinematic motion, or physically based motion simulations. Even perfectly capturing the motion of a human performer through motion-capture technology doesn't quite create character animation. It is merely the near "perfect" replication of a "realistic" movement, but without the soul and spark of a true character.

In addition to the visual statement and insights that any artist will give to a character, character animation is truly part of the story structure. It is the ability to communicate thought process and emotion, and the ability to establish the personality of the animated character with a strong artistic and dynamic point of view.

Picture composition, image structure, kinematic and cinematic language, acting skills, and animation principles are the ground training of a great character animator. These are the prevailing skills that are most sought after in a production environment. Animators need to develop the ability to visualize and strongly communicate a concept or the performance of a character. Computer training is useful in as much as it allows students to understand how to use new tools to transfer their knowledge into the digital realm.

Peter Docter
Pixar Animation Studios
doc@pixar.com

Not all animation involving characters can be called "character animation." True character animation gives you a look into the brain and heart of the character; you see it thinking and feeling. Production of quality digital character animation requires a very specific set of skills, including those of an actor, artist, designer, and choreographer. Sadly, there are only a handful of schools that teach these skills. Most computer animation training programs focus instead on computer and software literacy – skills of secondary importance to the digital character animator.

Ken Perlin
New York University
perlin@nyu.edu

I think it is going to be increasingly important for art departments and computer science departments at universities to join forces and create a continuum of offerings for students in subjects related to digital character animation. Within academia, there is currently far too much cultural separation along traditional lines, at a time when the technical and artistic challenges of this problem need to be addressed by a new generation educated to think both technically and artistically.

Curricula should include active collaborations that allow the student significant freedom to migrate between a technical and an arts orientation. This program should start with a core interdisciplinary course sequence that teaches the relevant principles of visual design and character description, and the basic technical foundations of kinematic, dynamic, and behavioral modeling. It should be taken by both students in computer science departments AND in art departments who have an interest in learning about digital character animation. This will give students a proper foundation from which to effectively specialize later on.

I believe that academic structures will migrate to this model (after some growing pains), simply because the jobs are out there, and the needs of this emerging industry demand this kind of thinking.

James Sayers
Sheridan College
james.sayers@sheridanc.on.ca

There is no question that a good digital character animator requires all the knowledge of the classical tradition and that of a computer animator. The demand for specialized digital character animators is bringing Sheridan College's Post-Diploma Computer Animation and Classical Animation Programs together. In recognition of this, the computer animation program is moving away from a more generalist approach to one that recognizes the distinction and need of specialization in the character and technical disciplines.

Human Factors in Virtual World Design: Psychological and Sociological Considerations

Moderator
Elizabeth Reid Steere
Royal Melbourne Institute of Technology

Panelists
Lynn Cherny
AT&T Research Labs

Mary Czerwinski
Microsoft Research

Tammy Knipp
Florida Atlantic University

Beth Kolko
University of Texas at Arlington

Elizabeth Reid Steere
Royal Melbourne Institute of Technology
elizrs@mediaone.net
members.xoom.com/elizrs/

During the Applied Virtual Reality full-day course at SIGGRAPH 97, Rudolph Darken commented to the effect that "VR designers strive for realism because it's easy to evaluate, and because we don't really know what's necessary or expendable for usability."

This statement struck a strong chord in me, since I had recently been invited to give a talk on psychosocial factors in multi-user graphical virtual world design at Microsoft Research. Darken's comment reinforced my belief that designers of graphical virtual words both needed and would welcome contributions to their field by those working in the areas of sociology and psychology. My aim in organizing this panel is to introduce virtual world designers to the ideas and design strategies suggested by a number of researchers within these areas. The panel focuses on methods of successfully designing virtual worlds intended as social or community platforms, with special consideration to design choices that are at odds with encouragement of user commitment and social interaction.

The members of the panel have diverse backgrounds in areas with which many technically oriented virtual world design professionals may not be familiar. Lynn Cherny and myself have extensive backgrounds in the study of human behavior in text-based community-facilitation systems such as MUDs and IRC. Tammy Knipp comes originally from a background in advertising and is now a well-known artist working in computer-generated media. Mary Czerwinski works in the User Interface Group at Microsoft Research. Beth Kolko has come out of a background in rhetoric and writing theory to address issues of communication and community construction online.

We would like to advance the idea that lessons in psychosocial design considerations learned from such diverse areas as advertising, art, rhetoric, and text-based computer-mediated interaction can be extremely useful when building graphical virtual worlds. Many of the concrete design methods formed out of experiences with other systems in which the illusion of reality is central are readily translated into graphical media. It is the psychosocial effect and affect that are important to each medium rather than the technological particulars of any one system. By examining what works in one system, many lessons may be learned about how to succeed in another. We hope that this panel introduces virtual world designers to theories and methods that may serve them in designing better systems.

Representations and Agency in World Design

Lynn Cherny
AT&T Research Labs
cherny@research.att.com

I discuss two principles that I believe are important for designers of multi-user worlds on the Internet: reliable shared representation of events and reliable representation of agency. The principle of reliable shared representations is related to the usability concept of WYSIWIS (What You See Is What I See). Users of a virtual world should be guaranteed that when something happens, everyone who witnesses the event sees the same thing (although perhaps with different degrees of detail), or else is positively notified otherwise. The representation of the activity is therefore a reliable representation. They can safely assume that their interface's perspective is a "true" one, and the evidence on which they base their understanding of events is guaranteed to be the same as other users'.

Related to the shared representation concept is that of reliable agency identification: users should be able to determine who produced an event and how it was produced within the virtual world. Events may originate with a user or with the system itself, or from some interaction between the two. Simple communication events (for example, what a user's avatar says) are the events most directly produced by a user.

But some communicative actions in some MUDs and chat systems are produced with menu commands that output speech-like utterances (for example, "*Mary smiles.*"). As far as the user witnessing these events is concerned, these are confusingly indistinguishable: it's all just text. Graphic systems in which avatars are controlled manually or by macro- or meta-commands that do path-finding, for example, may also confuse other watching users. "Spoofs," user utterances unattributed to a user or falsely attributed to another user (possible on some MUDs and Palaces), are the most problematic and potentially abusive violations of this principle.

The significance of reliable agency representation in virtual systems is at least twofold. First, it makes users accountable to each other for their actions. This benefits the community by reducing the likelihood of transgression against its norms. Second, the ability to identify the source (the who and the how) of an event makes it easier for new users to understand what's going on; this is a learnability concern that impacts community growth. Both reliable shared representation and reliable agency representation are important for the establishment of membership in a community: they allow acculturation based on reliable information about causes, and they reduce ambiguities generated by poor system design rather than human intention.

Beyond Designing For Usability

Mary Czerwinski
Microsoft Research
marycz@microsoft.com
www.research.microsoft.com/research/ui/marycz/home.htm

Virtual world design today presents many challenges for the builder: Should a 2D or 3D user interface be used? What's the appropriate depth and breadth of the worlds? What amount of end-user editing should be provided not only in the world, but also in the avatars and objects utilized in those worlds? How best should communication, emotion, and navigation in the world be supported through the UI, etc.?

The designers of these worlds need not only worry about the beauty and usability of these worlds, but also the social parameters that are responsible for fostering a sense of community and belongingness on the part of the participants. The latter issue is a particularly sensitive matter, and is at least as difficult as designing for ease of use. For example, no consistent, reliable rules have as yet been agreed upon for how to instill a sense of community on the Web, and only over the last few years has a sizable body of literature become available on the subject. Guidelines, experimental methods, and measures are sorely needed so that designers can build virtual worlds from a user-centered perspective. To that end, this panelist argues that what is needed is to run empirical studies, longitudinally if possible, exploring issues of usability and community in virtual worlds. Only through careful analysis, cataloguing, and comparison with knowledge available from sociology and psychology will useful constructs, measurement methods, and design principles emerge. Recent findings from our work studying virtual communities are described, as well as our preliminary design guidelines grounded in this empirical data.

Peak Experiences In Virtual Environments: A Sudden Surge Of Meaning

Tammy Knipp
Florida Atlantic University
tknipp@fau.edu

My primary interest in a discussion of virtual reality lies in investigating the underlying source, motive, or drive that creates a longing for contact with, or escape from, reality. To examine these motivational factors, I must broaden the context of virtual reality, in part by borrowing from the psychoanalysis of advertising.

Advertising appeals to unconscious defenses, modeled on such mental habits as projection, displacement, identification, and repression. The discussion I bring forth is not in the realm of manipulation or persuasion but the borrowing of methods from the virtual and imaginary worlds advertising creates. For example, marketers use advertising messages that appeal to our sense of risk-taking adventures. However, virtual environments that amputate fear from danger make the risk-taking encounter a leisure activity. If the concept of risk is defined as a fear of loss, or if something of value is at stake, then the element of fear must accompany the anticipation of danger.

Researching motivational factors, human behavior, and stimuli, I discovered characteristics to apply in the virtual realm. I began seeking a forum whereby technology would facilitate or perhaps instigate a "social happening" and create "peak experiences" while encouraging elements of laughter, humor, and play. As an electronic/video sculptor with a background in advertising design, I create CASE STUDIES comprised of 3D structures bridging and integrating the dimension of video imagery with the realty of the physical and social world. These constructed realities are virtually perceived from a haptic, kinetic language approach. Raising issues of belief and perceptions of trust, the demarcation between virtual risk and real risk (virtual reality and reality) breaks down. Bringing forth research and motivational theories, I present video clips illustrating characteristics of: laughter and play behavior; selective optimum stimuli; kinesic behavior and social language; and perceptions of risks with psychosocial factors in virtual environments.

Mapping Real Success for Virtual Worlds: The Rhetoric of Space and Interactivity

Beth Kolko
University of Texas at Arlington
bek@uta.edu

Rhetoric is the science of communication; the rhetorical tradition maps the relationship among a speaker, an audience, and a message. The overall lesson of rhetoric is that the context of a communicative act is key. If we examine virtual worlds as a particular kind of rhetorical situation, we see a whole new way of understanding what goes on in virtual spaces and how successful virtual spaces can be created. Rhetoric is a way of understanding the dynamics of written, spoken, and unspoken language; a rhetorical analysis shows how individuals signal a sense of self-identity, relate to one another, create or avoid conflict, adopt external and unspoken cues such as dress or mannerisms to position themselves in a larger context, and, finally, how they use place as an element in communication. If the purpose of creating commercial virtual worlds is to build spaces that invite participants to spend time (and money), then those spaces must be geared to successful community building; they must accommodate public and private lives; and they must encourage a variety of types of communication so that participants can express the complex sense of identity necessary to sustain interaction. Part of building such worlds is understanding how space and geography affect communication patterns.

Rhetorical theory grew out of considerations of how place could be tied to the purpose of communication. Space matters substantially precisely because when people communicate they have a sense of self within the context of language. Words are not disembodied from the person offering them, and this relationship must be acknowledged in the design of virtual worlds. The physical world is always present when we communicate – the world of the person and the world of the place. Physical space affects communication (for example, consider the Greek agora, a public space whose architecture was keyed to the goal of widespread public debate), and a speaker in any exchange, a virtual world conversation or a face-to-face exchange, will find that location affects speech: what is said, what is heard, what is ignored, etc. Location may not be everything, but it is crucial to communication, and, consequently, the success of a virtual community will in part hinge on the particular virtual geography of the landscape. Rhetorical analyses hold the key to effective design of space and place in virtual worlds.

Virtual Reality as Healing Art

Moderator
Galen Brandt
Virtual Healing

Panelists
Dorothy Strickland
Stetson University

Richard Satava
Yale University

Hugh Lusted
BioControl Systems, Inc.

Tom Riess
HMD Therapeutics

Rita Addison
Virtual Reality Artist/Consultant

Myron Krueger
Artificial Reality Corporation

Hunter Hoffman
Human Interface
Technology Lab

Galen Brandt
Virtual Healing
GalenB42@aol.com

What does it mean to be healed, in the sense of whole, and wholly self-expressed, human beings? How can we use virtual reality – virtual selves, relationships, art forms, spaces – to heal ourselves in body, mind, and soul? Who are the pioneers of cyberspace healing, and how can they help return us to wholeness? What is the most profound work in this exploding field, from classic vision to cutting-edge research, realworld healing to realtime art?

This panel looks at cyberspace as healing place, avatar as anima, virtual reality as healing art. For in its deepest promise and most profound practice, virtual reality – as both healing modality and visionary art form – is a transformative technology with extraordinary power to make and keep us well.

Suppose we can barely walk. Then suppose we put on VR glasses that "augment" the world and find we can walk again. Suppose we can barely move at all, yet find that, by using a VR biocontroller to reconnect to our "body electric," we can play computer games, or the violin. Suppose we have been brain-injured and our vision has been impaired, yet in using VR to show others what it is not to see, in making it impossible for them to be blind to our pain, we enlighten both them and ourselves. Suppose we have autism and cannot learn real-world skills, yet in simplified virtual worlds, we can learn to cross a street and use a fork. Suppose we are dying in Bosnia, and the "digital physician" who saves our life operates from Boston. Suppose we have forgotten how to play ... yet as virtual selves, we find we can float and fly, dance and dream as reborn children.

These are healings that virtual reality alone makes real. For in VR, as in no other technological or artistic practice, we both see and embody the virtual selves of our needs and dreams: we become what we behold. As we cross streets, climb mountains, play, fly, bare our kidneys, and make love and music as new virtual beings, we are giving ourselves positive, chemical messages about what is and can be real for us.

This is not metaphor. This is literally, neurochemically true. Consciousness creates the body. Your biochemistry results from your awareness. To give yourself a new message is to become that message, down to your neurons. In beholding ourselves as healed – in becoming our self-visualizations – we become the selves of our deepest and most healing dreams. Belief becomes biology; the technological, the transformational. This is a revolution in medical practice. And in bringing what mind/body pioneer Dr. Bernie Siegel calls "ethical hope" to millions, it is the realization of the truest and most artful promise of virtual reality.

VR lets us collaborate with machines (in itself a profound cultural healing) to enhance vision and visualization, empower the imagination, stretch empathy, relink mind and body, engender joy. In recovering these essential, birthright aspects of our humanity, we may not always be cured, but we can be healed. Could it be that this virtual healing is the "killer app" of our digital future – not because it kills, but because it heals?

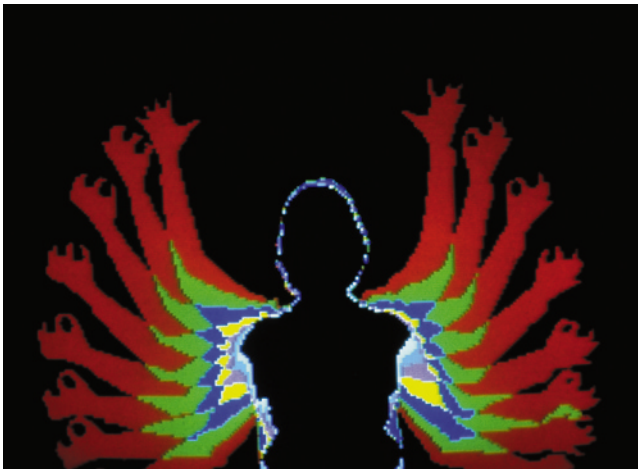

The healing art of Myron Krueger's VIDEOPLACE.

Virtual Reality as Healing Art

Dorothy Strickland
Stetson University
Virtual Reality Aids
dstrickl@stetson.edu

Nowhere is it more obvious that realities are judgments than when treating neurological injuries and abnormalities where the diagnosis of "normal" is relative rather than absolute. Recent literature describes how brain imaging techniques and genetic engineering are redefining sanity from perfect measures to thresholds of clinical significance. You and I may share not an absolute interpretation of the world, but instead, interlacing points of agreement along a continuum of reality. Where life is a web of perceptions, virtual reality can excel as a treatment aid because of its unique ability to use illusions to help measure, manipulate, and reconstruct reality.

Children with autism have shown special promise in accepting and learning via virtual illusions designed to overlap their realities with ours. When I first started using VR to treat these children, I thought I was helping them understand and adjust to our world. But if I in any way taught them how to move in this reality, they taught me much more about what reality is. Their senses randomly gathered a different subset of life than mine, and their highly selective memory system stored in ways still mysterious to me. One autistic child identified an arm not by its form, as we would, but by a stain on the sleeve that in his eyes perhaps matched a more dominant information pattern. The secret to seeing with their eyes, and to successful treatment, was often to suspend my previous judgments of how the world should be separated and sensed – for reality can be understood in many ways.

Richard Satava
Yale University School of Medicine
richard.satava@yale.edu

Medicine has discovered VR and is now realizing the revolutionary potential of this technology to perform surgery that is otherwise physically impossible. Today, all patient information can be digitally input and output. This is medicine's wake-up call to the Information Age – the movement from blood 'n guts to bits 'n bytes. By viewing, manipulating, and transmitting the "information equivalents" of patient organs, today's "digital physicians" can use VR to fly through a 3D rendering of the human bowel in a virtual endoscopy, simulate and practice virtual surgeries (including emergency procedures) in simulated disaster scenes, conduct actual procedures wearing "X-ray vision" glasses that superimpose digital images (such as CT scans) over actual organs, operate remotely via telepresence surgery using a robotic arm, access enormous medical databases such as the National Library of Medicine in mid-operation, and create "the operating room of the future." Soon all of us will walk through scanners in our doctor's office, yielding up instant, realtime, 3D holograms of our inner selves.

This is the new world order for medicine, in which surgeons have extraordinary opportunities to extend and enhance their skills beyond the frail limitations of the human body. The ultimate value of robotics, telepresence, and VR is not that they replace, but rather that they empower physicians to provide better care for their patients, while reducing costs and increasing access to treatments. The future of medicine is not about technology; it is about human caring.

An autistic child practices street crossing on a virtual street.

A virtual emergency room for medical education and training. Courtesy the HITLab, University of Washington.

Hugh Lusted
BioControl Systems, Inc.
lusted@biocontrol.com
www.biocontrol.com

Science fiction writers portray future worlds wherein we humans communicate in symbiosis with computers via neural interfaces to our "body electric." What is the real state of neural interface technology, and more importantly, who is going to use it? My work involves development of neural interfaces that use the electrical energy of the human nervous system as an instrument of virtual communication, control, and healing. The BioMuse "biocontroller" converts bioelectric signals from muscle contractions, eye movements, brain waves, and heartbeat into real-time electronic commands, giving hands-free, virtual control of electronic devices to even the severely disabled. Also a biosignal-to-MIDI converter, the BioMuse can turn even quadriplegic humans into living MIDI controllers who can play music.

This technology has given hope to many physically and neurologically challenged people who have had no access to computers, motorized wheelchairs, synthesizers, and other electronic devices because they could not manipulate the available input hardware (mouse, keyboard, joystick, etc.). My ultimate goal is to create an easy-to-use, low-cost, hands-free, "natural" computer interface that offers access to everyone. Future neural interface devices will benefit from advances in biosignal pattern recognition. Eventually, there will be a full brain-computer interface, although its vocabulary may be limited. Whatever form this technology takes, I hope to empower all humans to get in touch with their body electric to heal themselves.

A quadriplegic child using BioMuse to control a computer with her eye movements.

Tom Riess
HMD Therapeutics
74244.1521@compuserve.com

Seventeen years ago, I was stricken with degenerative Parkinson's disease, and have since developed related gait problems including severe akinesia (frozen movement) and dyskinesia (jerky, uncontrolled movement). These gait abnormalities are extremely debilitating and result in much of the morbidity (38 percent of subjects fall, 13 percent more than once a week) and social isolation associated with Parkinson's. I have devoted myself to researching "kinesia paradoxa": the little-understood phenomenon wherein by stepping over regular, visual "cues" such as stairs or evenly spaced floor tiles, a victim of Parkinson's who is completely unable to initiate or sustain gait can be transformed into a near-normal walking individual.

My objective has been development of a therapeutic VR device – proprietary augmented reality glasses – that evokes the same response of enabling gait by generating the virtual equivalent of these visual cues on demand, in a socially acceptable package, without obscuring the subject's view of the real world. Combining virtual reality, high-intensity micro LEDs, and chip circuitry to generate virtual imagery and project it onto a transparent screen contained within eyewear, I have built several effective prototypes. All meet universal requirements of space, portability, social acceptability, and hands-free control. I have tested these prototypes on myself and others with Parkinson's disease with very encouraging results.

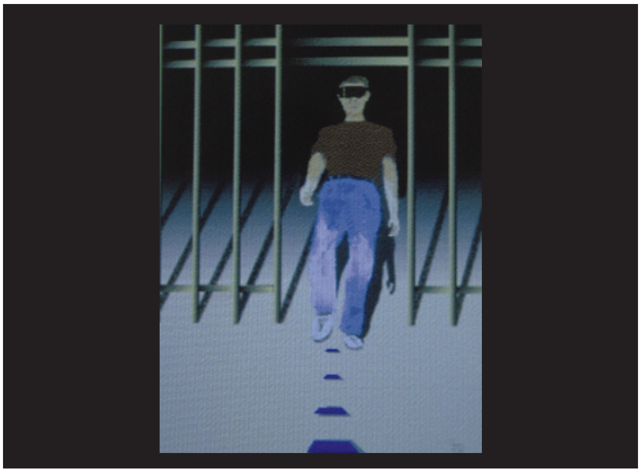

"Virtual Freedom." Painting by Tom Riess.

Virtual Reality as Healing Art

Rita Addison
Virtual Reality Artist/Consultant
raddison@earthlink.net

Virtual reality offers an unparalleled opportunity to be totally immersed inside another's daily perceptual sensorium, to make the invisible visible. My abiding goal is to use VR to create experiences in which participants cross their perceptual thresholds and enter a state in which wonder and awe await to nourish and resuscitate the human spirit. I am searching for the unique attributes of VR that can lead to heightened creativity, insight, and synaesthetically triggered evocation of memory. How do we build a virtual environment that is a responsive organism? What input/output devices can measure signals not only from our outer, physical bodies but also from our inner selves? I want to enable people to create their own immersive, interactive VR experiences. Using VR to tell our stories, to witness trauma in a new way, each of us can become the author of our own healing.

I realized this in a very personal way when a car accident in 1992 left me brain-damaged, perceptually impaired, and unable to continue my work as a psychotherapist and photographer. The need to communicate, to share one's inner experience and be truly understood, is a profound struggle among all head-injury survivors. We feel that neither physicians, family, nor friends comprehend what we're going through. Frustration, hopelessness, and increased isolation were my responses. Eventually my rage and despair fueled my vigorous investigation of VR, the creation of my CAVE installation "DETOUR: Brain Deconstruction Ahead," and the pursuit of my current projects.

Myron Krueger
Artificial Reality Corporation
myronkru@aol.com

The encounter between humans and machines is the central drama of our time. The use of computers has migrated from scientists doing science, to engineers designing products, to users operating applications. The next step is to people being people. The computer will not simply be used as an information appliance; it will interface to peoples' bodies and psyches in ways that are only now being explored.

The human brain evolved to support the body's way of knowing. Yet our culture reveres the intellect and immobilizes the body, making it a vestigial organ. This is a source of individual and cultural pain. In VIDEOPLACE, you move to change what you see, to recreate yourself and your world. When you move in new ways and generate situations for which you have no rehearsed expectations, you can have, just for a moment, a genuine experience, as when your body was new. Moving becomes an act of discovery; the perceptual becomes the conceptual. The mind and the body are reconnected, and everyone becomes an artist. By reinforcing movements that would otherwise be meaningless, VIDEOPLACE can motivate physical activity that sustains function and promotes healing in the physically impaired. By altering how people see themselves and how others see them, it can build self-esteem and develop empathy. These are personal and cultural healings; I believe positive good can result.

The virtual gallery in "DETOUR" showing Addison's photographs before her brain injury.

Tiny Dancer in My Hand, a live telecommunications performance in VIDEOPLACE.

Hunter Hoffman
Human Interface Technology Lab (HITLab)
hunter@hitl.washington.edu
www.hitl.washington.edu

Since 1989, the University of Washington's Human Interface Technology Lab (HITLab) has been investigating applications of virtual and augmented reality technologies in the treatment of medical and psychological disorders. Two recent projects focus on the use of immersive VR for pain reduction and phobia desensitization.

Severe burns are often excruciatingly painful during wound care, even with traditional morphine doses. Pain requires conscious attention and is exacerbated by anxiety. Immersive VR is exceptionally attention-grabbing. Pilot burn patients from my collaborative project with Dr. Dave Patterson's NIH-funded pain research team at Harborview Burn Center show dramatic reductions in pain when immersed in VR. Controlled studies are under way to verify the effectiveness of VR for pain management and to explore whether the amount of pain reduction depends on how present the patient feels in VR.

Psychologist Al Carlin and I are using "virtual therapy" to cure clinical-level phobias. Fear is in the mind of the phobic. VR can be used to help change the way phobics think so they can lead normal lives again. We put severe arachnophobics into virtual "Spider Worlds" designed to systematically desensitize them to their fear. To maximize their sense of presence, patients physically touch a mixed reality spider: part virtual, part hairy toy. After twelve hours of VR exposure therapy using this tactile augmentation, subjects showed a marked reduction in fear which transfers to the real world; our most severe phobic, Miss Muffet, now enjoys outdoor camping for the first time in seventeen years.

VR burn pain management. Courtesy G. Carrougher, Silicon Graphics, Inc. Paradigm Simulations.

A spider phobic receives virtual therapy. Courtesy Mary Levine, University of Washington.

Interfaces for Humans: Natural Interaction, Tangible Data, and Beyond

Moderator
Michael Harris
NCR Human Interface Technology Center

Panelists
Bill Buxton
Alias | Wavefront Inc. and Silicon Graphics, Inc.

William T. Freeman
Mitsubishi Electric Research Laboratory

Mark Lucente
IBM Research

Hiroshi Ishii
Massachusetts Institute of Technology

Michael J. Sinclair
Microsoft Research

Michael Harris
NCR Human Interface Technology Center
mh@mindspring.com
www.ncrhitc.com

Huge advances in interface modalities are evident and imminent. We demonstrate and explore some of the most interesting, promising, and clever of these, as well as their integration into powerful multimodal systems.

When users talk about computers, they usually describe the interfaces, because, for most users, the interface *is* the system. As Bill Buxton says, "The most powerful force in shaping people's mental model of the nature of the beast is that which they see, feel, and hear." It seemed to take forever for toggle-switch panels to evolve into today's WIMPs, and both are still visual/motor-based controls; in fact, switch panels were probably more haptically satisfying! "Keyboards only work for people who know the Roman alphabet. In 20 years, people will laugh at us for calling that technology," says Mark Lucente.

Now, thanks to exponential increases in commonly available computer power and versatility (and concomitant cost decreases), significant progress in interface modalities and their affordances can be perceived. In this panel, we emphasize demonstrable and practical stuff; we have hardware to monkey with, ideas to ponder and try.

This is a gadget-intensive topic, and we present gadgets galore. Input devices that can tell systems where users are looking, the gestures they are making, the direction and content of their sounds and speech, and what and how they are touching. Display devices that image directly onto the retina, high-resolution miniature LCDs, spatial sound generators. Some of these innovative transducers operate not just non-invasively but invisibly. "No one should ever have to see a computer. The complexity should be soaked into the world around you," says Lucente.

While humans are adept at sensory integration and data fusion, computers are far less so. It is clear (and probably has been since GLOWFLOW in 1968) that multimodal interaction is a seminal goal, and that achieving it is a formidable challenge. Now that computational power seems be catching up with algorithmic understanding, the panelists can report and discuss exciting progress in this area.

Our panelists have decades of experience in interface design. Their perspectives are theoretical and pragmatic, incremental and radical; their work is elegantly inspiring and often delightfully unconventional. All were considered visionaries, but now their visions are achievable, and even industry is paying attention. They are seasoned practitioners with their own viewpoints. All are articulate, and none is shy; the Q&A and discussion periods promise to be stimulating.

Interfaces to newborn technology are usually "close to the machine." Early automobiles had spark-advance levers, mixture adjustments, hand throttles, choke controls. As automobiles have evolved, their affordances have moved "closer to the user:" speed, stop, reverse. We're tracking a similar evolution in human-computer interaction (HCI) space. Perhaps interfaces are finally growing up?

A video game can be controlled by real-time gesture sensing and recognition. (William T. Freeman)

User Domains and Appropriate Design

Bill Buxton
Alias | Wavefront Inc. and Silicon Graphics, Inc.
buxton@aw.sgi.com
www.dgp.toronto.edu/people/BillBuxton/billbuxton.html

When people are asked to "draw a computer," about 80 percent of the time they draw the I/O devices. This says two key things:

1. The most powerful force in shaping people's mental model of the nature of the beast is that which they see, feel, and hear.

2. This same shaping influence is an accident of history, and hence a candidate for change.

So, in designing a system, I can design its physical manifestation in such a way that its affordances conjure up the mental model that I am trying to encourage. And the better I understand the application domain, the skills of the intended user, and the context (physical and social) where the system is to be used, the more appropriate the mental model I can develop. And, consequently, the more appropriate the affordances and design of the system. But this flies directly in the face of how we design systems today. Today we follow the "Cuisinart," "superappliance" approach to design, which more or less says that the same basic type of box suits all types of users. But different users and tasks may well require (often radically) different approaches to what constitutes "a computer," because the two key components - the display and the input transducer(s) - are affected. I attempt to show how this particular approach to design affects the design of computer graphics systems for animators and industrial designers, as reflected in some of our research at Alias | Wavefront and SGI, in the form of live demonstrations and video examples.

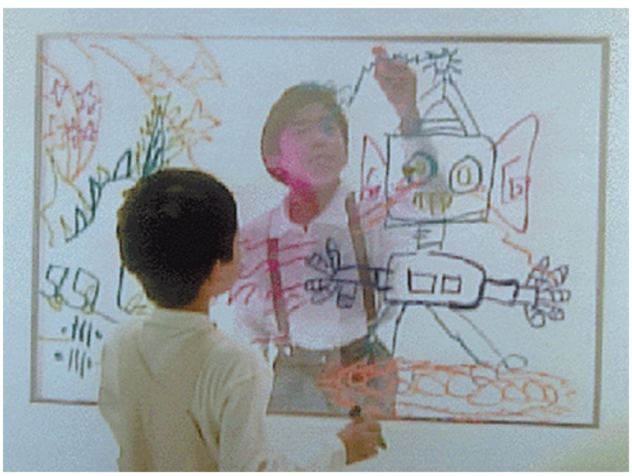

Computers Looking at People

William T. Freeman
MERL, a Mitsubishi Electric Research Laboratory
freeman@merl.com
www.merl.com/people/freeman/index.html

Computers can be used to interpret users' movements, gestures, and glances. Fundamental visual measurements include tracking, shape recognition, and motion analysis. For interactive graphics applications, these algorithms need to be robust and fast, and they need to run on inexpensive hardware. Fortunately, interactive applications can make the computer-vision problem easier. They can constrain the possible visual interpretations and provide visual feedback to let users adjust their inputs. I present several vision algorithms for interactive graphics and various vision-controlled graphics applications that use them: vision-based computer games, a hand-signal recognition system, and a television set controlled by hand gestures. Some of these applications can employ a special artificial retina chip for image detection or pre-processing.

Tangible Media

Hiroshi Ishii
Massachusetts Institute of Technology
ishii@media.mit.edu
ishii.www.media.mit.edu/people/ishii/

Eyes are in charge, but hands are under-employed. We live between two realms: our physical environment and cyberspace. Despite our dual citizenship, the absence of seamless couplings between these parallel existences leaves a great divide between the worlds of bits and atoms. "Tangible Bits" explores seamless interfaces among people, digital information, and everyday physical environments to go beyond eye-centric graphical user interfaces. We are designing "tangible user interfaces" that employ augmented physical objects, instruments, surfaces, and spaces as media to bridge virtual and physical worlds. We are making bits physically accessible through graspable objects as well as through ambient media in an augmented space. These interfaces emphasize both visually-intensive "hands-on" foreground interactions and background perception at the periphery of human senses through ambient light, sound, airflow, and water movement.

Clearboard-2 seamlessly integrates groupware and videoconferencing technologies. (Hiroshi Ishii)

Interfaces for Humans: Natural Interaction, Tangible Data, and Beyond

Natural Interaction

Mark Lucente
IBM Research
lucente@alum.mit.edu
www.research.ibm.com/natural/

Humans communicate using speech, gesture, and body motion, yet today's computers do not use this valuable information. Instead, computers force users to sit at a typewriter keyboard, stare at a TV-like display, and learn an endless set of arcane commands, all of which often leads to frustration, inefficiencies and disuse. We have created DreamSpace, a system that enables natural interaction through an intuitive yet richly interactive interface that "hears" users' voice commands and "sees" their gestures and body positions. Interactions are natural, more like human-to-human interactions. This information system understands the user, and - just as important - other users understand. Users are free to focus on virtual objects and information and understanding and thinking, with minimal constraints and distractions by the computer, which is present only as wall-sized 3D images and sounds (no keyboard, mouse, wires, wands, etc.) The multimodal input interface combines voice (IBM ViaVoice speech recognition), body tracking (machine-vision image processing), and understanding (context, and small amounts of learning). DreamSpace is essentially a smart room that employs a deviceless natural multimodal interface built on these emerging technologies and combined with ever-cheaper computing power. Future natural interfaces will allow information and communication anywhere, anytime, any way the user wants it - in the office, home, car, kitchen, design studio, school, and amusement park.

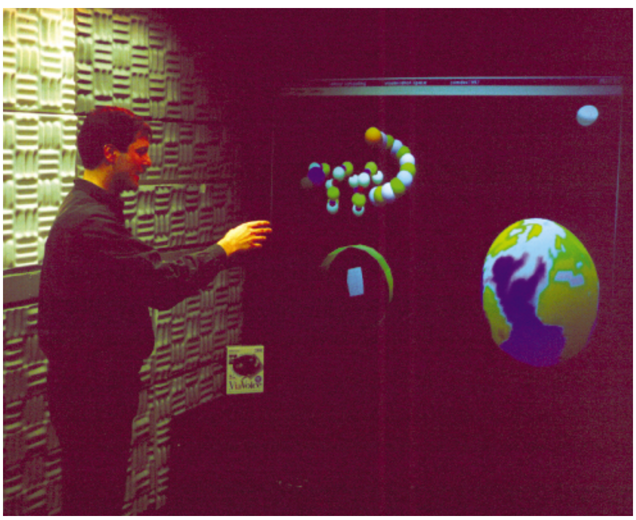

DreamSpace optimizes ease of use, enjoyment, and the organization and understanding of information. (Mark Lucente)

HCI Through Creative Plagiarization

Michael J. Sinclair
Microsoft Research
sinclair@microsoft.com
www.oip.gatech.edu/imtc/html/mike.html

To find creative HCI solutions that are both usable and affordable for a potential volume market, we can look at existing well-executed engineering efforts and discover a multitude of existing embodiments waiting to be re-purposed. From a steady diet of engagements with sponsors looking for low-cost commercializable outcomes, we have learned to investigate and exploit appliances from unrelated fields. As a result, we can demonstrate high-fidelity, low-cost solutions for 3D tactile feedback, digital panoramic photography, medical instrumentation, and low-cost 3D digitization. Significant gains can be realized through creative re-purposing. So: "Don't shade your eyes - plagiarize."

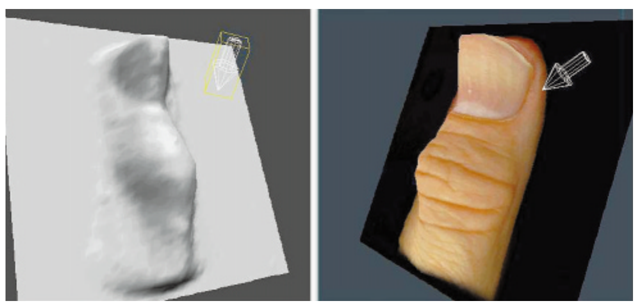

The Haptic Lens acquires 3D surfaces in real time, while applying a known force on the object or body part. (Michael J. Sinclair)

Are You There? Presence in Virtual Reality: What Is It About and Why Care?

Moderator
Mel Slater
University College London

Panelists
Nat Durlach
Massachusetts Institute of Technology

Lawrence J. Hettinger
Logicon Technical Services, Inc.

Randy Pausch
Carnegie Mellon University

Dennis R. Proffitt
University of Virginia

Mel Slater
University College London
m.slater@cs.ucl.ac.uk
www.cs.ucl.ac.uk/staff/m.slate

Virtual reality is supposed to be able to provide a strong illusion of being and acting in a simulated "other place." Is this realizable? What is required to enable someone to be "present" in the scenario created by an application? Can this degree of presence be measured, if it exists at all? How does it relate to performance? Does any of this matter anyway to application builders or users?

Immersion in a computer-generated (virtual) environment (VE) is a matter of fact. To what extent can the energy (light, sound, touch, force) comprising the environment surround the participants? To what extent can they act within it? For any given system, whether desktop or CAVE, these questions can be definitely answered. However, what are the psychological and behavioral consequences of these different states of immersion? Given any application running on a system, do participants come to experience the scenario as a "place" and act appropriately within it? The word "presence" has been used to describe this reaction, thus granting a virtual environment sufficient place-like characteristics that people can talk about "being there" and doing things "there." Presence can be thought of as an emergent property of a virtual environment system, indeed of any system, ranging from theater and movies to computer-generated displays, that portray alternate realities embedded within our everyday reality.

There are a number of important issues for discussion that follow from this:

- If there is such a phenomenon as presence? What use is it to application builders and users?
- Is there a scientifically acceptable, practical way to measure presence, independent of application and system characteristics, that can be compared across different people?
- What are its determinants? How important is realism in the sensory displays (visual, auditory, haptic) or conformity with physical laws such as gravity and collision response?
- Independently of presence, what are the benefits of immersion? What types of applications require immersion?

The panel considers these questions in the context of talks about presence and task performance, whether presence is a factor to be taken into account in the design of effective VE applications, the circumstances under which people will behave as if they were present, and the conditions under which natural perceptions of simulated scenes are attained.

Presence and Task Performance

Nat Durlach
Massachusetts Institute of Technology
durlach@cbgrle.mit.edu
rleweb.mit.edu/P-DURL.HTM

The concept of "presence" in a synthetic world (as well as the less-discussed concept of "absence" in the real world) is interesting in its own right. However, it is yet to be demonstrated that the concept of presence is useful in the sense of having a well-defined relationship to task performance. In other words, beyond the need to arrive at a satisfactory definition and satisfactory measurement technique for presence, for the concept of presence to be practically useful, it must be demonstrated that it can function as a useful intermediate variable between the parameters that objectively specify the task situation and the parameters that objectively describe performance in this situation.

Roughly speaking, presence can be regarded as a useful intermediate variable provided:

1 Knowledge about the value of the variable presence realized in a given situation increases one's ability to predict the value of task performance in that situation.

2 The gain in the predictability of performance achieved by determining presence is large relative to the effort required to determine presence.

3 This gain exceeds the gain that would be achieved by directing equivalent effort elsewhere.

In view of all the difficulties in defining and measuring presence, and the impossibility of showing that performance is a monotonic increasing function of presence (in fact, it can be shown that under some circumstances, an increase in presence leads to a decrease in performance), the likelihood that presence can play the role of a useful intermediate variable appears rather low.

Designing Successful Virtual Environment Applications

Larry Hettinger
Logicon Technical Services, Inc.
lhettinger@falcon.al.wpafb.af.mil

What do we want virtual environment systems to be able to do with respect to the human beings who use them? Do we want users to be able to perform difficult tasks, such as neurosurgery or flying airplanes, better? Acquire cognitive and perceptual-motor skills more thoroughly and rapidly? Enjoy themselves and forget, for the time being, their problems and difficulties in the "real" world? What do we want these systems to do? And do we need "presence" to do these things? If so, how do we know we need it, or do we simply assume that we do?

These are extremely critical issues for the future success and proliferation of this technology. I present aspects of a user-centered approach to designing VE systems that we have employed in our work with the U.S. Air Force and various medical organizations to illustrate the following points:

- At least in the applications that we have been pursuing, effective VE system design has been the direct result of mapping human performance requirements onto system characteristics (i.e., what information does the user need? How can we best make it available?).

- Presence, in and of itself, is and should not be a prime consideration in system design. Achieving the "behavioral goals" of the particular VE application should be. Sometimes this may result in a need for a configuration that produces a sense of presence in some users, but that is (and ought to be) strictly secondary.

- Presence is a higher-level characteristic of VE system design that results from various combinations of lower-level engineering characteristics. It can be useful in accomplishing the behavioral goals of the system, it can be irrelevant to the accomplishment of those goals, or it may interfere with the accomplishment of those goals.

Getting People to Behave as if they are Present

Randy Pausch
Carnegie Mellon University
pausch@cs.cmu.edu
www.cs.cmu.edu/~pausch

I have worked on "presence" from two ends of the spectrum: I have conducted scientific studies, with Dennis Proffitt of the University of Virginia, where we have measured task performance in immersive versus non-immersive displays, attempting to discern what aspects of these displays produces a better spatial awareness and spatial recall of the environment. In addition, we have looked at people's ability to perform tasks such as measuring the slant of a hill, and found that people's impressions in immersive (which we define as "head-tracked") displays mirrors real-world performance, but performance using displays does not. In this work, we do not care if people think they are present. We care if they behave as if they are present.

At the other end of the spectrum, my work with Imagineering/DisneyQuest has allowed me to see thousands of "naïve users" experience high-quality, first-time immersive experiences in a variety of technologies. From this, I have developed a set of design principles regarding what does and does not elicit a sense of "being there." These principles focus surprisingly on what one might think are tangential issues: the "pre-show," the guests' "goals" during the experience, and their interaction with other guests during the experience.

Evoking Natural Perceptions of Simulated Scenes

Dennis Proffitt
University of Virginia
drp@virginia.edu
www.people.virginia.edu/~drp/

The desktop terminal is not a window on the world. Contrary to what one might expect, many of our everyday perceptions are distorted, and these distortions are of a much greater magnitude when viewing occurs in the world or in a head-tracked, head-mounted display (HMD). For example, being immersed in a real or virtual environment causes hills to look steeper and buildings to look taller than they actually are. These perceptual biases are considerably smaller when the same scenes are viewed on a desktop terminal. One of the reasons for these differences is that, when immersed in an environment, the size of everything in the scene can be scaled to the size of one's body. When viewing a scene presented within a HMD, people spontaneously scale the size of unfamiliar objects to the altitude of their eyes. When viewing a desktop display, this body-scaling of size is impossible, since the altitude of the point of projection need not coincide with one's viewpoint, and thus, is unknown. We have found that when people feel that they are immersed in a scene, as assessed by a rating-scale measure of presence, they use eye-height scaling effectively. When they rate their sense of presence as low, they do not do so.

Another set of experiments showed that people overestimate vertical extents far more when immersed than when viewing a desktop display. In one condition, people viewed a "virtual desktop display" within a scene presented in an HMD. When observers moved close to the "virtual desktop," so that the image completely filled their field of view, the image that was projected into their eyes was exactly the same as the projection that was formed when they were immersed in the scene. The vertical extent bias was smaller in the "virtual desktop" condition than in the immersion one even though the pixels projected into the eyes were identical. What matters here is whether or not the person perceives the scene to be a small projection.

My talk compares immersive and non-immersive environments with respect to their ability to evoke natural perceptions of the simulated scene. Natural perceptions are not necessarily accurate, but are often biased in a manner that promotes everyday actions.

The Sorcerer's Apprentice: Ubiquitous Computing and Graphics

Organizer
W. Bradford Paley
Digital Image Design Incorporated

Panelists
Bill Buxton
Alias | Wavefront Inc., Silicon Graphics, Inc.

Steve Feiner
Columbia University

Hiroshi Ishii
Massachusetts Institute of Technology

S. Joy Mountford
Interval Research Corporation

Steve Shafer
Microsoft Research

Mark Weiser
Xerox PARC

W. Bradford Paley
Digital Image Design Incorporated
brad@didi.com

This panel will address Ubiquitous Computing and related research in an unusual, but appropriate way: many short presentations of small objects and interaction techniques. Six wildly creative leaders pursuing innovative interaction techniques will present very brief talks centering on one or more objects. Demonstrations may be given, with an emphasis on computer graphics.

These objects will help define the territory of how computers might be used when they are as plentiful and inexpensive as credit cards or keychains. The audience will then be encouraged to ask questions and suggest applications, potentially developing the panel into a 2000-person brainstorming session. Bring your favorite problem; take notes and take away the future.

Bill Buxton
Alias | Wavefront Inc., Silicon Graphics, Inc.
buxton@aw.sgi.com
www.dgp.toronto.edu/people/BillBuxton/billbuxton.html

One of the mantras of ubicomp is that computation be both transparent and yet everywhere. Of course, one can't take that literally. What it really means is that services should be available at the right time at the right location in the right form for the intended user and purpose.

This also implies a move from the one-size-fits-all approach to computation that dominates design today, an approach that matches precisely the design philosophy of the Quisinart. With appropriate design, we will soon see as few general purpose computers as we do Quisinarts.

The key here is where, when, and how computation is exercised, with the attendant divergence (as opposed to convergence) of computation. Thus, when walking through an animation studio, we will be able to determine if we are in the character animation department or the accounting department simply by looking at the tools being used (something not possible today, given that the computers are almost interchangeable from a design perspective.) Likewise, we will see a move to embedded systems, where the computation is integrated seamlessly into what appear to be common devices.

Thus, in cinematography, for example, we are about to see a return to "in-camera" effects, since the camera itself will simply be a computer cleverly disguised as a camera so that its operators (the cinematographers) can operate it using their existing skills. The result will be "what you see is what you get" cinematography, with CG and compositing done "in camera" along with the live action.

Likewise, digital still cameras will actually be graphics computers with integrated paint and compositing capability so that on a location scout, for example, you will be able to capture what you want how you want it. You will not have to take the images back to a host computer and load Photoshop in order to get your work done.

And finally, we come back to the lowly pencil and the oft-repeated request that I hear, namely: "Why can't a computer be as easy to use as a pencil?" Well, ask someone who has graduated from art college and has spent 13 years learning how to draw: "How easy is a pencil, really, to use?" Then duck. To imply that a pencil is easy is an insult to the skill of the artist. Yet, that is what nearly every graphics program does, by virtually ignoring this hard-won skill. And, more often than not, when one is allowed to draw, this is accomplished using a mouse, which is more like a bar of soap (and we all know that the only time to draw with soap is Halloween - not on a $50k workstation!).

It is arguable that despite the interval since Dick Shoup built the first paint program in the early 70s, nobody who has ever used an airbrush in a paint program has ever seen an actual airbrush, much less used one. This claim is based on the fact that nobody uses an airbrush in the physical world without a mask or frisket in the other hand. Yet, other than StudioPaint, no paint program supports this capability, and as a consequence, they thumb their proverbial noses at the artists they are supposed to be helping. In ubicomp, respect for such skills is at the forefront of design. The resulting system will be useless for accountants, but great for the artist. One size will not fit all.

So, computers might actually be designed to capitalize on artists' ability to draw. And traditional animators will be able to pose characters, for example, using their hard-won traditional skills, rather than having to learn a whole new way to articulate what they are already fluent in doing.

There has been no significant progress in how we interact with computers since 1982. It is about time something changed.

Steve Feiner
Columbia University
feiner@cs.columbia.edu
www.cs.columbia.edu/~feiner

Ubiquitous computing is often presented as the antithesis of virtual reality: embedding many computers in the world, as opposed to embedding the world in the computer. We believe, however, that these are complementary concepts rather than competing ones. This may be especially true of that form of virtual reality known as "augmented reality," which uses see-through displays to overlay a virtual world on the real world. As the technologies needed for see-through displays and high-performance computers get better, lighter, and cheaper, they can be used to create ubiquitous wearable systems. I describe what these systems might look like and how they might be used. One possibility is that of developing "hybrid user interfaces," in which multiple computers and displays, ranging from hand-held to head-worn to wall-sized, are used together to take advantage of the benefits of each.

Hiroshi Ishii
Massachusetts Institute of Technology
ishii@media.mit.edu
www.media.mit.edu/people/ishii

Tangible Bits is the new paradigm of HCI (human-computer interactions) for the era of ubiquitous computing. The goal of Tangible Bits is to change the "painted bits" of GUIs (graphical user interfaces) to "tangible bits," taking advantage of the richness of multimodal human senses and skills developed through our lifetimes of interaction with the physical world. Tangible Bits explores seamless interfaces between people, digital information, and everyday physical environments. Based on the vision of Tangible Bits, we are designing "tangible user interfaces," which employ augmented physical objects, instruments, surfaces, and spaces as media to bridge virtual and physical worlds. We are making bits physically accessible through graspable objects as well as through ambient media in an augmented space. These interfaces emphasize both visually intensive, "hands-on" foreground interactions and background perception at the periphery of human senses through ambient light, sound, airflow, and water movement.

S. Joy Mountford
Interval Research Corporation
mountfor@interval.com

Having computer appliances in every place or situation that I might want to access information may indeed be helpful and convenient. However, I want to make sure that such things are not invasive, either in a design or personal sense. I must maintain control over when such appliances choose to communicate with me, and also the manner in which they display their content. If they are to be truly ubiquitous for personal use, they must also be designed to be attractive and personal for individual use.

The main issue for computer graphics seems to be the potential lack of need for sophisticated 3D graphics to accompany small portable pieces of computers. High-resolution miniature graphic display technology will be critical to ubiquitous and portable uses, which will also drive acceptance by the design community for personalized uses. Current wearable computers are a far cry from acceptance for the fashionable uses of pervasive technology, so the design space is open for much graphics innovation.

The Sorcerer's Apprentice: Ubiquitous Computing and Graphics

Steve Shafer
Microsoft Research
stevensh@microsoft.com
research.microsoft.com/research/vision/stevensh

People are beginning to look beyond the now-classical "desktop PC" model of computing, and one of the things they see is computing that is built into the environment as a part of the building. At Microsoft Research, the form this is taking is the EasyLiving project, which uses cameras as the primary input device and includes an explicit geometric model of the environment, people's location and facing, and the location and "usability field" for key devices in the room (such as displays, cameras, and microphone/speakers). Such environments and the applications that will run in them introduce several problems for computer graphics:

- How to display to the user a view of what the system knows and does not know.
- How to handle multiple video streams in dynamic pathways within a single (possibly low-bandwidth) network.
- How to map a user interface onto graphical displays with form factors that vary by an order of magnitude or more.
- How to provide effective visual signalling for a wide variety of events that the system may want to tell the user about.
- What should be the graphical elements of a speech/vision interface such as the multimodal equivalent of a "dialog box"?
- How to predict what a camera should or should not be able to see in a given situation.
- How to let the user probe the state of the system and the world meaningfully.

We don't have any answers to these questions at Microsoft, but by laying out these issues as we see them, we hope the audience will be inspired to start coming up with clever ideas for future publications that we can read and learn from.

Behavioral Modeling and Animation: Past, Present, and Future

Organizers
Demetri Terzopoulos
University of Toronto / Intel Corporation

Xiaoyuan Tu
Intel Corporation

Panelists
Kiran Joshi
Walt Disney Feature Animation

Ken Perlin
New York University

Craig Reynolds
DreamWorks Animation

Toby Simpson
Cyberlife Technologies Ltd.

Would computer animators rather be graphical model puppeteers who keyframe the detailed actions of their characters, or would they prefer to direct intelligent, self-animating virtual actors? On the one hand, the animator has complete control over all aspects of the character's low-level motions. On the other hand, control is relinquished to gain greater convenience in the higher-level specification of a character's behavior.

Behavioral modeling was introduced about a decade ago in Reynold's "Boids" model, as a means of producing animated scenes containing many more characters than could practically be animated by hand. The behavioral modeling approach has today expanded to include sophisticated functional modeling of animals and humans, resulting in realistic, self-animating graphical characters.

This panel discusses the fundamentals of behavioral modeling and animation arising from knowledge of living systems and their environments. Artificial life models have evolved a long way from the comparatively primitive geometric models of traditional computer graphics. The panel reviews the state of the art and debates the promises and limitations of behavioral modeling and animation from multiple perspectives, including production animation, the interactive games industry, and the research community.

The Design of Characters with Complex Behavior

Craig Reynolds
DreamWorks Animation
cwr@red.com
hmt.com / cwr /

Behavioral control allows animated scenes to contain more characters than would be practical otherwise. The most exciting aspect of behaviorally driven animation, however, is the way these multi-agent systems form an environment in which complex global behavior can emerge from the interaction of relatively simple local rules. A well-tuned behavioral simulation amplifies an animator's effort. When everything goes well, the result is an engaging and visually rich scene full of unexpected details of motion. Poised on the boundary between chaotic dynamics and rigid control, the most enjoyable behavioral simulations operate in the life-like regime Langton called "the edge of chaos."

The crux of behavioral design is the art of tuning the dozens of parameters in a typical behavioral model. I advocate a toolkit approach to building autonomous characters: Starting with a library of simple general-purpose, reusable behavior modules, a character requires only some custom control structure to switch or blend between behavioral modules.

While crowd scenes for animated films are a significant application of behavioral animation, a more compelling argument can be made for its importance in interactive applications. Behavioral characters are reactive agents, and so are uniquely suited to provide believable interaction between human users and autonomous characters. A behavioral character designed to react to others of its kind can just as well react to the avatar of a human participant.

Behavioral Animation in Disney Feature Films

Kiran Joshi
Walt Disney Feature Animation
kiran@fa.disney.com
www.disney.com /

Over the years, Disney has evolved from traditional hand-drawn crowd scenes where only a few characters are animated to scenes of epic scale involving thousands of animated characters. From the herding system developed for "The Lion King" to the crowd animation packages used for "The Hunchback of Notre Dame," "Mulan," and "Dinosaur" features, we have refined the process of crowd animation. In a production environment such as ours, it is absolutely crucial that an artist, at all times, have absolute control over the visual outcome of a shot. The issue I address is how to gain control over the result of a procedural animation, i.e., the crowd.

While physics, dynamics and artificial intelligence may carry you 90 percent of the way, we need to achieve that final 10 percent. We therefore implemented a hybrid system, where a simulation can be post-edited to achieve a better-looking result. The system provides the means for both macro and micro control. In general, the simulator is used to obtain results as close as possible, which are then fine tuned in an editor. At the macro control level (i.e. the simulation dynamics), we can often provide more explicit control through "image maps," which map from pixels to state parameters. These maps can simply be drawn by an artist, and they provide an input parameter-set to the simulation that would be hard or impossible to achieve only programatically.

Afterwards, at the micro level, an editor can change virtually any parameter of any entity, pertaining to position and velocity, appearance, and behavior timing, thereby providing a mechanism to stage an entity against any visual requirements.

I show how we go from a layout drawing to the final animation and give the artist the control to achieve the final look.

Wildebeest stempede in "The Lion King."

Crowd in "The Hunchback of Notre Dame."

AI Modeling for Behavioral Animation

Xiaoyuan Tu
Intel Corporation
xiaoyuan_tu@intel.com
www.cs.toronto.edu/~tu/

The distinguishing feature of behavioral animation is that each animated character is governed by a model of how it should behave. Although the model can be as simple as a few behavioral rules, the interaction between the characters can generate elaborate emergent behavior. The "Flocking Boids" is a landmark example.

A good topic for discussion is the future or extension of behavioral animation models. On the one hand, it is interesting to investigate the realm of emergent behavior from the complex interactions of simpler behavioral entities. On the other hand, a natural extension to current models of reactive behavior is the modeling of cognition. I consider it the ultimate challenge to animation modeling that we may someday model a fully functional human. Imagine how differently an animated feature would be produced when the characters can react and reason like real human actors. The animator's role then will be like that of a director, and the virtual characters will improvise their parts based on the direction they receive.

To this end, the topic of artificial intelligence naturally enters the domain of graphics modeling. We are still a (very) long way from achieving this goal. However, this should not intimidate us from making initial steps, nor should this invalidate our early attempts. I advocate exploration of existing AI techniques and ongoing AI research for cognitive modeling in animation. The common goal of modeling human intelligence shared by AI and graphics researchers will surely prove beneficial to both areas.

Artificial Fishes in a Digital Sea

Artificial Life in Home Entertainment

Toby Simpson
Cyberlife Technologies Ltd.
toby.simpson@cyberlife.co.uk
www.creatures.co.uk/

A critical part of computer gaming in the future will be construction of believable artificial agents and rich, diverse, and self-consistent environments in which they can live – an application for which artificial life techniques are well suited. Artificial life is likely to be a key technology of the future, and many aspects of it are already finding their way into home entertainment in titles such as "Creatures."

"Creatures" allows users to interact with artificial autonomous agents whose behavior is controlled by genetically specified neural networks and biochemistry, and is currently the only commercial entertainment product to provide this. We believe that the success of "Creatures" demonstrates the value of such technologies in entertainment and the strength of the relationships that users are able to form with such agents.

We expect that by pursuing the process of using computers to model biological systems that can in themselves be intelligent, rather than attempting to make a computer intelligent, we will be able to achieve human-level intelligence in a machine by the year 2020. We believe that "Creatures," and now "Creatures 2," represent substantial steps in this direction – plausible artificial organisms whose behavior is emergent rather than programmed – living in rich, detailed eco-systems. It is likely that this approach will yield virtual realities that are so real that it may not be possible to tell the difference any longer.

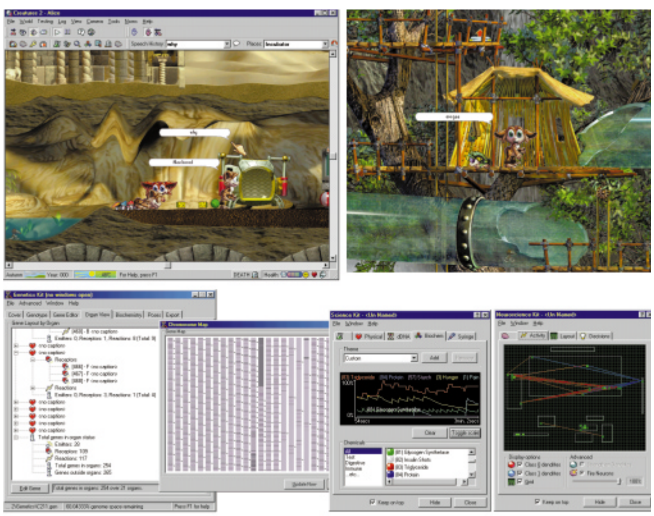

Creatures 2

What are the Limits of Behavioral Modeling and Animation?

Ken Perlin
New York University
perlin@nyu.edu
www.mrl.nyu.edu/perlin/

Animators freely tap into many (and often unexamined) intuitions and judgments in order to create their work. Even the most sophisticated behavioral modeling techniques cannot completely replace culturally and psychologically informed authoring techniques that talented animators employ to create linear animation (for example, why did a character raise his eyebrow and hunch his shoulders in just that particular way at that moment?).

Explicitly defined behavioral models will never be able to completely replicate such intuitions and judgments. Such behavioral models will always need to be integrated and leveraged with contributions from more traditional approaches that simply give animators a flexible tool with which to "sculpt" their intuitively based judgements.

So how do we blend behavioral modeling with the sort of hand-tuned work that animators and other skilled craftspeople are so good at? How do we do this in an interactive setting, when the animator is no longer present to modify a character's response to an evolving story? I think our most important challenge is to work out good ways to integrate behavioral and animation-compositing methods. This challenge is the focus of our Improv project at NYU.

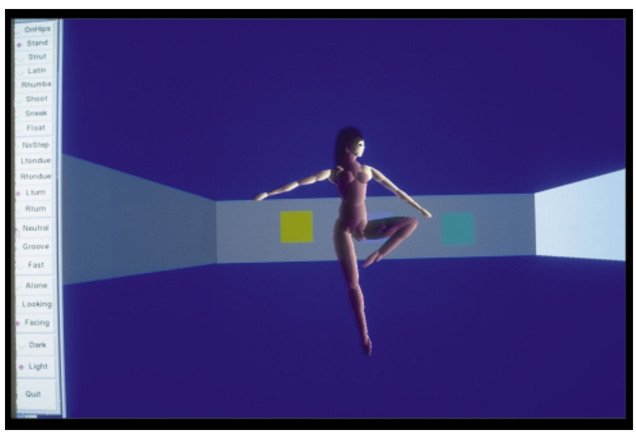

"Danse Interactif" by Ken Perlin, SIGGRAPH 94 Electronic Theater

Location-Based Entertainment: The Next Generation

Moderator
Randy Pausch
Carnegie Mellon University

Panelists
Trevor Bryant
Sony Development

Joe Garlington
Walt Disney Imagineering

Jon Snoddy
GameWorks

Jordan Weisman
FASA Interactive

Recently, location-based entertainment (LBE) has moved from speculation to reality. Following in the tracks of Virtual World Entertainment (formerly BattleTech Center) and Dave & Buster's, GameWorks and DisneyQuest are now open to the public. This panel discusses the factors that go into creating an LBE attraction, based on the industry leaders' real-world experiences.

We focus not on the business model, other than to the extent that it affects design and its effect on issues like throughput and demographics. Instead, our focus is on the process and content that goes into a creating a guest experience in an LBE venue. We'll also discuss what works and what doesn't! An underlying theme is that LBE attractions represent a new medium, and no one really knows how to author successful attractions yet.

The Meaning and Context of "Location"

Trevor Bryant
Sony Development
trevor@sre.sony.com

The key to defining and creating location-based entertainment attractions is understanding the meaning and context of the word "location." For me, the word is best characterized by the concept of community. Community as a descriptor of location then puts a requirement on the designer and the developer to be responsible, communicative, and sensitive to issues of neighborhood. In retail and entertainment industry jargon, location-based entertainment often means simply erecting a shopping mall facility, adding some entertainment, and dressing up conventional strategies with hype and marketing to be re-baptized as an LBE. This approach falls far short of providing an urban neighborhood with a sense of ownership and pride in a commercial but highly visible public-use facility. This sense of ownership is key to the maintenance and long term success of an LBE.

In turn, any attraction within that LBE should also reflect a sense of the community where it resides. What are the entertainment interests of the community? What are the cultural interests? Do the demographics require local repeat visitation or are they tourist-based? All of the questions regarding demographics, cultural interests, and community, I believe, are prerequisites for committing to an attraction concept. In creating Metreon, a Sony entertainment center, the attraction design team has placed the emphasis on understanding the interests of the neighborhood. Of course, innovation, technology, and hardware are clear prerequisites, as authenticated by Sony's endorsement. However, the accent is on entertainment concepts that reflect the community rather than on technology or hardware. I believe this position is helping Sony Development to redefine concepts of location-based entertainment and provide attractions that represent both corporation and community.

Shared Activity, Storytelling, and Creative Process

Joe Garlington
Disney Imagineering
Joseph_Garlington@cc.wdi.disney.com

Our LBE needs to attract our core audience: families with kids eight and above. To do so requires an understanding of why people come to our facilities and a process that develops products that will please them when they do. People come to enjoy themselves. They come in groups. They come to have fun together, to share an experience that draws them closer to each other. Our products entertain people singly, but more importantly they are the McGuffin that helps people enjoy the company of each other. Our games are either true group games or single interactor activities that others are encouraged to share. We want people to laugh and talk and bump and touch and never lose track of the other people they came with, no matter how absorbing the activity is.

Seven out of 10 movies fail. Yet movies have an hour or more to tell their story, are based upon centuries of story-telling experience, use a visual language shared by filmmakers and filmgoers, and require little of viewers other than to suspend disbelief and identify with the hero or heroine. In our venues, we have three to five minutes to tell our "story." We have no gaming tradition that works, since, though games have been around forever, the commonalities between them are small, (for example, find the common elements in chess, baseball, bridge, jax, jigsaw puzzles, cat's cradles, Tetris, etc.). There is no common language. And a guest coming to the attraction hasn't a clue about his or her role: Is this first-person or third? What's the interface? What are the rules? What strategy and tactics are needed?

My belief is that nobody is smart enough to guess how to set all the variables in an interactive attraction. The Hollywood "director-as-dictator-with-a-vision" paradigm doesn't work. Only by letting guests lead us to the solutions can we ensure that we provide games that will please them. In our process, we use designers to identify basic, fun things to do. If a baby does it or animals do it, it's probably a good place to start. We identify a potentially novel way of getting people to engage in this basic behavior, often by enabling the behavior with unexpected technology. Then we build the cheapest, quickest, lowest-tech mockup we can, and we playtest it. We listen to what our playtesters tell us, modify the mockup, and do it again, increasing the technology if required. We iterate the process clear through production (often to the consternation of our vendors) to ensure we don't get off track. No process is perfect, but we believe this process will lead to much higher success ratios than are normally found in entertainment development of this kind.

When Intuition Doesn't Work: The Need for Iterative Testing

Randy Pausch
Carnegie Mellon University
pausch@cs.cmu.edu

I designed and implemented the "playtesting" for the DisneyQuest attractions, which were all extensively tested and iteratively redesigned. Most LBE ride designers come from a background in film, theme parks, and other passive media. My background is in the design of interactive devices: computer GUIs, VCRs, bank machines, etc. Most LBE attractions are interactive, and the rules of storytelling and design are not yet understood for this medium. Surprisingly, the set of techniques used to design good VCRs are very effective when applied to interactive LBE rides. In particular, the only way to find out what works is to mockup the attraction quickly, test it, and make changes based on the test data, not the designer's intuition. Knowing how to extract what is in guests' heads during the experience is part science, part art, and part sociology. I discuss the techniques we used to find what worked and didn't, and the emerging patterns regarding storytelling, how to direct guests' attention, and how to get guests to interact with synthetic characters in virtual environments.

A Market Emerges: Selecting and Developing "Hit" LBE Attractions

Jon Snoddy
GameWorks
jon@gameworks.com

Though it is hard to believe, a viable LBE marketplace is finally emerging. A number of companies, both large and small, seem to have found combinations that actually make business sense. This is good news for the attraction developers who were lucky enough to survive the past few years of famine. As this new channel grows, it will have a giant appetite for product.

It has taken a surprising amount of time to figure out what kinds of experiences the public is willing to pay for. Early attempts were either too expensive or lacked enough excitement to generate the necessary repeat business. The successful experience today must deliver an exciting and unique experience to the player, as well as the audience, in order to work for the LBE center. It must deliver that experience at a development cost that allows for the relatively short life of the attraction. Though producing a hit no longer seems impossible, it is still pretty complicated. I discuss the product selection and development process in more detail.

Lessons Learned from Three Generations of "BattleTech"

Jordan Weisman
FASA Interactive
jweisman@fasainteractive.com

I have always been fascinated with the symbiotic relationship between interactive games and dynamic social environments. Virtual World Entertainment (originally BattleTech Center) established that interactive media could be integrated into a traditional three-act format and experimented with focusing the social dynamics to increase the value of the experience.

Growing from a single site to a chain of 28 sites worldwide over the next five years, our company developed three new versions of our LBE "cockpits" and experimented with various aspects of pre-experience, post-experience, and food/beverage and retail tie-ins. Having looked at both aspects, we are currently focusing on entertainment rather than location; we currently OEM our "experience content" to Dave & Buster's, for example. In my talk, I discuss many of the lessons learned in our decade of experience building technology-assisted person-to-person experiences.

Sublime and Impossible Bodies

Moderator
Sara Diamond
Banff Centre for the Arts

Panelists
Jane Prophet
Slade School of Art

Joshua Portway
Realworld Studios

Catherine Richards
The University of Ottawa

Douglas MacLeod
WurcNet Inc.

Arlindo Ribeiro Machado Neto
Catholic University of São Paulo

Critical Art Ensemble (Steve Kurtz)

Ahasiw Maskegon-Iskwew
SOIL Digital Media Production Suite

Sara Diamond
Banff Centre for the Arts
sara_diamond@banffcentre.ab.ca

Bodies loom large when we configure the potential of digital media and cyberspace. Bodies extended, connected, sublime, erased, implanted, deformed, defaced, empowered, degendered, buffooned, grotesque are the stuff, the background story, and the metaphors of animation, cyberspace, net chat. The ability to transcend and engage with identities centered in human bodies as experiential, social, and biological "real worlds" has fueled much creative and intellectual fantasy and engagement within cyber culture. Bioengineering and artificial life have allowed further possibilities, anxieties, and impossibilities to emerge in most recent times. Flesh-eating diseases, viruses on screen and off, and millennial ecstasy renew expectations and fear about our bodies within digital and popular culture. The panel also considers the sublime, a resonant cultural paradigm of the modern West. Is the body and its nature as "exalted, awe-inspiring" made so through the interventions of human repression and ordering?

This panel features artists, architects, animators, out-of-the-box experts, and scientists. It draws on case studies, artists' works, small games, popular culture concerns and practice, as well as engineering and architecture to open up questions of identity, anxiety, and the physical. It includes anime, systems that link the physical and cyber worlds, data body configurations, and AL programming.

Sublime and Impossible Bodies draws from two recent symposia at The Banff Centre: Flesh Eating Technologies and Death, Desire, the Dream, and the Machine, disinterring themes of artificial life and death, the crazy organization of knowledge about the human body, cyborgs, and delirium, as well as the rational scientific pursuit of ideal and imagined bodies; the day-to-day use of the net for dialogue, seduction, and erotic play. These events follow on a number of conferences, publications, and musings, off- and online which place the body at play within cyber culture, such as Bioapparatus (The Banff Centre, 1991), Immersed in Technology (MIT Press and The Banff Centre, 1995), Body Matter (ISEA 1995 Panel), The Cyborg Manifesto (Routledge, 1996), Gender and Technology (1997, Wexner Center), and many others.

The panel approaches these issues primarily through provocative projects and ideas that the presenters have unleashed on the world. It follows the blood lines of a hoary, "old" debate, within the short life of cyber culture, yet hopefully opens new veins.

Jane Prophet
Slade School of Art
jane@cairn.demon.co.uk

My presentation offers an investigation of the cyborg body that is grounded in "the meat." Through a discussion of my new artist's CD-ROM, "The Imaginary Internal Organs of a Cyborg," I intend to briefly explore notions of the sublime, switching our sense of the sublime from relating to the large scale to the sublime as the very small, as exemplified by micropscopy and nanosurgery. Extracts from interviews with surgeons and medical researchers are used to provide anecdotal evidence of the psychodynamics of the operating theatre, or the "theatre of operations."

Joshua Portway
Realworld Studios
josh@realworld.on.net

I'm never very sure about my body on the net. Sometimes, when I'm playing Quake, it's very concentrated, disembodied, but focused. But mostly it's the kind of tenuous, dispersed, planktonic sense of presence you get when you're drifting through the Web, a pleasant, scary feeling of dissolution. I imagine it must be similar to the experience of a sensory deprivation tank. Meanwhile my wetware is feeling jealous and wants something a bit less ... ergonomic. I've written a program that crawls the Web, following a train of thought, distilling out the consensus of the net. I left it thinking about "sublime and impossible bodies" last night, and this is what it came up with this morning: "lonely computer taste nothing." Not bad, eh?

Catherine Richards
Artist/University of Ottawa
richards@uottawa.ca

In the panel, Richards argues for one foot in the body and another in cyberspace. If we lose our footing in the "real world," we will potentially lose psychic and social stability. She explores bodily issues using feminist tools of analysis. Richards chooses 19th century technologies, transforming these with digital means and media. She will engage with the actual physical effects of these technologies on human bodies, as well as emotional and physical intimacy in cyberspace.

Douglas MacLeod
WurcNet Inc.
wurc@acs.ucalgary.ca

MacLeod considers the architecture of the body in cyber- and virtual spaces, including the translation of these into the real. He addresses the ways that architects and designers imagine bodies and the ways that virtual imaging can instrumentalize human space. He also speaks to the ways that virtual imaging transforms actual architecture and the movement of bodies within it.

Arlindo Ribeiro Machado Neto
Catholic University of São Paulo
arlimach@uol.com.br

In the last few years, some artists have brought forward a cultural discussion of the possibility of surpassing the human through radical surgical intervention (Orlan), through the interface between flesh and electronics (Kac), or with robotic prostheses to complement and expand the potentiality of the biological body (Stelarc). An important landmark of this current took place on November 1997, at an art center in São Paulo, when the Brazilian artist Eduardo Kac implanted in his ankle an identification microchip with nine digits and registered himself with a databank in the United States via the Internet. Kac's emblematic event seems to suggest that a biological mutation may take place in the next future, when digital memories will be implanted in our bodies to complement or substitute our own memories. This reading is clearly authorized by the associations the artist makes between the implant and the public exhibition of his familial memories, his external memories materialized in the form of photographs of his grandmother's family, which was entirely annihilated in Poland during World War II. These images, which strangely contextualize the event, allude to deceased individuals whom the artist never had the chance to meet, but who were responsible for the "implantation" in his body of the genetic traces he has carried from childhood, and that he will carry until his death.

"Will we in the future still carry these traces with us irreversibly, or will we be able to replace them with artificial genetic traces or implanted memories? Will we still be black, white, mulatto, Indian, Brazilian, Polish, Jewish, female, male, or will we buy some of these traces at a shopping mall? In this case, will it make any sense to speak of family, race, nationality? Will we have a past, a history, an "identity" to be preserved?

Ahasiw Maskegon-Iskwew
SOIL Digital Media Production Suite
ng1soil@dlcwest.com

I explore the cyclical imagination and the human body – ways that bodily and cultural experiences are mediated by race and language, and the role of memory (physical and cultural) and history in configuring identities, including false and fragmented ones on the Web and in net culture. Isi-pikiskwewin Ayapihkesisak (Speaking the Language of Spiders) as a Web work uses such a structure in creating overlapping worlds.

Aboriginal people are visible as part of Canadian culture, but far less so in the United States. I will also argue for initiatives in the computer world that would acknowledge Native American presence and practices in the United States and abroad.

Critical Art Ensemble
c/o Steve Kurtz
kurtz@andrew.cmu.edu
mailer.fsu.edu/~sbarnes

Two electronic body types currently haunt the labyrinths of cyberspace. The first is the virtual body – the body of electronic utopia in which wish-fulfilling life simulation becomes the desired state of being. While it is an interesting concept, this hyper-aestheticized body never seems to come to fruition and remains contained and stunted in the halls of the video arcade. The second type of body is the data body, which has two manifestations: a persona that an individual can control and a body that is beyond individual control. Some of the utopian promises of the electronic realm have been fulfilled in the former.

Sublime and Impossible Bodies

My online persona can send, retrieve, exchange, and store information (although there are limits) in a quick and efficient manner. For those who have this type of electronic access, it's a very nice luxury. However, the institutional data body is not so desirable. First, the online activity of this data body is recorded activity. Life becomes an electronic file. This situation may or may not be a problem; however, if an individual is doing anything subversive, transgressive, or threatening to electronic power vectors (and such threats can be as simple as living in poverty), in all probability problems will emerge, because punishment systems will be alerted to any activity beyond narrowband normalization. Another problem is that the data body becomes the original, and the organic body becomes the counterfeit. An individual's physical self is recognized as legitimate and authentic only if it is validated by the data body (the total collection of one's electronic files). To complicate this matter further, an individual's data body is largely not in his or her control; it's in the control of a variety of institutions ranging from the medical establishment to creditors to security agencies. This lack of autonomy over one's own being in the world is among the most serious problems in the electronic present.

Is Robust Geometry Possible?

Organizer
Kevin Weiler
Autodesk

Panelists
Tom Duff
Pixar Animation Studios

Steve Fortune
Bell Laboratories

Chris Hoffman
Cardinal Technology, Purdue University

Tom Peters
University of Connecticut

Computer graphics and its accompanying design and analysis applications strongly depend upon geometric representations and operations. Yet geometric computing involves accuracy issues that are more complex and difficult than simple number representation and calculation issues because it involves maintaining additional constraints between sets of values, exacerbating an already serious digital representation and calculation problem. The effects of these problems in applications can range from simple visual discrepancies in displays (such as “cracks” between polygons) to strange program behavior including outright crashes. Anyone who has tried to implement a geometry-based algorithm quickly becomes aware that there’s more to it than meets the eye. A large amount of programming time in geometric computing is spent devising ad hoc solutions to these problems, and much user time is devoted to avoiding error-prone portions of geometric programs. The problem can get even worse when erroneous geometric results emerge from the computer into the real world, such as in the manufacture of mechanical parts used in complex everyday objects like automobiles and airplanes. There, the problems may result in lost schedule time and rework during manufacturing, or may not even surface until much later, with potentially disastrous results.

Until recently, most of these issues have been privately debated among friendly colleagues but publicly swept under the rug by both academicians and practitioners. Typical geometric implementation strategies of using “fudge factors” and double-precision numbers offer only limited relief. Despite the relatively recent beginnings of research into this area, there has been little organized effort to provide general methodologies for either avoiding geometric errors or proving the correctness and accuracy of geometry-based computer programs. Yet the majority of 2D and 3D graphics, design, and analysis applications depend on geometric computing in some fashion.

The panel members have been selected from both industry and academia because they represent a variety of different approaches and perspectives on the geometric computing problem. In addition to having significant academic credentials in this field, most of them have either worked in or consulted in industrial settings solving practical problems.

Topics:

- The nature and causes of geometric computing and accuracy problems.
- Specific graphical and animated examples to illustrate the issues.
- The variety of approaches being used to attack these problems.
- Personal experiences and practical advice on geometric programming.
- Perspectives on both the past and future of robust geometric computing.

In addition, we hope to engage the audience to share their own personal experiences, perspectives, geometry horror stories, and practical advice on geometric computing implementation strategies.

Tom Duff
Pixar Animation Studios
td@pixar.com

Three important areas where we encounter robustness difficulties in rendering are local and global illumination, and scan-converting higher-order surfaces.

A variety of local illumination problems require transforming coordinates of closely spaced points from model to screen space, doing some computation on them, and transforming back to model space before taking differences between them, for example to compute tangent vectors to surfaces or local sampling rates for texture maps. Often, the transformations from model to screen-space have large condition numbers (two ways this can easily happen are if the client specifies a near-clipping plane unrealistically close to the viewpoint, or when a projection with a very narrow viewing angle is taken from a great distance), causing large uncertainty in the inverse transformations. Subtraction of quantities with large uncertainty leads to catastrophic cancellation and, often, worthless shading results.

Is Robust Geometry Possible?

In global illumination computations, we often find ourselves either dicing primitives in ways that may expose geometric robustness problems of the sort usually found in computational geometry or CAD applications, or performing large numbers of ray-intersection tests, as in ray-tracing and Monte Carlo radiosity techniques. Ray intersections are notoriously hard to compute correctly, if only because secondary rays always start at a surface and can easily cause a spurious intersection there. The usual epsilon-test method of handling these performs badly, especially for glancing rays, which may have a legitimate intersection quite close to the spurious one. Interval arithmetic techniques handle this problem very well and deserve to be more widely adopted.

Another area in which interval arithmetic is useful is scan-conversion of higher-order surfaces. Interval methods produce guaranteed, 100-percent robust, answers to numerical questions that are hard even to pose in a pointwise (non-interval) context and can be a powerful tool for geometric problems when algorithms are designed specifically to take advantage of the interval perspective. Retrofitting intervals into existing algorithms is mostly pointless, although many arguments against their use stem from assuming that this is the only possibility.

Steve Fortune
Bell Laboratories
sjf@research.bell-labs.com

Geometric algorithms are usually described in the conceptual model of the real numbers, with unit-cost exact arithmetic operations. Implementors often substitute floating-point arithmetic for real arithmetic. This leads to the well-known problem of numerical robustness, since geometric predicates depend upon sign-evaluation, which is unreliable if expression evaluation is approximate. An ideal solution to the problem of numerical robustness would be simple, efficient, and widely applicable. No such solution exists or is likely to exist.

The most attractive available solution to robustness is exact evaluation of geometric predicates. Exact evaluation simplifies reasoning about the correctness of algorithms and can reduce the number of special cases. Performance concerns can be addressed by the use of adaptive-precision arithmetic, which with careful engineering reduces performance cost close to floating-point arithmetic. Exact arithmetic is the only way to obtain trustworthy implementations of complex geometric algorithms (e.g., Boolean set operations on polyhedra).

Chris Hoffmann
Cardinal Technology, Purdue University
C.M.Hoffmann@worldnet.att.com

Geometric computations are a mixture of numerical calculations, such as determining the intersection of two surfaces, and logical deductions based on the numerical computations. When the computations are done with floating-point arithmetic, the errors incurred may lead to uncertain deductions. Worse, different numerical computations may answer the same logical question in two different ways, without anyone realizing it.

So, what is one to do? Of course, you could fix the numerical computation, or try a different deduction method! Both have been tried. To fix the arithmetic, one can use exact computations. That eliminates the robustness problem entirely, but it is too expensive. Fixing the process of logical deduction is harder to do and must be based on a ternary logic: "yes," "no," and "unclear." Many attempts have been made, but they fell short, usually by failing to account fully for the "unclear" outcome of a test.

There are also more radical proposals to repair the robustness problem in geometric computations: You could redefine the problem so that troublesome inputs are changed to "nice" ones. Or you could redefine what the result should be, so that wrong outputs are "right." Depending on the nature of your business, these solutions might work for you.

Current trends postpone "general" solutions in favor of targeted solutions for specific geometric problems. For example, if your numerical computation involves evaluating a univariate polynomial and testing the result for zero, then there are fancy techniques that make the "unclear" outcome almost as rare as being struck dead by a meteor. Better yet, they are not excessively expensive. And they do astonish the numerical analysts. Then again, that is just for evaluating a polynomial with one variable. Still, perhaps we can figure out how to leverage such small advances and accelerate what has been, for so many years, disappointingly slow progress.

Tom Peters
University of Connecticut
tpeters@eng2.uconn.edu

Scan-line techniques typify the relationship between geometry and graphics. Namely, the geometric vertices of a polygon are approximated by floating-point values within the resolution available. A discretization process then identifies all pixels within the interior of the polygon and shades accordingly. In complex scenes, topological anomalies often appear at the boundaries between adjacent polygons.

These topological anomalies arise from the failure to reconcile the concept of a boundary element with the realities of finite precision computation. By definition, even infinitesimally small perturbations of boundary points can cause them to fall into the interior or exterior. Even while topological connectivity of geometric models is represented exactly by symbolic data, this connectivity is determined approximately via floating-point values. The inherent disparities between exactitude and approximation can lead to semantic inconsistencies. In practical implementations, algorithmic tolerance values are intended to preserve semantic consistency between topology and geometry. Yet complete integration of suitable tolerance couplings remains a formidable unsolved software design problem.

New conceptual approaches are needed for semantic consistency between floating-point values for geometry and the symbolic representations for topology. Novel perspectives from mathematical semantics of programming languages afford some promise. The issues considered in defining neighborhoods within which point perturbations preserve equivalent model topology are of particular interest in animation and morphing. For instance, small geometric perturbations between successive animation frames can lead to undesirable self-intersections.

Kevin Weiler
Autodesk, Inc.
kjw@autodesk.com

Robust geometry requires more than simple numerical robustness. It also requires maintenance of constraints between sets of values, such as topological relationships. Exact arithmetic helps but does not guarantee that implementations of geometric operations will maintain these constraints unless the constraints are well specified and the operations are carefully implemented.

Since everyone "knows" geometry and how it should work based on their everyday experience, application programmers using geometric engines often mistakenly feel safe in making assumptions about what geometric constraints the underlying geometry packages support and maintain.

At the same time, geometric engine implementors rarely specify what constraints their packages create and maintain, and often don't know themselves how various combinations of geometric operations affect accuracy and consistency constraints.

But many steps can be taken to improve the situation. Further exploration of number representation methodologies is certainly required. We also need to develop a consistent vocabulary for explicitly defining accuracy in geometric computations as well as explicitly defining constraints and common sets of constraints between values (such as topological constraints). With proper definition and use, it may be possible for developers to specify, and for users to understand, the limits of their geometric computations.
In addition, most developers, even though they are not yet privileged to be using exact computations, don't even know how precision is affected by code they implement. We need better static and dynamic accuracy instrumentation directly in our compilers, and perhaps in our floating point hardware.

In the long run, entirely new paradigms may be required to handle geometry shape definition, manipulation, analysis, and visualization robustly with absolutely consistent semantics. But literally thousands of years of geometric theory and culture will not die easily, and in the meantime we have much work to do to put our current house, built from geometry and digital computer number representations, in better order.

Computer Vision in 3D Interactivity

Organizer
Mark Holler
Intel Corporation

Panelists
Ingrid Carlbom
Bell Laboratories, Lucent Technologies

Steven Feiner
Columbia University

George Robertson
Microsoft Research

Demetri Terzopoulos
University of Toronto/Intel Corporation

With microprocessor clock rates in excess of 350MHz, SIMD integer instructions commonplace, and shared memory multiprocessing available for under $3,000.00, integration of computer vision with 3D graphics is now more practical than ever. Tracking the user's head, hands, and body, and detecting gestures, is one obvious direction to explore to eliminate encumbering sensors and enable new modes of interaction. Another direction is using computer vision techniques to understand 3D structure and camera parameters in multi-view image-based scenes for the purpose of re-rendering the scenes as a user directs. Yet another is giving animated characters visual awareness of users and other characters to enable richer interactions. What will be the most compelling integration of computer vision with 3D graphics?

The panelists address a subset of the following questions: What information besides human-user attitude/gesture can be extracted from images to enhance 3D interactivity. What other input modes are compatible with gesture and when? Is computer vision technology good enough today to be applied in commercial applications? If not, when? Is there a set of computer vision software components that would be useful to people working in 3D interactivity. What are the best applications for image-based rendering. Is the compute load small enough to run on today's machines? If not, when? Does system architecture need to change? What about memory and bus bottlenecks when multiple-video input channels are added to a system nearly bandwidth limited rendering graphics. Is computer vision + 3D graphics a big enough combination to drive the need for multiprocessing? Are there other standards, performance improvements, or specialized functions needed in video and multi-channel video capture for computer vision applications?

Audio-Visual Tracking for 3D User Interfaces

Ingrid Carlbom
Bell Laboratories, Lucent Technologies
Carlbom@research.bell-labs.com

Interactive virtual environment systems combine graphics, acoustics, and haptics to simulate the experience of immersive exploration of three-dimensional virtual worlds. Most such systems require users to wear cumbersome sensors for input and display units, eye- and headphones for the visual and auditory experience. However, the long-term goal for 3D interactivity is an interface more closely resembling human-to-human communication, depending more on multi-modal, unencumbering sensor and display technologies.

Tracking is a key technology for hands-free (unencumbered) 3D interactivity. Tracking can be used to determine user position and orientation, as well as user actions, such as gestures, facial expressions, and lip movements. While visual tracking with cameras alone has met with some success, the robustness of tracking can be increased if combined with acoustic tracking using microphones. Integrated acoustic and visual tracking can drive visual and auditory input, as well as output, to enhance the sense of immersion in a virtual world.

Camera and microphone-based tracking can be both complementary and cooperative to achieve accurate user localization. Camera-based tracking is particularly useful in acoustically noisy or reverberant environments, or to continue tracking a user who has temporarily stopped speaking while continuing to move. Similarly, acoustic tracking information from a microphone array can be used to localize the person who is speaking when several persons are present. This is particularly important under poor lighting conditions. User localization enables foveated processing for more detailed analysis of a user's gestures and expressions, as well as focusing of microphone beams on a user for high-fidelity speech input.

Accurate localization allows visual and auditory output to be directed to the user. The visual focus can be changed to the user's location (e.g., perspective vanishing point opposite the viewer, gaze of an avatar directed to the user). Auditory display in the form of spatialized sound can complement and enhance visual cues to aid in navigation, communication, comprehension, and sense of presence in virtual environments. Maximum fidelity and minimum disturbance to others is achieved if the acoustic output signal can be steered towards the listener. With a known user head position and orientation, combined with loudspeaker crosstalk cancellation, it will become possible to produce 3D spatialized sound for a moving user with virtual loudspeakers.

Steven Feiner
Columbia University
Feiner@cs.columbia.edu

Augmented reality refers to the use of see-through displays to overlay graphics, audio, and other media on the user's experience of the surrounding world. To accomplish this so that virtual objects are spatially registered with physical objects, we must be able to precisely track the 3D position and orientation of the user's head. As cameras and the compute power needed to process their input rapidly decrease in size and cost, the prospect of using computer vision for tracking becomes increasingly brighter. I discuss some of the issues involved in tracking for augmented reality, and potential advantages and disadvantages of using vision-based approaches. For example, one significant distinction of vision-based systems is the rich nature of the raw sensor data itself. Unlike other tracking technologies, input from one or more cameras can be used to perform object recognition, to build up a model of the surrounding environment, or just to document the user's experience.

Mark Holler
Intel Corporation
mark_holler@ccm.sc.intel.com

As computer performance marches forward according to Moore's Law, entirely new application domains are enabled. Digital imaging is currently going through a spurt of growth and will soon be followed by digital video processing. 3D graphics performance in PCs is also going through rapid growth now as 3D graphics accelerators proliferate. In addition to Moore's Law, there has been the addition of Single Instruction Multiple Data Instructions to most microprocessors. These instructions perform four or eight operations on four or eight pairs of 16-bit or eight-bit integers in parallel, typically in one clock, enabling a number of image-processing functions used in computer vision to be accelerated by 2-4X. Optimized libraries to achieve this acceleration are available for download on the Web [Performance Libraries]. Support for symmetric multiprocessing in mainstream CPUs such as the Pentium II and operating systems such as Windows NT has also provided a quantum leap in compute power available for integration of computer vision with 3D computer graphics. Bradski (1998) has reported a four-degree-of-freedom, 30fps head tracker using under 30 percent of one Pentium II CPU in a multithreaded app where head position/orientation controls fly above a 3D model of Hawaii. The second CPU and an E&S RealImage 3D accelerator are fully utilized for 3D rendering.

Immersive VR using HMDs requires a participant to wear the display and most often cumbersome sensors on head, hands, and body. "Fish-tank VR" (non-stereo) using computer-vision-based head tracking offers a less immersive experience but still provides control of motion parallax while freeing the user from wearing hardware. Arthur et al. (1993) and Rekimoto (1995) have shown that fish-tank VR enables users to understand complex 3D scenes more accurately than when given just static views. Ware et al (1993) have shown that motion parallax is a stronger cue for understanding 3D structure than stereopsis, suggesting that fish-tank VR is more effective in providing 3D cues than a stereo display, in addition to not requiring the user to wear shutter glasses. The narrower field of view of typical fish-tank VR systems is less likely to produce motion sickness.

Intuitive navigation in 3D spaces fundamentally requires more input than a mouse can provide. The mouse provides two degrees of freedom simultaneously while full 3D navigation requires six degrees of freedom, or more if viewing is de-coupled from navigation. Hand-controlled devices with six degrees of freedom require more attention to control than may be available during a 3D interactive game. Computer vision can extract some or all of the degrees of freedom from head position and orientation to reduce the required attention to hand coordination. Head movement such as peeking around corners to produce view changes is very intuitive for humans because we do it all the time in the real world. Used conservatively, tracking also promises to lower the interactivity bar for young children because of reduced requirements for fine-motor control.

Computer vision is capable of extracting 3D structure information from stereo views or motion sequences. With the view morphing approach [Seitz, Dyer 1996], a full 3D model of the scene need not be extracted to produce the novel views. This information is useful in producing novel views of an image-based scene. One can imagine an interactive telepresence application in which trackers know the positions and head orientations of participants and morph available view images to achieve eye contact and motion parallax cues. We have demonstrated such a capability in our labs.

Performance Library Suite: MMX technology optimized libraries in Image Process, Pattern Recognition, Signal Processing and Linear Algebra can be downloaded from developer.intel.com/design/perftool/perflibst/index.htm

G.Bradski. Computer Vision Face Tracking For Use in a Perceptual User Interface, Intel Technology Journal, developer.intel.com/technology/itj/q21998/articles/art2.htm Q2, 1998

K.Arthur, S. Kelogg, S.Booth, C.Ware. Evaluating 3D task performance for fish tank virtual worlds. ACM Transactions on Information Systems, Vol 11, No. 3, pp 239-265, 1993.

J.Rekimoto. A Vision-Based Head Tracker for FishTank Virtual Reality - VR without Head Gear. Proceedings of the 1995 IEEE Annual Virtual Reality International Symposium, March 1995 pp: 94-100.

C. Ware, K.Arthur, and K.S. Booth. Fish tank virtual reality, in INTERCHI'93 Conference Proceedings, pp37-42, 1993.

S.Seitz, C.Dyer. View Morphing. Computer Graphics Proceedings, SIGGRAPH 96.

George Robertson
Microsoft Research
Ggr@microsoft.com
www.research.microsoft.com/ui/ggr/ggr.htm

Virtual Reality apparently attains its power by captivating the user's attention to induce a sense of immersion. This is usually done with a display that allows the user to look in any direction (like HMDs or CAVEs), and that updates the user's viewpoint by passively tracking the user's head motion. However, there are other forms of VR where immersion occurs. Fish-tank VR uses a desktop stereo display rather than surrounding the user visually. Desktop VR uses animated interactive 3D graphics to build virtual worlds with desktop displays and without head tracking.

Current HMD-based VR techniques suffer from poor display resolution, display jitter, and lag. These problems tend to inhibit the illusion of immersion. Fish-tank VR uses desktop stereo displays to solve display resolution and jitter problems. Desktop VR solves all three problems, but at the expense of losing stereo and head tracking. Studies have shown that head-motion parallax is a stronger depth cue than stereopsis. Hence, adding head-motion parallax to a Desktop VR system could bring it quite close to fish-tank VR capabilities. Computer vision can track the user head motion without the user wearing any tracking sensors. This has additional benefits of eliminating fatigue and making it easier (and more desirable) to use, thus enabling everyday or extended use.

Computer vision enables other capabilities that may make 3D interactivity more effective and enjoyable. Adding awareness to our systems becomes possible. The system can know whether the user is present, whether the user is facing the screen, whether the user is engaged in some other activity (like talking on the phone or to another person in the room), and what the user is looking at on the screen.

Combining computer vision and 3D does involve solving some problems. The devices (cameras) are not expensive and are becoming ubiquitous. In the near future, the standard PC will likely include a camera. However, computer vision is computationally expensive. We currently use multiprocessors, which are a bit more expensive. We are nearing a point when computer vision and 3D interfaces can be effectively integrated and enable a number of exciting new interface capabilities.

Demetri Terzopoulos
University of Toronto/Intel Corporation
dt@cs.toronto.edu

Interactive 3D virtual worlds populated by autonomous characters with realistic behaviors rely on perceptual information processing, especially computer vision, so that the characters can sense one another and the user. I review the state of the art of perceptual modeling for behavioral characters and discuss how new vision algorithms promise to couple interactive characters much more closely to the user.

Ray Tracing and Radiosity: Ready for Production?

Organizer
Jacquelyn Ford Morie
Blue Sky | VIFX

Moderator
Richard Hollander
Blue Sky | VIFX

Panelists
Chris Wedge
Blue Sky | VIFX

Gonzalo Garramuno
Digital Domain

Grant Boucher
Station X Studios, LLC

Bob Powell
Rhythm & Hues Studios

The computer graphics community has a 25-year history of developing and constantly improving rendering tools, and at this milestone it is time to examine where and how these advances have been implemented in commercial production.

Ray tracing and radiosity (and hybrid methods that combine these two techniques) are rendering methods that produce the most realistic images possible in computer graphics. These techniques, however, require extremely long time-frames to compute. Production for feature films, special effects, and commercials is typically bound by deadlines that may not allow for the luxury of using these rendering techniques. Does the advantage of producing better imagery justify the problems involved in their use? Does production really need them?

There are still relatively few production companies that use global rendering systems for film and commercial work. This panel provides an opportunity for SIGGRAPH 98 attendees to hear from some practitioners who have had actual experience using ray tracing and radiosity in film production. Panelists discuss the issues involved in these techniques (including such things as ease of use compared to other popular software, user interface concerns, correlating with standard lighting terminology, and rendering times), show examples from actual production work, and consider what they have learned. Attendees also hear from a production company that has chosen to use these methods only rarely because they feel that other rendering techniques provide all the necessary tools they need to create any look.

Finally, the panel explores where production rendering is likely to go in the future, as both hardware and software continue to advance.

Chris Wedge
Blue Sky | VIFX
chris@bluesky-vifx.com
www.bluesky-vifx.com

Research and production are strange bedfellows, especially in the world of commercial computer graphics animation. One is not possible without the other, but they take place in very different environments. For example, ponderous ray-tracing algorithms and delivery deadlines don't mix well at all. Normally, ray-tracing is too slow, which translates into "too expensive," for production. In some high-end production environments, however, optimization and careful production planning are used to force them to coexist.

We examine current uses of advanced rendering techniques in production. Techniques that make ray tracing possible under tight schedule and budget constraints are explored. The future of high-end rendering is also considered.

Grant Boucher
Station X Studios, LLC
grant@stationXstudios.com
www.stationXstudios.com

Only the most advanced forms of mathematics and algorithms properly simulate the lighting and rendering models necessary for seamless effects work within live action films, which is the hallmark of the modern special visual effect. Rather than focus on finding hacks and work-arounds for simulating proper lighting models, the Station X development team has instead taken the approach that ALL of our work will be generated with the latest rendering technologies, and we will implement these in a production environment by deploying more horsepower and through careful optimization of the algorithms themselves. Fortunately for us, the faster CPU technologies like the 600mhz DEC Alpha processors are actually less expensive and easier to maintain than the legacy MIPS CPU-based technologies that remain the current industry standard.

Ray Tracing and Radiosity: Ready for Production?

Because of this, and because our core rendering software (an adaptive ray tracing solution based on the commercially available LightWave 3D renderer) has been optimized over the course of eight years, Station X has been able to deliver high-end photorealism for the same or lower hard costs as other industry-standard rendering solutions. Early in 1998, our team delivered a full radiosity-rendered commercial for Kirin Beer. This commercial was accomplished in a normal production environment and schedule (five weeks) and required only a modest increase in facility rendering resources. We believe the time has come for the entire special effects industry to embrace more physically correct rendering. Today's savvy special visual effects audiences now demand the uncompromising quality that such technologies and techniques afford us.

Gonzalo Garramuno
Digital Domain
gga@d2.com
www.d2.com

My company, since its inception, has relied heavily on the use of numerous third-party applications to achieve their goals, with no preference for hardware platform or rendering method. Ray tracing has been used and is still used extensively on both commercials and features, and the first attempts at using radiosity within the company have been very successful. Those techniques have also been used and will probably continue to be used, sometimes by themselves and sometimes integrated with other rendering solutions.

Bob Powell
Rhythm & Hues Studios
bpowell@rhythm.com
www.rhythm.com

Rhythm & Hues is widely known for its photorealistic rendering look in creating some of the best special effects in our industry, using our own proprietary renderer. This renderer has been in development with a team of supporting engineers for almost 10 years. During this time, the renderer has been used on such award-winning projects as "The Race for Atlantis" theme park ride, the movie "MouseHunt," and the Oscar-winning animals in the movie "Babe." However, at our studio, ray tracing is rarely used in our lighting efforts, even though our ray tracing is easy to use and runs well within standard production timelines. It's just that we do not feel the need to use ray tracing to achieve our lighting goals. The cost of ray tracing is usually not viewed as effective for the results. Other approaches not only work well, but are easily as effective in our production pipeline and yield the desired results. Although ray tracing is an important part of our renderer, it is by no means the most important or widely used technique available to us.

Characters on the Internet: The Next Generation

Moderator
Steve DiPaola
Darwin Digital - San Francisco

Panelists
Barrett Fox
infoplasm

Athomas Goldberg
New York University

Mark S. Meadows
Construct

Celia Pearce
Celia Pearce & Friends

Steve DiPaola
Darwin Digital - San Francisco
steve@darwindigital.com
www.darwindigital.com

The Windows-based desktop metaphor. Text and graphical user interfaces. Multimedia displays of moving images and audio. These three concepts constitute the language we use to communicate, to educate, to entertain with our computers and the Internet.

And yet in our daily lives we communicate and engage in a totally different way. We talk with our friends and relatives. We watch their facial expressions, read into their pauses, their vocal inflections and hand gestures. This is the language, the syntax, in which we are all truly experts: communicating and engaging interactively with people, with characters. Characters that emotionally engage and entertain us through films, plays, television, cartoons, and comics. Characters that inform, educate, and try to influence us, such as teachers, sales people, and business colleagues. Characters that have personality and spirit.

There is a real schism between the metaphors and interfaces we use with our interactive systems and those we use in our ubiquitous life. The high-end computer animation industry is now quite mature. It has both the knowledge and techniques to create computer-animated characters that can communicate with and engage an audience, almost at the level of their live-action actor counterparts. Some of this knowledge and experience has been successfully transferred over to Internet-based characters. But with few exceptions, character animation is still mimetic to the linear style associated with film, TV, and comics.

We are now at a seminal point in time where it has become possible to combine the emotive and communicative qualities of characters with the interactive, programmatic, and alternative narrative technologies of the Internet. Characters we can talk and listen to with speech recognition and synthesis. Characters who exhibit the illusion of life and cognition via artificial life/intelligence algorithms, information filter and retrieval capabilities, and behavioral models. These technologies can be combined with emerging communication and narrative metaphors such as multi-user worlds, interactive and participatory performances, and interaction between human-directed avatars and computer-controlled agents.

Confronting the challenges head on, all of the panelists have been associated with character technologies for many years and are currently working with or on practical innovative Internet-based character systems. The goal of this panel was to demonstrate these systems through interactive discussions with real-time characters and computer-controlled agents, giving the audience a practical glimpse of where Internet-based character systems are heading in the next few years.

Barrett Fox
infoplasm
pajamas@sirius.com
www.infoplasm.com

Character animation, as a multidisciplinary art form, has a long history and is well explored. Examples abound of characters imbued with the ability to richly emote and communicate. But characters stand at a nexus caused by the convergence of a staggering array of emerging technologies. Just as the advents of live theater, audio, video, and the computer demanded advancements in the craft, the Internet and technologies that it enables stand to present us with dramatically new creative challenges.

At the heart of this nexus is the ability to marry the power of a programming language with the latest techniques for character visualization. A character can be the visual manifestation of the latest artificial life algorithms, exhibiting stunningly subtle and complex behaviors. Or it can react dynamically to the laws of a physical simulation. Coupled with an endless procession of other technologies such as voice recognition and synthesis, streaming, multi-user environments, real-time puppeteering and information filtering and retrieval, characters can become effective interfaces, agents, and representatives for us. Indeed few other constructs can simultaneously embody so many disparate new ideas as meaningfully and cohesively.

Because of the facile nature of the digital arena, a character can manifest itself in many different environments. Whether it be rendered animation frames, real-time 3D, or dynamically controlled sprites, characters can appear in applications, virtual environments and on our computer desktops. Characters with expressive, interactive behaviors can exhibit autonomy and intelligence while containing messages, hyper links, files, and even viruses. The possibilities are fractally complex.

And while the characters we create will, undoubtedly, always be a simplification of ourselves and our environments, this very simplification, when executed creatively and thoughtfully, can deliver worthwhile insights. Because they are manifested in this new multifaceted medium, they may be a most appropriate instrument to comment on technology's effect on the world. Couched in an effective artifice, they can be a magnificent lens, simultaneously distorting and revealing unknown aspects and consequences of the bug-eyed juggernaut that is the Information Age.

Athomas Goldberg
New York University
athomas@mrl.nyu.edu
www.mrl.nyu.edu/athomas

The IMPROV Project at NYU's Media Research Lab is building technologies to produce distributed responsive virtual environments in which human-directed avatars and computer-controlled agents interact with each other in a lifelike, believable fashion in real-time. There have been a number of different approaches to this problem, which range from self-actualizing computer organisms to various adaptations of robot-control theory. In the former case, results have been limited to only the simplest "creatures" with only elemental skills, goals, and motivations. In the latter case, intelligent behavior is defined as the ability to navigate space or pick up objects, which ignores the depth and breadth of that which make human (and non-human) behavior interesting, namely the kinds of complex motivations and behaviors that are based as much, if not more, on emotional experience as they are on logical problem-solving. Therefore, the focus of our research has been on creation of authoring tools that will allow artists to carefully craft the personalities and behaviors of Web-based interactive characters based on their own understanding of human nature and dramatic involvement, characters who respond to human participants and each other in ways that reflect the goals and intentions of the artist, while always maintaining the "illusion of life."

Mark S. Meadows
Construct
pighed@construct.net
www.construct.net www.subnut.com

First, SOUL. There needs to be something to say. A designer without a message builds a character without a soul. This usually means that the designer ends up distilling some part of someone's personality, body, mannerism, dress, etc. Their [or someone else's] soul.

Second, TECHNIQUE. The technical defines the technique, and the technique defines the message. Any good designer can make a character, but the difference between quality character design and mediocre character design lies in the designer's abilities to work within the limitations of the medium. This is the case with any art, be it photography in the 1920s, painting in the 1800s, or architecture in the Middle Ages. Technology is the sysiphean stone. The design is the hill. Pushing it over is the pleasure.

Third, PROPORTION. The classic rules of proportion should never be forgotten, should never be blindly followed, and should serve as a point of evaluation to the designer. These can be used to develop stereotype and individuality.

Fourth, THEME. A consistent vocabulary helps things along. Picking thematic choices is a good idea, and relying on some consistent metaphor for the design of multiple characters helps make design decisions that might otherwise be overlooked.

Removing the Fourth Wall: Creating Interactive 3D Worlds

Celia Pearce
Celia Pearce & Friends
celia@cpandfriends.com
www.cpandfriends.com

Let me shock you all by saying that the future of 3D graphics on the Web is in the hands of ... not artists, or animators, nor engineers ... but WRITERS. Here's why: Characters and story are one and the same. And story is what writers do. The challenge as a writer is how to develop new story systems that allow for audience interaction. To do this, we have to begin to think of a story as an algorithm – no longer a linear construction, but a dynamic evolving system – a narrative ecosystem, if you will. If we can deconstruct classical story structure to its basic components, we can begin to create a mathematic model for dynamic narrative. In my current work with LEGO's Wizards group, and on other current work, I am beginning to do hands-on experimentation with some of the ideas I spoke about in last year's "Narrative Environments" panel. I'd like to share a couple of revelations I've had in this process:

1. **Story is about relationships between characters.** As we develop algorithms for interactive characters, we might want to reconsider the individual character "personality engine" in favor of a "relationship engine" that translates into a relationship matrix that maps the interactions between characters.

2. **Characters are driven by intention.** The foundation of story structure is character goal or intention. Thus, it's important to model the individual character with a strong intention set that is designed to fit into a larger relationship matrix described above.

3. **Keep it simple.** For a long time, people were looking at long-form narrative for models. I've recently become enamored of short-form narrative models, especially cartoons. Cartoons are short and simple, and generally revolve around an ongoing conflict between characters. They do not rely heavily on dialog for their humor or narrative impact. And they use a simple animation style that is very forgiving of low-resolution solutions.

Now: how do you get the player "in on the act?" Once you create a dynamic narrative system, as described above, it becomes much easier to bring the player into the story: she simply becomes another node in the relationship matrix. If she enters the world with an intention, just like an autonomous character, she can then participate in an evolving story in real time. An object-oriented dynamic narrative system such as this is thus scalable: you can add or subtract both autonomous characters and players and the story will continue based on the matrix parameters. In this way, you can begin to create self-evolving, self-generating stories. The neat thing about this is that it also eliminates the need to keep creating new content. No more "webisodics." The story just creates itself. This is good news for lazy writers like me, because it means once I set the system up, I can go on to something else!!!