Provocative speculation on where computer graphics will take us in the new century.

Panels

Jeff Jortner
Sandia National Laboratories

Contents

Panels

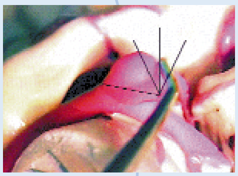

Special Sessions

Research

SIGGRAPH 99 Panels offer provocative speculation on where computer graphics will take us in the new century. Panelists from Asia, Europe, and the Americas survey a very wide range of issues, from digital protection of intellectual property to advanced techniques for practical visualization of massive datasets. Their insights and debates illuminate surprising possibilities. What's next depends on how the worldwide SIGGRAPH community responds.

Naturally, several Panels focus on issues that affect the large visual effects industry in California. Examples from recent feature films illustrate how animators and engineers simulate and control crowd behavior. Managers, producers, and creators explore visual effects production issues and how they apply today's research to tomorrow's problems.

Other Panels are already moving beyond tomorrow to think about the future of experiential art or interweaving reality with virtual reality to establish another paradigm: mixed reality.

Future

The Panels Committee is grateful to the organizers and panelists for their ideas, reflection, collaboration, courage, and humor, and their willingness to share their discussions with all of us who work in computer graphics and interactive techniques.

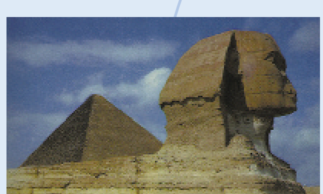

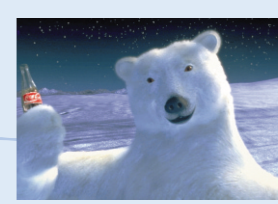

Panels Committee

Chair

Jeff Jortner
Sandia National Laboratories

Subcommittee

Joanna Alexander
Zombie

Rebecca Allen
University of California, Los Angeles

Aliza Corson
SIGGRAPH 2000 Panels Chair
Walt Disney Feature Animation

Clark Dodsworth
Digital Illusion/Osage Associates

Branko Gerovac
Massachusetts Institute of Technology

Andrew Glassner
Microsoft Research

Robert Judd
Los Alamos National Laboratory

Alyce Kaprow
the new studio

Committee

Organizer/Moderator
Nick England
University of North Carolina at Chapel Hill
nick@cs.unc.edu

Panelists
Industry experts selected
as late as possible.

Hot Topics in Graphics Hardware

Hardware

Graphics hardware gives life to interactive applications. It enables compelling, intriguing, and exciting experiences. The complexity of graphics hardware designs is staggering. Chips with many millions of transistors and many gigabytes/second of bandwidth are commonplace. The speed of innovation and new product introduction is equally staggering. And the end result is graphics performance that was unthinkable a few years ago.

In keeping with this rapid pace, this panel explores up-to-the-minute topics with leading hardware designers. Earlier in the week, top designers and developers from around the world attended a small, highly focused meeting, the SIGGRAPH/Eurographics Workshop on Computer Graphics Hardware. Some participants in that Workshop were invited to be panelists in this session, joining a few others who were selected based on recent exciting hardware product announcements.

The result is a "just-in-time" panel aimed at discussing the very latest developments and issues in graphics hardware design.

Rather than presentations by each of the participants (possibly repeating information that may be found elsewhere), there is only one presentation. In it, the moderator summarizes the most significant work presented in Hardware Workshop Papers and a special "Hot 3D Systems" session. Then, panelists comment on what they found to be the most interesting and controversial topics from the conference.

In addition to answering questions from the audience, the panelists offer their perspectives on some of the following issues:

- What is the future of "high-end" graphics in light of the market dominance and increasing performance of PC graphics cards?
- What applications will continue to demand workstation configurations (computer systems designed in conjunction with graphics) vs. add-on cards? Why? What are the bottlenecks?
- What features are needed for real-time visual simulation applications? Will add-on cards satisfy these? Is the market important enough to bother with?
- Is there any demand for features and flexibility beyond OpenGL/Fahrenheit? Is everyone comfortable with this?
- Is there any need for hardware to handle anything but polygons?
- Are dedicated geometry processors necessary? When and where?
- What Intel architecture changes are needed for greater polygon rates? For greater texture mapping rates? For new features?
- Where do graphics hardware designers come from? Should we be doing something to create more of them? Where does graphics hardware knowledge come from? Can we do anything to improve this?
- What's the coolest thing on (or not on) the trade show floor? Why? Does it do something new or does it do the same stuff faster and cheaper?
- What do you think about parallel processor designs?
- Plus more questions on newly announced designs.

The aim of this panel is to provide attendees an up-to-date snapshot of the state of the art in graphics hardware. Co-location of the Graphics Hardware Workshop provides an unparalleled opportunity for SIGGRAPH 99 attendees to get a low-latency, high-bandwidth view of what is happening right now in this exciting field.

How to Cheat and Get Away With It: What Computer Graphics Can Learn from Perceptual Psychology

Organizers
Victoria Interrante
University of Minnesota
interran@cs.umn.edu

Daniel Kersten
University of Minnesota

Panelists
David Brainard
University of California, Santa Barbara

Heinrich H. Buelthoff
Max Planck Institute for Biological Cybernetics

James A. Ferwerda
Cornell University

Pawan Sinha
University of Wisconsin

What psychology can offer computer graphics is essential insight into the fundamental workings of the human visual system: how we interpret what we see. Such knowledge provides a foundation for the deeper understanding of the science behind the art of effective visual representation. With a better understanding of how we interpret visual input, we are better equipped to determine how to design methods for representing objects and information so that they can be easily and accurately understood.

What computer graphics can offer psychology is a methodology for creating carefully controlled, physically valid representations of simple visual scenes. Psycho-physical experiments that investigate the effect on perception of a particular kind of stimulus (for example, the effect of luminance contrast on perceived depth) require careful control of all unrelated external factors, and graphics offers an easy way to do this. It also offers the ability to separately examine physically inseparable effects, such as the impact on depth perception of shading and shadows resulting from direct and indirect illumination.

David Brainard

Limits on the apparent realism of displayed images include not only the quality of the modeling and rendering software, but also factors inherent to the display itself. Typical displayed images are presented monocularly, do not track eye and head position, and are limited in size, resolution, light level, spectral composition, and dynamic range. We seek to understand whether and how display limits influence our perception of displayed images. To do so, we measure human performance both with free viewing of real objects and with displayed images of the same scenes.

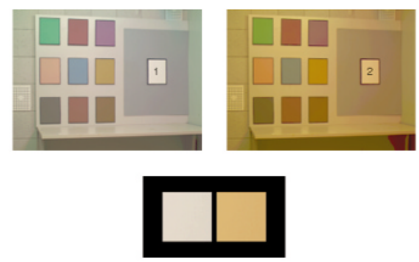

The same scene taken under two different illuminants, illustrating that the light reaching the eye changes drastically as the illuminant is changed.

In our experiments, observers judged the color of surfaces under different illuminants, and we used the results to quantify how well the observers' visual systems adjusted to the illumination changes (the degree of color constancy). Our displayed images were carefully rendered versions of acquired hyperspectral images of real objects, and we used the same observers for both real and displayed conditions. Nonetheless, we found quantitative differences in the degree of color constancy. We are currently assessing more systematically what factors lead to this difference in performance.

Heinrich H. Buelthoff

In the Max Planck Institute for Biological Cybernetics in Tübingen we study human perception and action using state-of-the-art computer graphics and virtual reality technologies. Our research is strongly motivated by the idea that in order to understand and interact with the world around us we do not need to build a full representation of it in our head. Our psycho-physical work on object recognition showed, for example, that if two views of a novel object were learned, recognition was better for new viewpoints oriented between the two training views than for viewpoints lying outside the training range. These effects were at odds with most of the current theories of object recognition and helped establish the image-based approach to object recognition. This theory proposes that objects are represented as a collection of views rather than explicitly related parts or 3D models.

This framework can be used not only for the recognition, but also for synthesis of 3D graphic objects and might help the 3D graphics community to build better and more efficient image-based rendering and immersion systems. One good example for what perception has to offer computer graphics comes from our psycho-physical studies of face recognition. Human subjects can recognize a face under many different viewing and illumination conditions even if this face has been seen only once in a photograph. A similar performance in synthesizing many novel views of a face from a single photograph has been achieved by Blanz and Vetter in our lab with an example-based face synthesis system that learns how images are transformed during viewpoint changes.

The virtual reality lab at the University of Tüebingen with a projection of a virtual city on a 180-degree screen.

We found also that for the animation of biological motion it is much more important to present an accurate sequence of 2D images than an accurately animated 3D model. For example, the recognition of point-like walkers in a stereoscopic display is unaffected even when the objects' depth structure is scrambled. Furthermore, the observers seem perceptually unaware of any depth anomalies introduced by the scrambled depth structure. Apparently, expectations about a familiar object's 3D structure overrides the true stereoscopic information. Therefore it might not be necessary for an efficient 3D stereoscopic display system to rely on accurate 3D models.

In our "Virtual Tüebingen" project (www.kyb.tuebingen.mpg.de/bu/projects/vrtueb/index.html) we built a virtual model of our hometown to study the relative importance of 3D geometry versus texture map information for the purpose of navigation. We can also learn how such information is transferred from training in virtual environments to real environments. This will help us to build better training simulators in the future.

James A. Ferwerda

We create images to provide information to human observers. In many cases, realistic images are desired because it's felt that they provide higher quality information. Physically based image synthesis methods can produce accurate simulations of light reflection in scenes, and within the graphics community, it's been taken as a matter of faith that this will lead to more realistic images, but is this necessarily the case? Do physically accurate images contain higher quality visual information, and is this information worth the added computational expense?

We use the information that vision provides to accomplish important tasks. Walking, reading, recognizing faces, using a tool all require different kinds of visual information. There's an intimate link between the kinds of visual information we look for in a scene and the kind of task we're trying to accomplish. Vision researchers have been interested in the relationship between tasks and visual information for decades, but until recently it's been hard to do meaningful experiments because of the difficulty in manipulating the visual features of complex scenes in controlled ways. With the development of physically based image synthesis methods, vision researchers now have a powerful new tool for studying how we use visual information to accomplish the tasks we need to do.

These developments have led to a new line of research in visual perception that asks two fundamental questions:

1. What is the visual information that specifies properties of objects in a scene?

2 How well do we use that information in performing the tasks we do?

Ideal observer analysis within a framework of Bayesian inference allows these two questions to be studied systematically.

Victoria Interrante

Most of us have little trouble distinguishing between a picture that "works" and a picture that doesn't. In the case of visualization, which is specifically concerned with the communication of information through images, one can use controlled experiments to objectively quantify the relative effectiveness of a particular method or approach in terms of the extent to which it facilitates the performance of a relevant task. But the essential design process (the art of developing a new representational strategy or simply creating an effective visual representation from known ingredients) remains largely guided by experience, intuition, and creative inspiration.

At the heart of my research in graphics and visualization is a quest for answers to the questions: To what extent is there a definable science behind the art of effectively conveying information through images, a theoretical basis for knowing what kinds of things to try, what should work and why? How can we best extract critical insights from research in visual perception and apply these ideas to important open problems in image generation?

When we look at a real object in the physical world with our two eyes, we derive an immediate, intuitive understanding of its 3D shape from the information provided by a multitude of complementary cues. Some of these cues arise as a direct consequence of the geometry of the viewing arrangement, which we get "for free," while others may be influenced by such things as the viewing context, the direction of illumination, the magnitude of the luminance contrast between light and dark regions, the properties of the surface reflectance, and the presence and characteristics of the surface texture, which we tend to regard as modeling parameters. Considerable progress has been made in the development of algorithms for photorealistic rendering, with the goal of enabling us to create computer-generated images of simple, controlled scenes that are nearly perceptually indistinguishable from pictures taken by a camera. Equally important, however, are efforts to devote attention to understanding the principles behind effectively "setting the scene," determining how to represent an object so that essential information about its 3D shape and location in space can be easily and accurately perceived. Research in perceptual psychology provides valuable insight into how we interpret shape from attributes such as shading, contour, and texture.

Daniel Kersten

There are two key elements in defining the problem of visual perception. The first is that useful information about the world, such as the shape, material, illumination, and spatial relationships of objects, is encrypted in the image. Second, the encryption process, of going from a description of the world to an image, is not, in general, reversible. Any single source of image information is usually ambiguous about its causes in the scene. Seeing is the process of decoding the image information.

Three-dimensional computer graphics simulates the process of encrypting scene information into the image. By creating images from synthetic scenes, we can gain insights into the constraints used by the visual system to decode image information and begin to bridge the gap between the simple images of the laboratory and complex natural scenes. Computer graphics modeling and animation tools provide the means to generate stills and animations that produce strong perceptual interpretations, but they are theoretically indeterminate.

Pawan Sinha

Introspection suggests that we have a good sense of the spatial relationships between some important entities in the visual world. For instance, knowing the position of the sun in the sky, we can confidently predict where on the ground the shadow of a flagpole will fall or where one might expect to see the sun's reflection in a lake. Given the confidence with which we can make these predictions, it would seem likely that images that have been artificially altered to destroy these relationships would be perceptually easy to detect. Experimental results suggest the contrary.

Even under prolonged viewing conditions, most observers are largely insensitive to anomalies introduced into real images of outdoor scenes using digital techniques. Often these anomalies were such as to render the scene physically impossible.

Our first attempts to account for this apparent failure of the human visual system to encode the spatial correlations between natural entities revolve around the simple idea of locality of spatial processing at a given moment. Thus, for instance, at any one time, we can either perceive the position of the sun in the sky or the patterns of shadows on the ground, but rarely both together. This "non-locality" of the two entities (sun and shadows) might be part of the reason for the visual system's inability to detect and code their mutual spatial correlations. This general idea suggests that anomalies within patterns that are small enough for all their constituents to be processed simultaneously will be perceptually evident. This is, indeed, what we find in our experiments with digitally modified patterns of shadows on images of human faces; subjects immediately report the anomalous nature of such images. An interesting result that is beginning to emerge from these experiments is that the notion of spatial locality might be more than just a general preset limit on processing area; it might, in fact, be tied up with a rather high-level notion of objects.

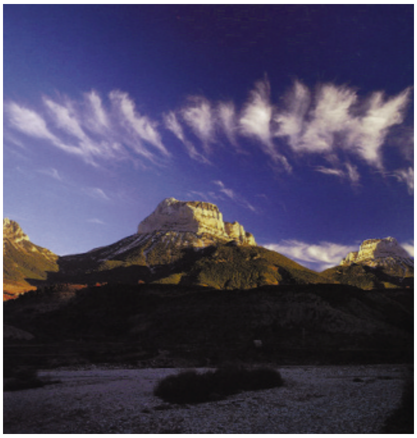

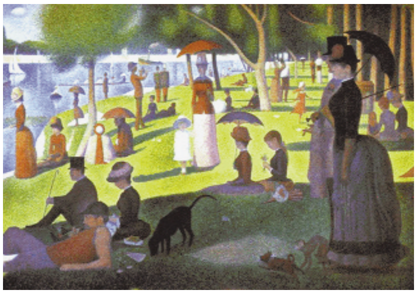

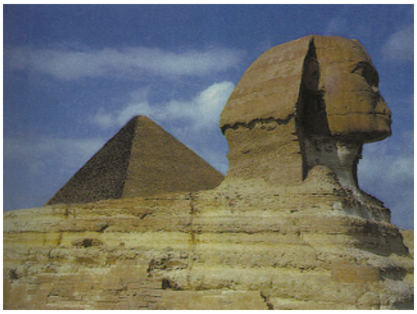

Images with some digitally introduced illumination inconsistencies.

CG Crowds: The Emergence of the Digital Extra

Moderator
Juan Buhler
PDI
jbuhler@pdi.com

Panelists
Jonathan Gibbs
PDI

Christophe Hery
Industrial Light & Magic

Dale McBeath
Pixar Animation Studios

Saty Raghavachary
DreamWorks SKG

Crowds have been used in film production more and more in recent years. Beyond behavioral simulation, what are the challenges faced when creating a crowd system for use in a feature film? How usable is a pure behavioral system, and how much "manual" control must be provided when the goal is to create crowds that help tell a story? How do you simulate and render the behavior of thousands of characters? This panel compares some solutions implemented for a number of different feature films.

Topics include:

- Pure behavioral simulation vs. keyframing each character: Where is the happy medium? For example, what are the limits of a pure behavioral simulation when working in a production environment, where crowds have to be "art directed?"
- How much control should a crowd system have? Will the animator be able to specify individual motion? Will the motion be simulated?
- How general can a system be? Is it possible to deal with problems on a shot-by-shot basis, or should a system try to solve every possible case?
- How should the motion of each individual character be generated? Are purely procedural approaches feasible? If cycles are used, what must be considered for their generation? What about foot registration and compliance with the terrain, for example?
- What user interfaces are best?
- How do you render hundreds or even thousands of individuals?
- Comparison of our experiences (*Star Wars: Episode I*, *A Bug's Life*, *The Prince of Egypt*, *Antz*).
- Future work.

Jonathan Gibbs – Antz

At PDI, we broke the crowd problem down into three parts: low-level motion (object deformations), high-level motion (paths through the world), and rendering. We chose to build our system around the use of hand-animated motion cycles. While it is possible to generate low-level motion procedurally, current techniques do not come close to the wide range of quality animation produced by talented character animators. Once these cycles are created, they can be procedurally modified to produce a wider range of motion, and to automatically fix certain problems that arise when the cycle is placed in a new environment or on a new character.

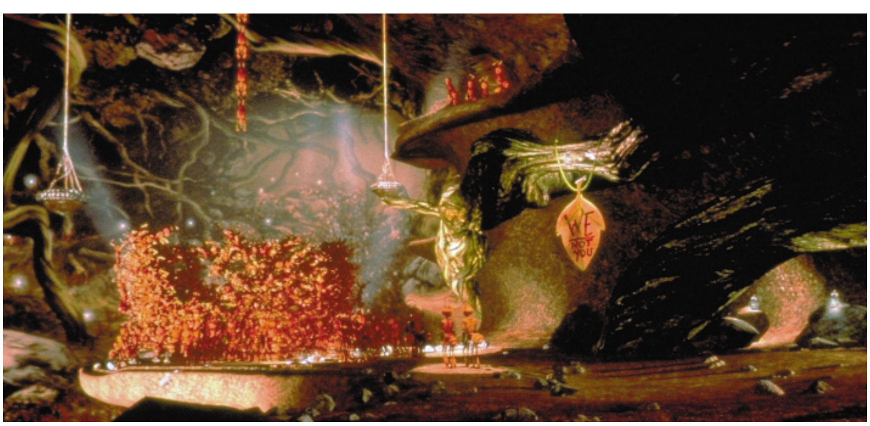

To compute the high-level motion, we use a crowd simulator that computes the global paths and selects from a set of actions for each character. It is important that the crowd simulator respect the motion in the cycle when it creates the global path. This allows subtle variation in speed and direction to be built into the cycle, and to influence the global motion of the character.

Due to its procedural nature, control is a big issue in crowd animation. We found it convenient to control most aspects of the simulator through maps (density maps, obstacle maps, flow maps) and to make all controls animate-able over time. Since there will always be shots that require the crowd to do something highly specialized, the simulator has a plug-in architecture that allows custom code and controls to be added on a shot-by-shot basis.

The final piece of the crowd puzzle, and the one least mentioned, is rendering. Rendering geometry can get very expensive very quickly. Since many characters are sharing the same (or similar) cycles, utilizing data coherence, geometry instancing features, and on-the-fly deformations in the renderer can reduce rendering times substantially.

Christophe Hery – Star Wars Episode 1: The Phantom Menace

For the third act in *Star Wars Episode 1: The Phantom Menace*, we were confronted with the issue of rendering images with potentially thousands of animated creatures. We needed high-level control on the crowd to the point where it could be easily intermixed with each hero character, especially when there were close-ups in the foreground of the scene. For directorial purposes, the choreography tool had to give immediate feedback and support a preview mechanism. On the rendering side, we had to deliver a method that would still compute the complexity of our maps and materials while keeping our costs in memory and CPU time at a manageable level.

After some investigation, we decided it was best to base all performances on a library of animated cycles. It also became clear that we wanted to take advantage of the delayed-read archive and primitives-generation features of RenderMan. Additionally, its ability to determine the level of detail in each scene was an important consideration. Animation cycles would still be created in Softimage. They would go through our normal creature pipeline until they were to be baked into sequences of rib files.

The choreography work was done in Maya in such a way that a particle represented each creature, effectively turning the pdb file into a container specifying the position, orientation, and variation of the color/material and rib file used for each entity. We wrote several support plug-ins, in order to help us pre-visualize our performances in near real time from within the Maya interface.

Dale McBeath – A Bug's Life

Yikes! There have been many technical designs and solutions to crowd animation put forth over the years, but it's when you're faced with the prospect of dealing with a cast of thousands that the sweat starts to pool. Pixar split the problem in two: technical and animation. As Supervising Animator for our crowd team, it was my goal to have each individual ant remain unique in activity, yet have the whole colony able to act and emote as a single supporting character.

Saty Raghavachary – Prince of Egypt

Crowds in *Prince of Egypt* were created with a mixture of elementary behavioral simulation and a procedural character placement system.

We took a very pragmatic, bottom-line approach to generating crowds when we developed the procedural system. Our approach is essentially sprite-based. We rendered each character as a sprite and then composited the sprites to generate crowd images. This approach allowed us to exercise separate control over characters and their placement, walk/run cycles, color and lighting, and rendering.

Moderator
Richard Hollander
Rhythm & Hues
reh@rhythm.com

Organizer
Jacquelyn Ford Morie
Rhythm & Hues

Panelists
Thaddeus Beier
Hammerhead Productions

Rod G. Bogart
Industrial Light & Magic

Doug Roble
Digital Domain

Arthur Zwern
Geometrix, Inc.

3D Tracking in FX Production: Blurring the Lines Between the Virtual and the Real

As image making has matured and expanded, especially visual effects, so has the need to blend real-world data (exposed imagery) with virtual data or computer-generated imagery. Doing so requires the effects artist to extract many things from the imagery. Exactly what needs to be extracted is always an expanding list but includes at least the path that the real camera took while exposing the imagery, the exact details of the dynamic lensing (focal length, f-stop, focus point, and all the distortions), and 3D data in the environment within the imagery. All this from a sequence of 2D frames!

This panel explores some of the current techniques that are used in this extraction process, ranging from the brute force methods to highly sophisticated image processing techniques. Topics include:

- Auxiliary cameras (monkey cams)
- How motion control (robotic camera systems) integrates with real-world applications
- How does one measure the accuracy of such a process?
- When do you need to do 2D tracking vs. 3D tracking?
- Lens distortions and their effect on the various techniques discussed.

Thaddeus Beier

One of the many problems encountered when combining computer graphics visual effects with live action is tracking one to the other. And one of these tracking problems is camera tracking, matching the computer graphics virtual camera to the film camera so that the synthetic elements of the scene move as if they were part of the original elements.

Many people have written tools to assist in this process. Hammerhead's is called ras_track. With this tool, the user first has to measure the position of 3D tracking points in the set, then track these points in 2D on the scanned film. With accurate 2D track points and 3D positions, and some simple approximations of what film does, it is straightforward to calculate a camera position and orientation. What makes the problem more interesting is that it is usually difficult to get accurate 2D track information and 3D position information, and some camera systems do not conform to simple models.

Two-dimensional track information is sometimes difficult to acquire accurately because points are obscured or go out of frame in part of the sequence, because lighting changes dramatically, and because the camera moves so much that the witness points change their appearance. Three-dimensional motion data is often hard to obtain. In many cases, it is impossible to obtain because no measurements were taken when filming, either because the objects were difficult to measure or because it wasn't known at the time that this was to be a visual effects shot. Very often, when the film is scanned, one finds that the carefully measured and recorded 3D points are out of frame, because the camera setup wasn't predicted correctly.

Real cameras cannot be accurately simulated with a pinhole camera approximation. This is especially true with wide-angle lenses. Interestingly, lens distortions can vary quite a bit depending on focus (and zoom, for zoom lenses). All of these things make camera tracking more of an art than a science. I maintain that it will always be so, that the decisions that tracking artists make are necessary creative decisions, and the most useful tool is one that provides as many controls as possible at all levels of the process.

Rod G. Bogart

Industrial Light & Magic has spent many years integrating computer-generated characters into feature film backgrounds. The early techniques relied heavily on motion control or fixed-camera positions. As computer graphics has become more accepted, directors and cinematographers have demanded freedom in camera planning and movement. This increased flexibility on the set or location shoot has resulted in a range of automated and assisted camera-matching software solutions. There is no single method for all shots. Simply tracking corners and texture details followed by prolonged iteration is fine when it works, but when it does not, the system must allow user intervention and modification. The result is a continuum of techniques from hand tweaking to 2D stabilization to object animation matching to full 3D data recovery and camera reconstruction.

Recent projects at ILM have made use of all these techniques. On the film *The Mummy,* we provided CG prosthetics to replace portions of the face and upper body of the actor playing the title role. On set, he was fitted with LED-laden objects plus additional marks on his skin. In post production, these were tracked to produce 2D data. The known 3D object was matched to the rigid points and the 2D data of the additional markers on the flexible areas were used to stabilize the final result. Similar lighted markers were used on the set to help reconstruct extreme camera moves. These careful preparations greatly improve the quality of the final camera and object matches.

By contrast, in *Star Wars Episode I,* literally hundreds of shots required matched camera moves to incorporate the CG elements. At this scale, all possible camera situations can occur. These range from simple pan/tilt, to dolly shots, to shooting wild. That volume and diversity of shots exposes many problems with tracking 2D data. Points often can be seen only for short periods. The camera derivation must be able to hand off or weight points as they become obscured. Because many CG characters enter or exit the frame, lens distortion must be taken into account at tracking time. Some camera moves must be split to accommodate whip pans part way through the shot. Otherwise, iteration techniques will fail. Although 3D tracking and reconstruction provides excellent results in controlled situations, we have found that these new techniques must be used in conjunction with user guidance and adaptations of 2D methods to accomplish large-scale feature film effects.

Doug Roble

I have been involved with extracting 3D information from images since 1990, when I realized that some computer-vision techniques could have a significant impact on special effects in movies and video. At that time, if a computer-generated effect needed to be inserted into a live-action scene, the camera was typically locked off or tediously hand tracked. In fact, this was a notorious tip that an effect was on screen: a previously dynamic, moving camera would stop and the effect would occur. One of my projects at Digital Domain has been to write a complete 3D/2D tracking package. It started as a program that required survey information to accurately calculate the location of the camera in a frame of film and is currently a large package that can tell you just about anything about the scene, provided the camera moved enough.

There have been a lot of challenges along the way. First, how do you accurately track a pattern on an image in 2D? The pattern should be tracked to sub-pixel resolution as it moves and should be followed as it rotates, scales, and skews. And after you think you've got these problems solved, the director lights the object that you are tracking on fire! How do you handle that? Next, the accuracy of the 3D track is of vital importance. If you are doing a set extension or placing a digital character in a scene, any shift, even sub-pixel, is unacceptable. A 3D track that looks good at video resolution can completely fail at film resolution. How do you deal with the error that naturally occurs with this constraint?

Now we can track without a survey of the set. How accurate is this, and what limitations are there? This technology lets us come up with a complete representation of a scene, the camera motion, the motion of moving objects, and a 3D reconstruction of the scene. Error also creeps into this process. How can it be minimized?

Three-dimensional tracking and scene reconstruction has enabled amazing effect shots that were previously considered impossible or too expensive. The Academy of Motion Picture Arts and Sciences recognized this in 1999, awarding two technical achievement awards to 3D tracking technology.

The future of this technology is very interesting. It will be combined with image-based rendering techniques and lighting analysis to provide artists with even more information about a scene.

Arthur Zwern

For years, 3D camera tracking for match moving was performed only by skilled specialists using manual iterative guesswork techniques. Then a range of software-based solutions to aid in this process was introduced. By tracking a handful of fiducials within the shot, these tools use photogrammetric techniques to effectively solve for extrinsic camera parameters (camera pose in six degrees of freedom at each frame) and often some intrinsic parameters (focal length and distortion). In some cases, the fiducials must be manually placed and surveyed targets, while in other cases visual features in the shot itself are tracked. However, each of these processes still requires significant human interaction to achieve useful results.

So, what's next? Automation! Already, we are seeing tools that eliminate the most burdensome task in today's camera trackers, which is managing selection and tracking of visual features. These tools scan every frame of the shot to extract the mathematically optimal visual features within the shot and robustly track those features for as long as they are visible. Typically, hundreds of features are tracked instead of just a handful, which improves results dramatically. If a feature becomes occluded or leaves the field of view and then becomes visible again, it is automatically identified as the original feature. By continually reacquiring new features in an automated control loop, shots of unlimited length can be tracked effectively. Under many conditions, the result is complete automation with sub-pixel match move accuracy.

Three-dimensional camera tracking is just one component of a major trend now beginning to blur the lines between 2D footage and 3D CG. For instance, every one of those thousands of features that get tracked in automated camera trackers also gets its 3D position in space determined by the very same algorithms. With today's computational geometry techniques, these 3D points can be converted into polygonal meshes and textured with pieces of the original footage. The result is not only the camera path, but a complete 3D model of the scene.

If we project ahead a few years, it is conceivable to imagine a film made entirely by scanning live shots into 3D, scanning talent into 3D and animating it with motion capture techniques, directing this CG talent at "post," and rendering the final result so well that few people can tell the film was built entirely in CG. Impossible? The technology is being developed today.

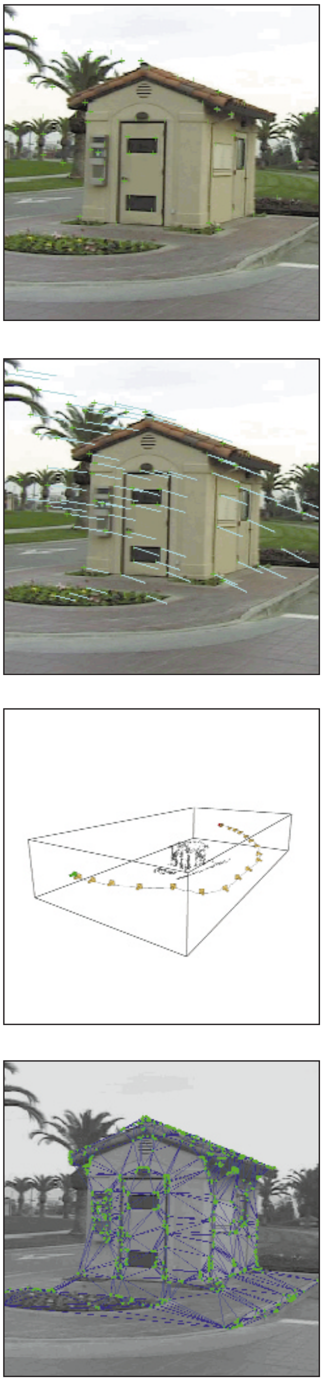

Steps in Geometrix's fully-automated SoftScene™ process (top to bottom)
1) Automated feature extractions
2) Automated feature tracking
3) Camera pather extraction
4) Scene geometry and texture extraction

Natural and Invisible Human Interfaces

Moderator
Michael Harris
Bear Systems, Inc.
mh@michaelharris.net

Panelists
Hiroshi Ishii
Massachusetts Institute of Technology

Caleb Chung
Giving Toys, Inc.

Clark Dodsworth
Digital Illusion/Osage Associates

Bill Buxton
Alias|Wavefront, Inc.

Huge advances in interface modalities are evident and imminent. This panel demonstrates and explores the most interesting, promising, and clever of these modalities, and their integration into exciting multimodal systems.

Because this is a gadget-intensive topic, the panel presents 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, and spatial sound generators. Some of these innovative transducers operate both non-invasively and invisibly. No one should ever have to see a computer. The complexity should be suffused in the world around you.

Panelist perspectives are theoretical and pragmatic, incremental and radical; their work is elegantly inspiring and often delightfully unconventional. All were formerly considered visionaries, but now their visions are achievable, and many industries are paying attention. They are seasoned practitioners with their own viewpoints. All are articulate, and none are shy.

Michael Harris

When users talk about computers, they usually describe the interfaces - because, for most users, the interface is the system. 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, although both are visual/motor-based controls. And switch panels were clearly more haptically satisfying! Now, thanks to exponential increases in commonly available computer power and versatility (and concomitant cost decreases), significant progress in interface modalities and their affordability can be perceived.

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. Computational power seems to be catching up with algorithmic understanding.

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 affordability have moved "closer to the user:" speed, stop, reverse. We're tracking a similar evolution in human-computer interaction space. Perhaps interfaces are finally growing up?

Hiroshi Ishii

interfaces

Tangible Interfaces
People have developed sophisticated skills for sensing and manipulating their physical environments. However, most of these skills are not employed by traditional graphical user interfaces (GUIs). Tangible Bits, our vision of human-computer interaction, seeks to build upon these skills by giving physical form to digital information, seamlessly coupling the dual worlds of bits and atoms.

Guided by the Tangible Bits vision, we are designing "tangible user interfaces," which employ physical objects, surfaces, and spaces as tangible embodiments of digital information. These include foreground interactions with graspable objects and augmented surfaces that exploit the human senses of touch and kinesthesia. We are also exploring background information displays that use "ambient media:" ambient light, sound, airflow, and water movement. Here, we seek to communicate digitally mediated senses of activity and presence at the periphery of human awareness.

Natural

The goal is to change the "painted bits" of GUIs to "tangible bits," taking advantage of the richness of multimodal human senses and skills developed through our lifetime of interaction with the physical world.

The musicBottles project presents a tangible interface for interaction with a musical composition. The core concept is that of using glass bottles as containers and controls for digital information. The bottles represent the three performers (violin, cello, and piano) in a classical music trio. Moving and uncorking of the bottles controls the different sound tracks and the patterns of colored light that are rear-projected onto the table's translucent surface.

Caleb Chung

Interfaces as Pets
Computer-human interface (moving from computers to humans) is difficult to bring off. Better to go the other way: start with humans. Begin with what people want around them: "nurturing," "fun." Remember that 80 percent of communication is non-verbal. Don't try to make interfaces friendly. Instead, start with friendly things and make them smart!

Imagine a personal digital assistant (PDA) that acts and reacts like a "virtual pet." It has attitude, character – qualities and cues that make you want to interact with it. It has an interesting personality. It has its own agenda. It can become a friend.

Remember when you were six? Your imagination brought your simplest toys to life, and the world around you was limitless. Those interfaces were driven by the human imagination. Open thinking let you make intuitive leaps to invention, expanded your imagination. You were free to create "on top of" your toys, following the most basic human/animal cues.

Toys! Toys have "user friendly" down pat. Toys teach minimalism in physical design. You can't use expensive parts, four circuit boards, etc. The best toys support and encourage imagination-driven open-ended play – true intelligence!

The best interfaces are transparent, unobtrusively observing humans and responding to their needs. The model "personal assistant" is the valet of 18th century British culture. Just tell it what you want done, without concern for its feelings, but always with a sense of play. You needn't be precise. And no one has time to learn to speak alien languages.

Send the message that something is alive, and we'll attribute intelligence to it. This is the true natural interface. And it's not that difficult, if we but try. Today's amazingly powerful computers can't even tell if you're there with them. Let them observe what humans naturally do, then do that!

Clark Dodsworth

Universal Studios' expansion project, *Islands of Adventure,* recently opened in Orlando. It uses roughly two orders of magnitude more digital infrastructure than the original – a step toward putting ubiquitous digital sensor and effector intelligence into the entire built environment, indoors and outdoors. The guiding notion is that a theme park should be aware of everyone who enters, learn a few facts about them, and then provide a customized user-experience. Every bit of the park, including the landscaping and robotic fauna, should behave or respond interestingly, engagingly to you, and then react with tailored nuance to the next person or family. That notion is not unique to the theme park industry; it's in the strategic plans of the consumer electronics industry, the toy business, the automobile business, and it's important to a few alert individuals in the computer industry.

The task is to create intelligent devices and environments that are designed to adapt to humans and augment the human experience, rather than ones designed to be easily manufactured and then adapted to by humans. In industry, the driving force is competition: parity products need to differentiate themselves. In the computer industry, which holds the biggest rewards for such adaptive interfaces and human-centered design, that driving force is largely quiescent. As intelligence and the software behind it migrate to common objects, the computing world has far more to learn than to teach.

Bill Buxton

If the user is conscious of using a computer, that is a strong indicator of a design failure. Another way of looking at this is to ask: Of the total number of brain cycles expended in performing a task, what percentage are consumed on operational issues compared to content-specific ones? If it is greater than about five percent, we most likely have a failure of design.

It borders on banal to state that we live in an ever-more-complex world, and much of that complexity is due to the previous generation of technology. It seems equally obvious that the basic litmus test of future designs should be: Does it enhance our ability to cope with that complexity? I view well-designed technology as a cognitive (and often social) prosthesis. It is a means to render tractable problems that would otherwise be overwhelming.

From a design perspective, there has been literally no progress since 1982 in the computers used by the majority of the population. And we still live in a climate where it is acceptable for over 90 percent of computer science students to graduate without ever writing a program that is used by another individual, much less be graded on their ability to do so.

Well, the 1980s are over. And the status quo in design and education is just as dated as the music of the Bee Gees, sideburns, and bell bottom trousers. We look back on them with a bit of quaint nostalgia, coupled with horror that we ever found them acceptable. It is time to grow up.

Invisible

Moderator
Christian Rouet
Lucas Digital Ltd.
cr@lucasdigital.com

Panelists
Keith Goldfarb
Rhythm & Hues Studios

Ed Leonard
DreamWorks SKG

Darwyn Peachey
Pixar Animation Studios

Ken Pearce
PDI

Enrique Santos
Digital Domain

Paul Yanover
Disney Feature Animation

Research and Development for Film Production

Visual effects and animation film studios are always starving for a new kind of moving imagery that has never been seen before, that, ironically, will inevitably be quickly surpassed. The rate of change, the high quality that is required by the film medium, the reasonably long pre-production phase of most feature films, and the public's insatiable need to see complex visual imagery contribute to research and development of "cutting-edge" tools and techniques that are often ahead of any other fields in the entertainment industries.

At the same time, production is rolling like a well-oiled engine, or at least, that's the hope. Not only must new tools be innovative while advancing the state of the art, but they must also be "reliable" enough to be used in production with careful management of the risk caused by their introduction in the production pipeline.

However, these amazing visuals will soon become visual commodities that need to be produced in massive quantities. So our tools must be more "efficient." Thanks to the geometric progression of digital technology, and unlike other non-digital fields, algorithms that used to be too slow or too complex to handle are rapidly becoming attainable. But how can the R&D effort be organized to face these challenges?

Christian Rouet

At ILM, most of the R&D effort is concentrated in a central department that provides services to the various features in production. Every show is an opportunity to improve and validate some aspects of our digital production pipeline in a way that provides long-term benefits to the company. The philosophy of the R&D department is to push the limits of the technology in two ways: by focusing internally on issues that don't have commercially available solutions, and by developing strategic relationships with key vendors to help them push their products to satisfy our needs. The R&D department almost never develops small shot-specific custom code. Instead, it provides an open architecture to technical directors that allows them to customize our applications or architecture with appropriate plug-ins or scripting capabilities, while guaranteeing stability and long-term evolution.

There are three major types of R&D projects:

1. Visually driven projects are usually scheduled during the pre-production phase of a feature film project in order to provide the innovative tools needed to invent new visuals. Examples of such projects include our skinning system, facial animation system, digital fur system, and natural-effect renderer.
2. Efficiency-driven projects are scheduled in a similar way, but with the goal of improving an already-known technique to make it significantly more efficient and produce more volume faster and cheaper. Based on production experience with earlier versions and the increased volume that is required by the new generation, it is often necessary to redesign and entirely rewrite an existing tool in order to break major efficiency bottlenecks.
3. Architectural projects have global side effects on most applications, like implementing or supporting new APIs, file formats, interactive frameworks, or computer system infrastructures. The strategy in their implementation is to anticipate and design the long-term architecture in R&D, break down the work into mid-sized tasks, and interleave execution of these (horizontal) architecture milestones with (vertical) production-driven development.

Software applications are infinitely more malleable than digital imagery. What an algorithm can do to an image is only limited by the programmer's understanding and creativity. The engineer who is motivated by the invention of new imagery can have as much of an impact on the final content as the concept or digital artists themselves. In fact, the synergy of having both the engineer and the artist brainstorm about an idea produces results that widely surpass the sum of their individual contributions.

Film

Keith Goldfarb

Rhythm & Hues believes:

Play slow, win slow; play fast, lose fast.
In-house software is essential. It is economical, flexible, effective, and educational. Purchased software is none of these things.

The simplest move is the best move.
Most productions do not need special-case tools.

Always remember: keep the balance.
Special-case software is, unfortunately, sometimes necessary.
When it is, try to fit it into the big picture.

With only one group, you will win.
Software design and development is a group endeavor.
The users need to be very involved in the software process.

There is a thin line between thick and slow.
Software needs to be designed, but not at a snail's pace.
Get the software into production as soon as possible.

Only amateurs try to come up with "fancy" moves.
Although it is fun to come up with new solutions to old problems, it is neither efficient nor fair to do so.

Rhythm & Hues Studios is well known for its work on the Coca-Cola "Bears" commercial campaign.

Ed Leonard

Development and integration of new software techniques has become an integral part of the creative process of making a feature-length animated film. Each new film strives to push the visual boundaries beyond those set by previous films. DreamWorks had the unique opportunity to design an organization and working structure for our software group from the ground up. Partnering with technology vendors, funding external research, and developing an in-house toolset were all part of the evolution of the DreamWorks software plan.

Ultimately, the success of the DreamWorks solution relies upon a blending of production-driven creative development and solid engineering practices. Our software teams are organized around the functional components representing the Core, Pipeline, Studio Tools, and Production Software groups. A key part of our overall software strategy is to rotate developers through the groups to maximize the experience and talent base. New creative techniques are driven by production need, and a phased development approach takes advantage of various developer skill sets at the different project phases.

Creative software R&D is driven toward creating compelling visual results during a visual development software phase. The Core, Pipeline, and Studio Tools groups help engineer these creative techniques into solid technology solutions for the movie pipeline.

Darwyn Peachey

At Pixar, we primarily use in-house proprietary software for animation production. We generally decide to "make it" rather than "buy it" because that gives us important flexibility in responding to really tough R&D challenges posed by each new film project. We've been fortunate to attract a world-class R&D team combining top researchers with fine software engineers (often in the same person), so we are in a good position to innovate. In some cases, the commercially available software products are so good and so well suited to our work that we do choose to buy rather than build. Alias AutoStudio is used for creating geometry in many of our character and set models, and Interactive Effects' Amazon paint system is used for painting textures and backgrounds. Both companies are responsive to our suggestions, which makes it possible to have some of the benefits of in-house software without the associated costs.

But in most of our modeling, animation, shading, lighting, and rendering, we resort to in-house technology so that we can completely control the results and devise new and better ways to solve problems.

We follow a fairly centralized model of R&D, in which an R&D department outside of production attempts to respond to current production priorities in all projects. At the same time, the R&D department is responsible for cross-project and long-term planning. We try to solve particular production problems by extending and revising our software in ways that are generally useful. We try to avoid introducing software changes that are "hacks" specific to a single film. Research is usually, though not always, motivated by our requirement for a particular film. During preproduction, we analyze the needs of the project and plan R&D activities to address them. This is made easier by the long life cycle of a feature film: roughly three years from inception to completion of production. However, the R&D department typically has to balance the needs of three or four film projects that are at different stages in their life cycles.

Ken Pearce

In-house software development has been fundamental to PDI since the company was founded in 1980, when the only way to create computer animation was to program it yourself. Nineteen years later, although our mechanisms and disciplines for software development have changed, we still use home-grown solutions for the vast majority of work that we do. Most of us like proprietary software for the level of control and efficiency it provides and we generally need some compelling reasons to buy off the shelf.

We believe long-term software developers need to be close enough to production to understand, collaborate, and deliver appropriate solutions, but not so close that frenetic schedules and demands for features make it impossible to design and engineer for the future. At PDI we try to maintain this balance not through a heavily managed development process but rather through direct, collaborative relationships between programmers and animators.

Enrique Santos

At Digital Domain, software development has been aimed at integrating the best commercially available solutions with in-house tools that give the technical edge to our 3D artists. The challenge has always been how to use limited resources to create solutions that meet tight production schedules.

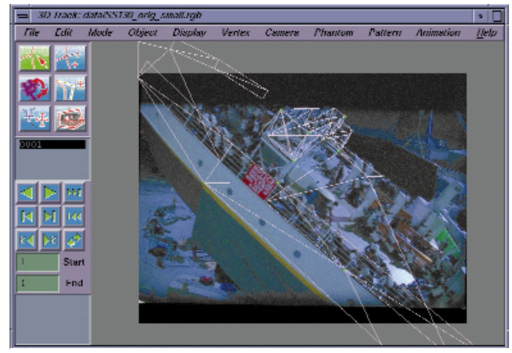

Screen from Track, Digital Domain 3D tracking tool awarded a Scientific and Technical Academy Award in 1999.

We believe that in-house tools should provide features that we don't see being developed into commercial packages, provide a technological advantage, and increase efficiency and productivity. The choice of buy vs. build is not always made on cost considerations. At times, the need for flexibility and control over the application supersedes the desire to buy off-the-shelf solutions.

Software development at Digital Domain falls into two categories: infrastructure and specialty tools. Infrastructure includes tools such as 2D and 3D format converters and editors. Specialty tools include 3D tracking, cloth, fur, and natural phenomena simulation.

Paul Yanover

At Disney Feature Animation, we aim to create technology that facilitates and enables filmmaking. At its heart, our challenge is the inherent discontinuity between research and development driven by the specific needs identified in shots on a particular film and creation of new tools that are applied across multiple films. Our goal is to satisfy both of these and nourish a culture that accepts the dynamic tension that this situation creates. At its best, we are successful in driving innovation that is both short-term and long-term depending on the source. Production-driven solutions often provide the underlying prototypes that in turn become long-term tools and departmental standards. Alternatively, we have developed long-term tools at the outset with a deep understanding both of our needs and available technology.

Our philosophy remains one that is dedicated to applying technology to the screen and pushing the use of available technologies to their practical limit. We believe strongly in sound software development and work hard to bring our systems to the hands of the many, many classically trained and technically naive artists who comprise some of our greatest departments. That means spending the time and money to build software that really works over the long haul and can be supported and used by large groups of non-technical artists. In this sense we operate very seriously like an in-house software company rather than just hacking things together to get out shots.

Visualizing Large-Scale Datasets: Challenges and Opportunities

Organizer
Kwan-Liu Ma
University of California, Davis

Moderator
John Van Rosendale
US Department of Energy

Panelists
Stephen Eick
Visual Insights/Lucent Technologies

Bernd Hamann
University of California, Davis

Philip Heermann
Sandia National Laboratory

Christopher Johnson
University of Utah

Mike Krogh
Computational Engineering International Inc.

While we are seeing an unprecedented growth in the amount of data from both computational simulations and instrument/sensor sources, our ability to manipulate, explore, and understand large datasets is growing less rapidly. Visualization transforms raw data into vivid 2D or 3D images that help scientists reveal important features and trends in the data, convey ideas, and communicate their findings. However, massive data volumes create new challenges and makes previous visualization approaches impractical. The new generation of visualization methods must scale well with the growing data volumes and cope with other parts of the data analysis pipeline, such as storage devices and display devices.

To accelerate the development of new data manipulation and visualization methods for massive datasets, the National Science Foundation and the US Department of Energy have sponsored a series of workshops on relevant topics. The workshops have generated a concept, Data and Visualization Corridors, that represents a combination of innovations on data handling, representations, telepresence, and visualization. In the next few years, we expect to see more human and financial resources invested to solve the problem of visualizing large-scale datasets as more demanding applications emerge.

In this panel, the findings and results of the workshop on Large-Scale Visualization and Data Management (Salt Lake City, Utah USA) are reported. The panel and the audience collaborate to answer the following questions:

- How large is large?
- Where do large datasets come from?
- Can current graphics and visualization technology cope with the volume and complexity of data produced by tera-scale calculations or high-resolution/high-volume data collection devices?
- How much of the data do we need to see, and how do we find what we need to see?
- What are ideal data representations that can enable more efficient visualization?
- How much processing power, storage space, bandwidth, and display resolution do we need?
- How much visualization computing should we do at runtime, when the data are being created, vs. at postprocessing time?
- Is computational steering a reality?
- Are there common visualization solutions for scientific, engineering, medical, and business data?
- What can the visualization software industry offer now and in the near future?

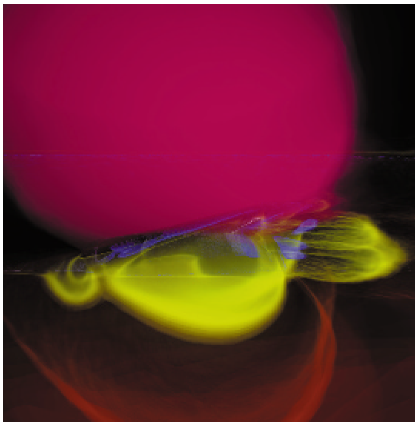

Volume visualization of data from high-lift analysis revealing flow structures surrounding an aircraft wing. The data contains over 18 million tetrahedra, and the visualization was generated using a parallel computer. Image provided by Kwan-Liu Ma, University of California, Davis.

Stephen Eick

The amount of data collected and stored electronically is doubling every three years. With the widespread deployment of DBMS systems, penetration of networks, and adaptation of standard data interface protocols, the data access problems are being solved. The newly emerging problem is how to make sense of all this information. The essential problem is that the data volumes are overwhelming existing analysis tools. Our approach to solving this problem involves computer graphics. Exploiting newly available PC graphics capabilities, our visualization technology:

1. Provides better, more effective, data presentation.

2. Shows significantly more information on each screen.

3. Includes visual analysis capability.

Visualization approaches, such as ours, have significant value for problems involving change, high dimensionality, and scale. In this visualized space, the insights gained enable decisions to be made faster and more accurately.

Datasets

Bernd Hamann

We are now reaching the limits of interactive visualization of large-scale datasets. This is to be interpreted in two ways. First, the sheer amount of data to be analyzed is overwhelming, and researchers do not have enough time available to "browse" and visually inspect an extremely high-resolution dataset. Second, the rendering-resolution capabilities of current rendering and projection devices are too "coarse" to visually capture the important, small-scale features in which a researcher is interested.

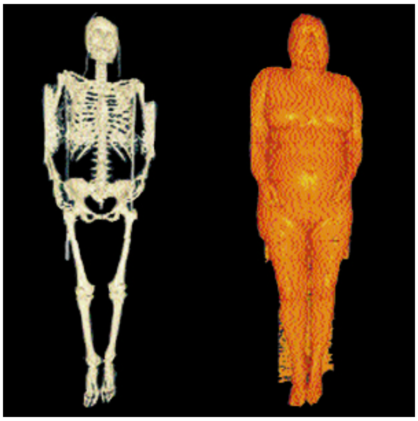

Ray tracings of the bone and skin isosurfaces of the Visible Woman. Image provided by Christopher Johnson, University of Utah.

Two active areas of research can help in this context: multiresolution methods used to represent and visualize large datasets at multiple levels of resolution, and automatic methods for extracting features defined a priori, and identifying regions characterized by "unusual" behavior. Multiresolution methods help in reducing the amount of time it takes a researcher to "browse" the domain over which a physical phenomenon has been measured or simulated, while automatic feature extraction methods assist in steering the visualization process to those regions in space where a certain interesting or unusual behavior has been identified.

Using a full system design approach, visualization requirements are compared with technology trends to quantify the visualization system requirements. Researchers are exploring data reduction and selection, parallel data streaming, and run-time visualization techniques. The system performance goals require each technique to consider a balanced combination of hardware and software.

In summary, multiresolution and automatic feature extraction methods both serve the same purpose. They reduce the amount of time required to visually inspect a large dataset. We should investigate in more depth the synergy that exists between these two approaches. For example, one could envision a coupling of these two methodologies by applying feature extraction methods to the various levels in a pre-computed multiresolution data hierarchy, which would lead naturally to extraction and representation of qualitatively relevant information at multiple scales.

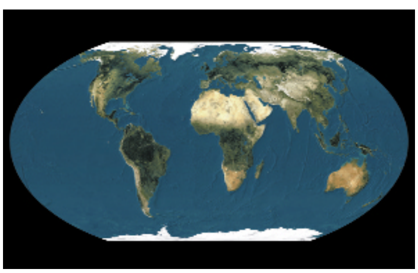

Earth surface image mapped onto a 9-km digital elevation model and rendered (in parallel) at 3000x2400 resolution using a Wagner IV equi-arial projection. Image provided by Tom Crockett, ICASE.

Philip Heermann

The push for 100 teraflops, by the US Department of Energy's Accelerated Strategic Computing Initiative (ASCI) Program, has driven researchers to consider new paradigms for scientific visualization. ASCI's goal of physics calculations running on 100 teraflops computers by 2004 generates demands that severely challenge current and future visualization software and hardware. To address the challenge, researchers at Lawrence Livermore, Los Alamos, and Sandia National Laboratories are investigating new techniques for exploring the massive data.

The leap forward in computing technology has impacted all aspects of visualizing simulation results. The datasets produced by ASCI machines can greatly overwhelm common networks and storage systems. Data file formats, networks, processing software, and rendering software and hardware must be improved. A systems engineering approach is necessary to achieve improved performance. The common approach of improving a single component or algorithm can actually decrease performance of the overall system.

Challenges

Christopher Johnson

Interaction with complex, multidimensional data is now recognized as a critical analysis component in many areas, including computational fluid dynamics, computational combustion, and computational mechanics. The new generation of massively parallel computers will have speeds measured in teraflops and will handle dataset sizes measured in terabytes to petabytes. Although these machines offer enormous potential for solving very large-scale realistic modeling, simulation, and optimization problems, their effectiveness will hinge upon the ability of human experts to interact with their computations and extract useful information. Since humans interact most naturally in a 3D world, and since much of the data in important computational problems have a fundamental 3D spatial component, I believe the greatest potential for this human/machine partnership will come through the use of 3D interactive technologies.

Within the Center for Scientific Computing and Imaging at the University of Utah, we have developed a problem-solving environment for steering large-scale simulations with integrated interactive visualization called SCIRun. SCIRun is a scientific programming environment that allows interactive construction, debugging, and steering of large-scale scientific computations. SCIRun can be envisioned as a "computational workbench," in which a scientist can design and modify simulations interactively via a dataflow programming model. It enables scientists to modify geometric models and interactively change numerical parameters and boundary conditions, as well as to modify the level of mesh adaptation needed for an accurate numerical solution. As opposed to the typical offline simulation mode, in which the scientist manually sets input parameters, computes results, visualizes the results via a separate visualization package, then starts again at the beginning, SCIRun "closes the loop" and allows interactive steering of the design, computation, and visualization phases of a simulation.

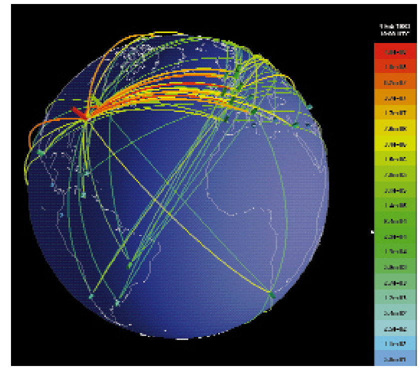

One frame from an animation showing world-wide internet traffic over a two-hour period. Image provided by Stephen Eick, Visual Insights/Lucent Technologies,

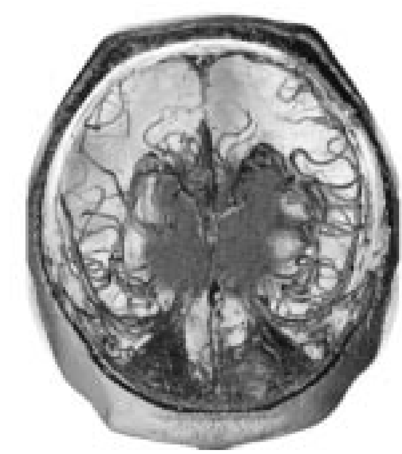

Volume visualization of cerebral arteries in the brain to isolate a large aneurism on the right side of the image. Image provided by Christopher Johnson, University of Utah.

Mike Krogh

The DOE's ASCI program brought about the recent advent of terascale supercomputers. These machines, and the calculations that they perform, are magnitudes larger than what is typically found in industry. A reasonable question is: Is commercial visualization software a viable option for large data visualization?

Visualization was identified in the landmark 1987 NSF report "Visualization in Scientific Computing" as essential for analysis, understanding, and communication of scientific data. In a second report issued in 1998 by the DOE and NSF ("Data and Visualization Corridors"), visualization was once again identified as an essential technology without which "we risk flying blind if visualization does not keep pace with simulations."

Visualization is still correctly identified as essential to computational science. However, even after 12 years, the 1998 report indicates that further research is required for scientists to understand the data they are generating. So in this context, how does a commercial package measure up to the demands of terascale computing? In particular, is such software a viable option for supercomputer users and their management? What features do users want? What, if anything, do they have to be willing to sacrifice? For the software provider, what hurdles must be faced? Where are the overlaps between mainstream and bleeding-edge requirements? And where does the real research lie?

Opportunities

Moderator
Wes Bethel
R3vis Corporation
NERSC/LBNL
wbethel@r3vis.com

Panelists
Carl Bass
Autodesk, Inc.

Sharon Rose Clay
SGI

Brian Hook
id Software, Inc.

Michael T. Jones
Intrinsic Graphics Inc.

Henry Sowizral
Sun Microsystems, Inc.

Andries van Dam
Brown University

Scene Graph APIs: Wired or Tired?

Following widespread adoption and use, the scene graph model has proven to be a popular and powerful development tool because it enables rapid creation of portable and efficient graphics applications. Unfortunately, not all applications fit within the boundaries imposed by a scene graph model.

A class of software tools based on a scene graph model has assumed an increasingly visible position in the set of resources available for developers. The basic model implemented by scene graph APIs is that a "scene" is built incrementally using an API, then scene-centric operations, such as rendering and picking, are implemented within the underlying infrastructure. From a developer's point of view, the scene graph model encapsulates and hides numerous complexities of, and specificities of, implementation-dependent resources, such as texture management and the underlying graphics hardware. With this stratification of resource management, it is possible to achieve both portable and efficient rendering applications.

The purpose of this panel is to examine, in a non-partisan, technical fashion, issues related to scene graph technology. Panelists define the scene graph model in general terms, identify the high-level goals this technology is designed to address, and then address a more specific set of issues from their respective areas of expertise. The scope of these issues ranges from architectural considerations to applicability for a given class or type of application and alternatives to the scene graph model.

Andries van Dam

What is Scene Graph and When is it Appropriate?
A scene graph (aka hierarchical display file) is a retained special-purpose data structure, typically extracted from an application data structure/database, from which the graphics system can efficiently render/refresh a scene and do first-level interaction handling. Such a "retained mode" is particularly advantageous if the scene doesn't change considerably between consecutive frames, which is the case, for example, for walk-throughs and fly-overs of largely unchanging scenes, and if certain kinds of performance optimizations are done to drive the underlying immediate mode graphics (hardware) pipeline efficiently. Recent scene graphs have generalized the more traditional ones (for example, PHIGS) with advances in such areas as viewing models, geometry compression, rendering capabilities, large-model optimization, and multi-processor optimization.

Sharon Rose Clay

Using Scene Graphs to Allow Differentiated Hardware to Produce a Differentiated Application
SGI has been using scene graph technology for 10 years, through three generations of differentiated graphics architectures, and across a diverse product line, to abstract away hardware details and underlying optimizations, allowing applications to get maximum benefit from underlying hardware without being specifically customized to that hardware. While a scene graph may take many shapes and forms, its quintessential property is a level of abstraction for the data the structure holds. This level of abstraction can contain semantic information about usage intent and structural information that can then be used to optimize various operations on that data for a particular type of solution.

As SGIs' sophistication regarding scene graphs has matured, so has our architectural approach and the problems that we are trying to solve. Originally, our focus was primarily on enabling software developers and customers to get the most from their differentiated desktop or high-end hardware in a particular application area, such as real-time visual simulation, highly dynamic scene modeling, and image processing. Building on past success and lessons learned in both scene graph applications and establishing OpenGL as an industry standard graphics API, we are now focused on developing increasingly sophisticated scene graph architectures with a broader level of standardization than ever before.

Henry Sowizral

Viewing Scene Graphs As Source Code
Scene graphs work well for scene composition: for placing, orienting, grouping, and procedurally manipulating objects within a scene. This domain-oriented approach makes the application programmer's tasks much simpler.

Unfortunately, this same domain orientation very frequently results in scene graphs that are poorly structured from a rendering perspective. Such poorly structured graphs require the underlying software system to perform significant optimizations. A scene graph's structure can impose subtle ordering constraints. These constraints can preclude possible parallelization or content reordering that could provide better rendering performance. Moreover, because scene graphs tend not to have a strong spatial coherence, spatially-based operations such as picking, collision detection, and culling become rather inefficient.

If we view scene graphs as "source code," then a "compilation" can reorder and optimize the scene description into a far more efficient form, possibly even generating system-specific code that exploits underlying hardware resources such as multiple processors and render pipelines. A compilation can localize and appropriately schedule state-change operations and enable concurrency. By building ancillary data structures with better spatial coherence, the compilation can also accelerate other operations such as picking and collision detection. Careful semantic design of a scene graph API can simplify the programming task and enable efficient rendering.

Wes Bethel

Scientific Visualization and Scene "Grafting"
Scene graph models and scientific visualization systems can harmoniously coexist, despite a push for optimizations that contribute to dilution of flexibility and extensibility. Taking a step back, we can reference the similarity between traditional dataflow visualization systems and scene graph models. The dataflow model describes a visualization "program." A scene graph model describes the composition and layout of a collection of geometry and a viewpoint in space. The "root node" of a dataflow program is the data "sink," which is usually a renderer. The "root node" of a scene graph is owned by "the system," with "data" pushed out to the leaves of the scene graph. In both, there is a well-defined and (a mostly) deterministic traversal path. Both systems hide complexities from the user, sometimes at the expense of extensibility and flexibility.

As a visualization person, I must be concerned with data modeling and dynamically changing the contents of the scene graph. For example, as users navigate through data, the subset of data visible to them changes. Using "straight scene graph," since "all the data" is owned by the system, render-time view frustum culling can be used to manage spatial complexity by breaking a big problem into a set of smaller problems. Alternatively, user code can be "grafted" with the scene graph model to provide data based upon the current view. The former method is feasible only when "all the data" fit into memory at render time. The latter method represents a harmonious blending of technology, potentially at the expense of real-time frame rates. Such blending of appli-

Tired

cation and system resources represents one approach to implementing large-model visualization, coupling simulation with visualization, and computational steering. Herein lies the conflict between the scene graph system's inclination to optimize and the developer's need for flexibility.

Scene graph is interesting as a framework for visualization due to its fundamental properties: defined traversal/execution mode and optimized rendering. Visualization clearly benefits from the fastest rendering possible, but at the same time, it requires a flexible and extensible framework. Given the disparate needs of a wide range of developers and users, one point of view is that there is no "one true scene graph design" that will satisfy the needs of the graphics and visualization community, and that there will always be "outliers" that bend and warp the scene graph metaphor beyond a "container for rendering."

Carl Bass

Every application worthy of that moniker, has a complex object model, but one that is probably not optimized for graphics display. In my experience, there are several characteristics that determine whether or not scene graph display technology is appropriate for a particular use. Is this a new application, or are you retrofitting an existing application? Is there enough memory to allow for the duplication of data that is associated with scene graphs? How well does the scene graph accommodate editing operations and data changes common to your application? What hardware functionality is being obscured by the scene graph's attempt to be easy to use? How difficult is it to replace a particular scene graph when better technology enters the market?

Michael T. Jones

Innovation Above and Beyond the Scene Graph
Scene graphs complement low-level hardware-abstraction APIs by providing compound structures and common graphics utilities to application developers. They can add convenience, efficiency, portability, and structure to programming projects and ease the use of differentiated hardware benefits that have created a following among application developers. These virtues, admirable as they are, cannot hide the fact that scene graphs are merely the second rung on a ladder that is too short to reach the goal of continued innovation.

The graphics toolmaker's quest to fuel advancement is no longer about discovering how to define cameras, traverse spatial data structures, provide concurrency and synchronization, or allow for extensibility and customization. Nor can the path be one of repetition and refinement. Restating past successes in new languages or programming paradigms is treading water, not swimming. Viewing the development landscape from the peak of successful scene graph systems shows that there is a further, and much greater, mountain to climb.

The character of this next level of graphics development tools derives from the intrinsic structure of graphical applications and the context of their deployment. Designing tools with this view illuminates inherent flaws in the scene graph approach. Solutions to these flaws exist only at a higher level of abstraction: the graphical application platform. This framework approach leverages all that is good in scene graphs while avoiding their weaknesses and adds a third rung to the graphics development ladder to create a powerful development environment that succeeds where scene graphs alone must fail.

Brian Hook

Is Scene Graph Technology Good for Game Developers?
Top-tier consumer 3D action games compete on many different fronts, including content, performance, gameplay, and striking visual displays. Unfortunately, high performance and truly stunning graphics are often mutually exclusive and require compromises and techniques that may not be readily apparent to scene graph API authors or applicable to a wide range of applications. In our particular genre (first-person 3D action games), performance needs dictate that we must pay heavy attention to the issues of occlusion, visibility determination, and level of detail, while at the same time the need for leading-edge graphics requires that we present realistic surface textures, geometry, and lighting, most likely using techniques that have never been implemented.

A general-purpose solution to a specific problem can not compete with a specific solution to that specific problem, assuming all else is equal (talent, hardware, and resources). Applications can make compromises of performance vs. quality vs. generality that are far more fine-grained and appropriate than similar decisions made by a general-purpose library. While scene graph APIs can allow applications to arbitrarily extend their features into unknown areas, doing so (via subclassing or callbacks, for example) subverts many of the claimed benefits of a scene graph API. This begs the question: If I'm spending the majority of my time going through back doors because the API doesn't have the features I need, what is it offering me in the first place?

Get Real! Global Illumination for Film, Broadcast, and Game Production

Moderator
Stuart Feldman
Discreet Logic
stuartf@discreet.com

Panelists
Craig Barron
Matte World Digital

Scott Lelieur
StudioDVP/Design Visualization Partners, Inc.

George Murphy
Industrial Light & Magic

Dave Walvoord
Blue Sky | VIFX

Experienced industry experts who are on the front lines of achieving realism in production review the pros and cons of the various approaches to attaining photo-realism. One of the main themes is the use of global illumination rendering techniques to achieve realism. This is an issue that is taking on new prominence for digital content creators as new rendering applications incorporating global illumination algorithms and increasingly more powerful computers arrive on the production scene.

Stuart Feldman

Global illumination provides the ability to create highly realistic and natural-looking lighting effects without the need to spend an inordinate amount of time trying to attain these effects with software acrobatics. Global illumination can accommodate a more natural interface for lighting a scene because it is based on the true physics of lighting (photometry). Production and lighting designers can intuitively light a scene in a virtual set as they would in a real one. Working with photometry also facilitates duplication of lighting conditions in situations that involve compositing live-action shots into virtual environments.

In spite of these benefits, adoption and use of global illumination software in production has been slow in coming. Global illumination solutions are often perceived as being too difficult to use in the production workflow and/or too slow computationally for practical application. The panel presents a number of first-hand experiences with global illumination algorithms in production and discusses the need and practicality of using other approaches to attain realism.

Craig Barron

Working in the field of realistic special effects in the entertainment industry for nearly two decades, I've overseen the vast changes in visual effects technology. Sometimes, it's nearly impossible to simulate realistic landscapes and environments with traditional rendering techniques that have to be intercut with shots of photographed reality.

Matte World Digital's digital artists often make custom texture maps to simulate realistic environmental phenomena, like weathering and other textures. The texture maps copy the undescribed lighting effects when traditional rendering is lacking in realism.

The ever-increasing need for a heightened degree of realism in creating our CG environments led my company to use radiosity rendering for the first time in 1995 to create various establishing shots of an old-style Las Vegas for the Martin Scorsese film *Casino.* Since then, this technique and other global illumination schemes have become an important tool for simulating reality on the computer when seamlessly blending live-action footage with computer-generated imagery.

Close-up detail of radiosity mesh on Tangiers sign.

From the film *Casino* – live-action set, shot in Las Vegas.

The final composite completes the scene with a 3D environment.

Get Real! Global Illumination for Film, Broadcast, and Game Production

George Murphy

While advances in CG lighting models are currently capable of creating very accurate simulations of real-world lighting effects, the price for this imagery is heavy processing and slow turnaround. In the harsh reality of limited budgets, tight deadlines, and highly burdened rendering resources, facilities such as ILM continue to use a variety of cheats and tricks – techniques that are more efficient at simulating realistic, real-world lighting effects while offering the appearance of more accurate CG lighting models.

These visual cheats have been at the heart of our ability to solve the CG challenges we've had to face in a commercial production environment and draw on a hybrid of 2D and 3D digital solutions. However, as the CG components of scenes have begun to comprise more of what actually makes up a scene, it has become more and more difficult to believably fake real lighting models. The nature of our CG cheats has become more sophisticated as we rely more heavily on a body of experience and skill.

The demands of modern visual effects, along with a more sophisticated viewing audience, have required the images we create to be more accurate than ever in their simulation of reality. Still, the computationally intensive nature of radiosity and even standard ray tracers – at least for now – keeps us developing our toolkit of CG cheats and illusions as a viable solution for creating the appearance of reality.

Dave Walvoord

The goal of *Bunny,* a seven-minute animated film, was to create a film with a completely original style. The director's vision was for a warm, almost photorealistic look, but one rooted in the foundations of a storybook style. A consequence of this vision was the progressive development of Blue Sky's proprietary renderer, CGI Studio, which made achieving the vision possible.

By using radiosity, a global illumination rendering technique that mimics the most subtle properties of natural light, the artists at Blue Sky were able to create a dimensionality and organic realism that was facilitated by computer graphics technology, instead of being dictated by it.

While using radiosity brought a natural complexity and warmth that would have not otherwise been possible, it presented new issues not normally encountered in computer graphics production. There was an instant conflict between physics and art. The physics of natural light did not always agree with what we as artists had in mind for the scene. We found ourselves using many of the same tricks cinematographers use on live-action sets, but with different expenses inherent to the world of computers and the algorithms we were using.

Scott Lelieur

Most of the discussion regarding realism in computer graphics revolves around the question of using process-intensive and accurate rendering algorithms such as radiosity and ray tracing vs. using less accurate techniques that are more efficient. This is a trade-off of quality vs. efficiency. It presupposes that the production will support frame-by-frame rendering. There are, however, types of productions that cannot support frame-by-frame rendering, especially in the gaming and television broadcast industry.

First-person games require the environment to be rendered as the game is played, usually at a rate of 20-30 frames per second. With a few notable exceptions, television programming is produced in a timeframe that is too short to support many visual effects. As a result, television production has increasingly begun to rely on real-time graphics techniques to solve visual effects problems for live and live-to-tape styles of programming.

At StudioDVP/Design Visualization Partners, Inc., we have been creating real-time visual effects for television for the past four years. Generally known as virtual sets, real-time production for television requires the same level of quality for the viewing audience as traditional types of effects, but it has to be done without the benefit of post-production. Global illumination algorithms like radiosity offer some distinct advantages for real-time production. They allow us to create extremely convincing effects by pre-computing the lighting in the scene.

illumination

Organizer
Lucy Petrovich
University of Arizona
lucy@u.arizona.edu

Panelists
Maurice Benayoun
Z-A Production

Tammy Knipp
Florida Atlantic University

Thomas Lehner
Stadtwerkstatt

Christa Sommerer
ATR Media Integration &
Communications Laboratories

Experiential Computer Art

Interactive computer art has a rich history, beginning in the 1950s with cybernetic sculptures. These large-scale, computer-controlled installations involving sound, light, and responsive sensors were developed with financial support and technical assistance from major corporations and research centers. The Renaissance of the 1960s brought together artists, engineers, and scientists in organizations like Experiments in Art and Technology to create an onslaught of electronic art, including interactive work. Artists such as Rauschenberg, Wen-Ying Tsai, and Seawright all participated.

Slowly and steadily, this new art form evolved unbeknownst to the general public and barely known to the art community in most cases. Within the last decade, interactive computer art has matured. It is now at the forefront of contemporary art. This artwork manifests itself as a dynamic process of experience instead of an aesthetic object, hence the title of this panel, "Experiential Computer Art."

Lucy Petrovich

I am working with interactive narrative from a feminist perspective, exploring women at the mid point of their lives, their reflections and aspirations. As the women reflect on their lives, so too the program is designed to reflect on the viewer's actions. This becomes a responsive system that tracks the viewer's path and begins a dialogue between the viewer and the program. The tracking is twofold: creating and nurturing (properties often attributed to women). By exploring the interaction, viewers create a unique path through the data, and if they choose, the program helps viewers explore areas that remain undiscovered.

Women of a Certain Age (1998), interactive narrative, Lucy Petrovich, University of Arizona

Maurice Benayoun

For many years, I have been exploring the potential of virtual technologies as tools for metaphorical architectures of communication. After having tried, in the early nineties, with the Quarxs, a 3D computer animation series, to say how the truth can escape even science, I thought that virtual reality was the perfect tool for interrogating art and the audience. The "big questions" (starting with *Is God Flat?*) were, for me, the opportunity to make a statement about the relationship between space, spectator individuality, and meaning. If architecture is a space to live, virtual reality can be a space to read. In *Is God Flat?* the audience explores an infinite world of bricks. Participants can build their own labyrinths in search of the flat, one-dimensional images of God(s) from art history.

Is the Devil Curved? was the second of these "big questions" in which, by digging corridors in the sky, we meet a living creature who attempts to seduce us, casting sounds of physical pleasure and changing his (her?) shape and behavior in order to gain a wider audience. So-called intelligent agents and artificial life were used to provide a virtual experience of full potential and not pre-written story.

World Skin (1997), Maurice Benayoun, Jean-Baptiste Barrière, Ars Electronica/SGI/Z-A

In 1995, the *Tunnel Under The Atlantic* allowed people from the Centre Georges Pompidou (Paris) and the Museum of Contemporary Art (Montreal) to meet each other by digging tunnels into pictures of their common past. Golden Nica Interactive Art Ars Electronica 98, *World Skin*, "a photo-safari in the land of war," is a CAVE work where the audience removes parts of memory, taking photos in the virtual world. They leave the virtual land of war with the printed shots, material witnesses of the lost memory.

Tammy Knipp

Is interactivity changing the content of the arts? "New media" challenges us to rethink and re-examine our basic presumptions about art, content, interaction, and our perceptions of reality. With the age of electronic media, "interactivity" has come to be refashioned. The prefix "inter" references "between, among, and within;" the word "activity" is defined as "energetic action or movement;" and the term "interact" means "to act on each other." In the context of new media, interactivity may be defined as the energetic action or movement between multiple realities or within a mediated environment.

There are three described realities outlined by Max Velmans, author and reader in psychology: the physical reality, the psychological reality, and the virtual reality. My interest lies with the energetic action or movement among the perceived boundaries that frame these multiple realities and the psycho-social interaction resulting from the encompassing levels of mediated experiences.

In my work (*CASE STUDIES*), performance-like installations integrate the reality of the physical, social, psychological, and virtual worlds. Each *CASE STUDY* provides an analytical view whereby participants and viewers become subjects from an observational view. It is from this perspective that the psycho-social interactive dialogue and kinetic languages can be observed, thus supplying the final constituent element of the artwork. The multilevel interaction among realities, viewers, and participants reflects content and, perhaps, provides an element of meaning with enticing, critical, and provocative thought.

Thomas Lehner

During the weeks of 28 August-13 September 1998, a large-scale communication sculpture controlled via the Internet (*ClickScape 98*) transformed the Danube area near the Nibelungen Bridge into a clickable public space.

Electronic visitors transmitted their visual, audible, and text messages online as interventions in the urban landscape of Linz. Every mouse click had an effect in real space. Real-time spectators experienced live how the lights of the EA Generali Building were switched on and off, heard acoustic greetings cross the Nibelungen Bridge, and saw messages from cyberspace scroll across the front of the Stadtwerkstatt in illuminated letters.

ClickScape 98, Thomas Lehner, Stadtwerkstatt

Christa Sommerer

We are artists working on the creation of interactive computer installations that use the audience's participation as essential input for creation of image structures within the systems. Instead of traditional object-oriented artworks, we aim for a new type of image generation that uses genetic programming, mutation, selection, and evolutionary image processes to create process-oriented artworks.

These artworks are no longer pre-fixed and pre-programmed by the artists but instead are non-deterministic, non-linear, and multi-layered. Visitors to such an installation play a significant role in the development of the artwork. Only if they agree to become part of the system will they understand that there are no pre-defined solutions to be found within the artwork but that instead they, through their interactions, essentially determine what they see. Visitors create their own artworks, which are essentially reflections of their inner expressions, expectations, and interactions.

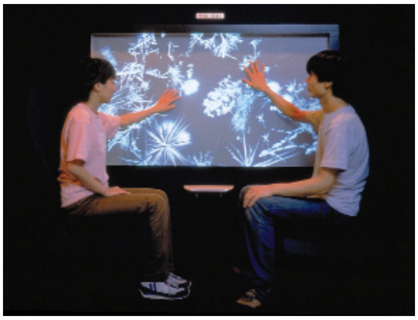

HAZE Express, interactive computer installation, Christa Sommerer. Developed for IAMAS International Academy of Media Arts and Sciences, Gifu, Japan.

Moderator
George Suhayda
Sony Pictures Imageworks
geo@spimageworks.com

Panelists
William T. Douthitt
National Geographic

Syd Mead
Syd Mead, Inc.

Rob Minkoff
Director

Jay Redd
Sony Pictures Imageworks

Stuart Sumida
California State University, San Bernardino

Bill Westenhofer
Rhythm & Hues Studios

Syd Mead

William T. Douthitt

Visual Effects: Incredible Effects vs. Credible Science

Can visual effects balance scientific truthfulness, art, and storytelling?

Scientific research, exploration, and journalism have given the public access to an extraordinary amount of information and imagery. As the work of directors and visual effects artists is flashed onto screens 25 feet tall and 60 feet wide, and viewed by billions, every nuance can be compared to the information and images at the audience's fingertips. The minute our audience says, "I don't believe it," we've lost our credibility and their attention.

As a result, directors are holding visual effects supervisors to an even higher standard. The fantastic gains in software and imagery quickly become old news. Digital artists and software developers scramble to deliver a better product.

How should we educate our digital artists to meet the challenge? Should science drive the art of storytelling and film? Is it okay to break the rules of science for creative purposes?

The quality of our images has also raised an ethical question in a world drenched with media. The line between journalism and entertainment is blurred. What seems real is not always truthful. Since the audience cannot always distinguish between "realism" and truth, do we have an obligation to be real and truthful in our images? Can the work of directors and visual effects supervisors educate and entertain?

Syd Mead

"Credible science" is a moving target of best guesses reinforced with observable facts and fashionable assumption. Populist reality is perception-dependent on individual factors of cerebral constitution, the external senses combined with emotional and cultural influences. Perception is fluid and as amorphous as the visual illusions that the special effects industry creates. The "reality" that we assume as our base reference is actually quite unstable and has been proven to actually contradict direct observation. Add in enough Prozac, and lemon drops may start to taste blue.

Entertainment, from Greek amphitheater to the latest state-of-the-art RGB projection cinema, has always been challenged by the gap between believability and pretense, made fragile by visual "style" and theatrical fashion.

The conflict between special effects and "credible science" is a matter of sensory choice, based on individual resident memory that constitutes "reality." Today, with the populist audience getting dumber by the hour, "credible" becomes dependent on a broad-based "last-seen-and-remembered" artifact. We have gone beyond the belief threshold that used to assure us that "seeing is believing." As we pursue ever more "incredible effects," the "credible science" flees before us into unexplored dimensions. The challenge is not to compete with the reality of "credible science," but with the capacity of the popular audience to match what we say is "real" with what they think reality is. "Credible science" has, in my opinion, very much less to do with it.

William T. Douthitt

National Geographic strives to present an accurate and honest portrayal of the world. The covenant we keep with our readers is to show them what actually takes place.

In the early 1980s, we acquired cutting-edge technology to help with pre-press operations. Enthused with the potential of those tools, editors of that era made a modest shift in the content of a picture, to better fit our cover format. The decision caused criticism and passed into publishing folklore ("When National Geographic moved the pyramids."). Yet the compelling features of digital manipulation tools can't be denied, any more than our photographers could be denied use of color film.

When we employ digital techniques, we carefully acknowledge them. The key question for us is how digital manipulation serves the story while maintaining honesty with our audience.

The digitally created effects in the film industry are awe-inspiring. Enthusiasm for their potential is evidenced by the speed with which the industry has accepted these new tools. Yet in many of the films, it often feels that the key elements (compelling story and honesty with the audience) are lost in the desire for spectacle.

The missions of a reality-based print publication, and a fiction-based film industry are very different. Yet both share the common need to provide a believable, "real" experience, through compelling storytelling and honest communication with the audience.

George Suhayda

Rob Minkoff

Story has led the director in his work. We can't dismiss science. It supports story, and sometimes suggests it.

It's important for directors to look carefully at the scientific research so they can make the most informed choice possible. As simple as *Stuart Little* may seem, a lot of time was spent finessing him into a believable, animatable character. Hours of instruction in anatomy, movement, fur, and cloth went into creating him. Without all this preparation, our character would not have the impact he does, and this insures his employment in future features.

Education and observation of nature can't be stressed enough, but it has to begin with the director. Only then will the crew understand that the director will hold them to a higher standard. A simplistic approach to filmmaking doesn't hold up with the advent of CGI. Overwhelming schedules and release dates make it a challenge to pay attention to detail. The alternatives are not as well-thought-out stories and characters.

Rob Minkoff

Stuart Sumida

Faithfulness to reality and the need to entertain are not mutually exclusive, as reality is often far stranger and more compelling than fiction. So we need to know the reality of what we model digitally.

There is an enormous gulf between available funding for anatomists and other scientists who are pursuing computing and research and the resources available to those who use technology for entertainment and imaging. And there is far less primary data than is popularly thought.

The answer to these challenges is cooperation, and that depends on scientists' natural desire to be scientifically literal and digital artists' willingness to use scientific models that may be both practical and entertaining. However, universities and museums must acknowledge that collaboration with corporate entertainment is as professionally valid as the old "publish-or-perish" mentality; and studios must recognize scientists as a useful pool of talent. Progress will accelerate if scientists become more than "hit-and-run" consultants and become deeply involved in aspects of story, writing, and design. The benefits of cooperation will be healthier academic and entertainment industries, more educational materials available for a new generation of artistic and scientific students, and positive public relations.

Stuart Sumida

Effects

Bill Westenhofer

Bill Westenhofer

Reality. An interesting subject in a profession that portrays the unreal, surreal, or too real to have been simply filmed in camera.

Humans have an innate ability to extrapolate things that are familiar onto the unknown. This willingness to extrapolate gives us room for the artistic license we often need. No one has ever seen an animal talk, yet we are familiar with how humans talk and how animals move their mouths, and it is possible to extrapolate one onto the other in a believable way. As part of the storytelling process, visual effects supervisors need to balance reality with imagery that conveys the visual and emotional impact required. We should rely on the real world wherever possible, but less out of a need to be faithful to science and more because reality itself can be interesting in its complexity. This requires us to be experts in the extrapolation process.

Education is the key. Simple observation is not enough to determine why something "doesn't look right." Understanding laws that govern what we see enables us to make better judgments on how to apply our "artistic license." We can then create superior effects that build on the natural world in a logical yet visually interesting way.

Jay Redd

Visual effects is both an art and a science. It is the melding of these two approaches that challenges us to go beyond just bringing to life that which doesn't exist in the real world. It pushes us to continually define and redefine what the "real world" is and is not.

Re-creating that which is "familiar" or "real" often leaves the audience uninterested and bored. Conversely, creating that which is "fantastical" or "unreal" often leaves the audience unbelieving, dismissive.

Our responsibility is to define and create reality itself – certainly in the context of a feature film. As the storytelling process becomes less and less inhibited by technological and creative limits, our ability to create believable, convincing worlds becomes greater and greater.

The coming together of artists and technicians yields perhaps the world's most collaborative art form – the feature film. Like a brush to a canvas, CG to a film is just another tool. How we use that tool is a combination of art and science, education and experience, technical accuracy and aesthetic prowess.

The game here is who can produce the most convincing reality, fully utilizing both science and artistry. Your eye is your camera. Your brain is your film.

Jay Redd

How SIGGRAPH Research is Utilized in Games

Moderator
Chris Hecker
definition six inc.

Organizers
Alex Dunne
Game Developer Magazine

Alan Yu
Game Developers Conference

Panelists
Seamus Blackley
Microsoft Corporation

Peter Lincroft
Ansible Software Inc.

Casey Muratori
Gas Powered Games

Michael "Saxs" Persson
Shiny Entertainment

Game developers are always looking for methods to improve performance and the visual appeal of their titles. To that end, the results of research presented at the SIGGRAPH conference each year are closely analyzed by developers of interactive games for their applicability. In this panel, leading game developers will explain what advances in game graphics they have made using SIGGRAPH research, and discuss promising graphics technologies that game developers should keep an eye on.

Chris Hecker, definition six inc.

For most of the game industry's lifespan, game developers have simply struggled to get anything at all onto the screen for the player to interact with. These images usually consisted of simple bitmap graphics, or some completely hacked pseudo-3D images. The academic community studied the game industry during this time, rather than vice versa, and this attention usually came in the form of user interface analysis and psychological studies on players. These days, game developers routinely mine the wealth of academic research for ideas, algorithms, and even implementations of solutions to problems in graphics, physics, and artificial intelligence. With consumer-level desktop computers rivaling workstations for performance, and perhaps more importantly, with the game industry making billions of dollars every year, academics have also taken notice of the game industry, and have started producing research that's more applicable to the kind of real time interactive problems that comprise a game. This positive feedback loop means today's games are incorporating today's research more and more frequently, and in some cases are even suggesting problems for researchers to solve in the future. The panel will discuss some of the past uses of academic research, and look towards the future.

Peter Lincroft, Ansible Software, Inc.

I have studied the SIGGRAPH proceedings every year since I was introduced to them in 1988. I remember using an old SIGGRAPH paper for reference when devising my BSP-based data structure for the 3D models in X-Wing in 1989. Another paper inspired me to develop a span buffer technique which I used to speed up my first "high-resolution" game, TIE Fighter. Yet another was the basis for my ordered dithering technique.

While I have learned some useful techniques from SIGGRAPH, the biggest benefit I get from attending SIGGRAPH, and reading the proceedings, is the energy boost. To see all of the exciting work being done in the field, gets me really excited about what techniques we might be able to bring to the PC gaming world in the future. Curved surfaces, fractals, wavelets, physical modeling, radiosity, particle systems, noise functions, genetic algorithms; all of these seemingly esoteric topics are relevant to computer gaming applications.

Historically, there has always been a gap of years between when something is presented in SIGGRAPH, and when it is first used in a computer game. TIE Fighter was one of the first computer games to use Gouraud shading, a technique which was nearly 20 years old at the time. Within a year or two, games were using techniques similar to Phong to capture specular highlights, narrowing the gap to about 15 years. This gap is narrowing exponentially, and will soon be closed. In my opinion, the R&D work being done by computer games programmers will soon be in step with, or even in advance of, the research presented at SIGGRAPH.

Research

Seamus Blackley, Microsoft Corporation

Using SIGGRAPH Papers in the Game World
I used to know very few people in the game industry who had read a SIGGRAPH paper. These days, however, it is very common for an engineer who is faced with a new problem to consult the SIGGRAPH literature before he does anything else. This is as much a result of the improved capability of game platform hardware as it is a shift in the focus of the literature. While researchers are becoming increasingly interested in the types of problems faced in real-time environments, so also is consumer hardware becoming capable of exploiting an increasingly large portion of the published work.

Over the course of my career in the games industry, I have seen the adoption of many techniques, all novel in the game space, arise from the implementation and optimization of published SIGGRAPH work. These include, among others, terrain rendering, large-model simplification, surface specification, IK, dynamics, and lighting. This has saved me and other developers a bunch of time and money, and made possible some leaps in technology that many thought could not be made. It is clear that the next few years will see a continuingly strong and productive confluence between the research we see in the proceedings, and the experience people have in their living rooms. I can't wait!

Casey Muratori, Gas Powered Games

The scope of computer game technology has grown significantly over the past decade. Whereas the game developer of yesteryear concerned themselves with low-level hardware trickery, the game developer of today must tackle far less concrete problems. Game developers are quickly finding out that the research community has already tackled several problems common to game programming, and offered up a variety of solutions. The application of these solutions, unfortunately, comes with its own set of problems. Developers may find that applicable research algorithms are too time intensive for use in a real-time environment. They may also find that producing a robust implementation is too costly for a production schedule. Finally, they may find that the algorithms they implement have not been thoroughly tested under the types of circumstances common to real-time interactive environments. Although there have been many successes, it is clear that the game industry still has a long way to go before simple ad hoc solutions disappear and research-oriented implementations are the norm. This will likely take a considerable shift in thinking by both programmers and managers alike.

Michael "Saxs" Persson, Shiny Entertainment

With computers becoming faster as we speak, it is increasingly important for game programmers to be aware of methods for using all that extra juice. What everybody wants for their money is increased realism, in the form of better visual quality, increased natural motion, and lifelike physics. The annual SIGGRAPH conference is an indispensible source of algorithms and the best source I know for the latest trends in visual computing.

More than adopting the techniques I find there, I have used the SIGGRAPH Conference Proceedings to inspire and energize my own research into new algorithms. I have specialized in scalable character systems for the last three years, and the SIGGRAPH conference has provided a tremendous amount of information on how to get those characters to come alive with better motion control systems, which I will be finishing after my current project, Messiah.

It used to be that most SIGGRAPH papers really weren't practical to implement into games due to limited computational power, but now, with a little creativity, a lot can be implemented and added to the gaming experience.

Visual Storytelling

Moderator
Graham Walters
Pixar Animation Studios
graham@pixar.com

Panelists
Sharon Calahan
Bill Cone
Ewan Johnson
Tia Kratter
Glenn McQueen
Bob Pauley
Pixar Animation Studios

When cinematic storytelling is at its best, the visual imagery furthers the narrative by establishing the world in which the story takes place and setting the emotional tone as each sequence unfolds. While this is a well understood practice in traditional filmmaking, the language of digital filmmaking is just now being developed. For digital filmmaking to move beyond imitation of traditional filmmaking and become an art form unto itself, the medium's aesthetic development will need to be given the same weight as development of the technology.

Using *A Bug's Life* as a case study, the panelists examine this issue by presenting the creative goals that drove the project and discussing the creative and technical directions they pursued to accomplish these goals.

The panel has two goals:

1. To expose the audience to the design process that goes on behind the scenes on a digital film to create visual imagery that best supports the storytelling goals of a film.

2. To initiate a dialogue on how the formal visual language of traditional live-action filmmaking compares to the just-now-developing visual language of digitally made films.

The panelists represent each of the major design disciplines. All have formal art training and have worked in various capacities in the non-digital design industry before eventually gravitating to computer-animated filmmaking. They discuss how visual imagery supports storytelling in traditional media and how they have adapted traditional techniques and devices to work in the digital world.

Storytelling

Moderator
Jennifer Yu
PDI
jen@pdi.com

Panelists
Ken Bielenberg
Apurva Shah
PDI

Jim Hillin
Walt Disney Feature Animation

Eben Ostby
Pixar Animation Studios

Neville Spiteri
Square USA

Function and Form of Visual Effects in Animated Films

Computer-generated visual effects, generally associated with live-action projects, have increasingly become a key component of animated features. Nearly all animated films now have visual effects: crowds, explosions, and natural phenomena such as water and fire.

The advent of computer-generated feature films, and the increased visual complexity of traditionally animated films, have driven the use of computer-generated visual effects in this medium.

This panel focuses on the forms of visual effects (what are they and how are they developed?) and their function in animated films. Panelists with experience on live-action, traditionally animated, and computer-animated films discuss strategy, challenges, definitions, issues, and creation of visual effects in animated films.

Jim Hillin

Digital effects are becoming a more and more prominent part of film production, whether live-action or animated. Otherwise live-action films have digital characters, props, and settings in addition to the classic explosions and laser beams. Deciding what is and isn't an effect isn't as interesting as exploring how live action, animation, effects, and filmmaking in general are converging. Practically speaking, an "effect" is a visual element that you must add outside of the normal process of shooting the film (or of animating it).

Because making an effect adds work and cost to the production in a highly visible manner, you must take extra care to be sure that it's visually necessary for the story you're telling. A large part of deciding this necessity comes in the overall design of the production. A production's effects strategy, then, is often settled quite early and in accordance with the artistic direction and production design. Because changes to that design can necessitate changes to the effects, the effects strategy must remain integrated with the whole production strategy, more than many other areas.

It is in effects design, perhaps, that the greatest differences between live action and animation become apparent. In live-action digital effects, the visual design must essentially match the real world, or rather, the footage that's already been shot. The challenge lies almost entirely in achieving that match and that believable level of realism. Animated effects, on the other hand, must be designed from scratch, and the degree of stylization or realism must be decided and approved before the shot can even enter production. Animated effects can, therefore, require more overall work than live-action effects. Otherwise, live-action and animated effects share the same processes and techniques. Although they require separate consideration in the planning and design stages, when it comes to doing the actual work there's no difference between the two. Once again, the more interesting question is this: With the rapid growth of digital effects in all kinds of film, where do we draw the line between live action and animation?

Function

Eben Ostby

At Pixar, we approach effects production as an adjunct to the design-model-shade-layout-animate-light-render production pipeline. While most of the production can be compartmentalized into these processes, certain operations don't fit into this scenario. An effect is any such element, and its creation usually involves several departments or processes. Very commonly, effects use procedural animation, simulation, and unusual application of lights, shaders, or renderer functions.

They almost always start with the question: How do we do this...? Effects provide an emotional boost to the story. They emphasize key moments in the film, amplify story points, or add to the visuals. Design direction for effects comes both early, in the design phase of the film, and late, in editorial. What's unusual about effects in animated films is that the look of the objects involved are rarely considered separately from their motion. The visual qualities of surface and motion are bound together.

Effects are an important and distinct part of the making of an animated film. They are a fascinating area to investigate both because of their unusual visual character and because of the multidisciplinary approach that is used to implement them.

Grain pour for *A Bug's Life.* This image shows part of an effects shot that simulated grain pouring from a bottle. Layers of grain flowed from the bottle and growth of the pile was regulated by animated depth maps. The effects work was done by Christian Hoffman. Image above © Disney/Pixar.

Apurva Shah

When considering effects in computer-animated films, people often ask: "Isn't the whole thing an effect?" This is, naturally, not the case. At PDI, we took a process-oriented approach to defining what an effect is in *Antz*. In essence, an effect is anything that doesn't fit into the normal production pipeline of modeling, layout, motion, and lighting.

Another way to consider this is through the function that effects play. The more obvious effects in animated films add visual flash or thematic emphasis, much as in live action. This class of effects includes explosions, fire, water, and other natural phenomena, and even digitally created crowds. A more subtle class of effects that comes for free in live action, such as dust and stuff in the air, must be explicitly created in animated films.

Beyond the issue of defining an effect is that of designing one. Although *Antz* is graphically stylized, for example, it is physically realistic. This results in effects with the same level of complexity as you would get in a live-action film. However, there are differences in the development process itself. In a computer-animated film, the effects department is intimately involved from the earliest stages of visual development. They can provide input in every decision that affects the integration of the effects into the film. Visually, the design for the effects proceeds in lockstep with the production design of the film as a whole. Photorealism, typically the obvious choice for live-action effects, is not necessarily the goal for animated effects.

Effects, then, whether obvious or subtle, ideally help tell a story and hold it together visually. In computer-animated films, all effects must be explicitly created and designed in a style consistent with the production design of the whole film.

Form

Neville Spiteri

The scope of effects for a computer-generated movie depends, of course, on the style of the film. Effects need to seamlessly integrate with the rest of the movie, help emphasize a certain story point, and generate a certain mood. They create a visual effect in the true sense of the term. Consequently, the effects work is a primary factor in establishing the style, look, and level of realism of the animated film.

In live-action movies, the effects typically need to be as photorealistic as possible so that they are perceived by the audience as events that have been shot through the director's lens. This realism has that specific quality that comes from the film camera, a look that we are so familiar with and fond of. In computer-generated films, the level of realism or unrealism is an artistic choice. The effects help define this level and artistic choice.

The starting elements for effects in computer-animated films differ from those for effects in live action. In a computer-animated film, you work with the digital information of the entire 3D cast and set, not just the scanned live-action plates and motion-controlled camera moves. However, the underlying processes and techniques employed to generate effects are similar to live-action effects, and the greater the photorealism of the movie, the more similar they become.

Recent animated films include effects that clearly complement the chosen level of realism and artistic style. In the future, we will see movies that convey different levels of realism from photorealistic to highly stylized. If the computer-generated characters and environments are closer to photorealism, then the effects must be more photo-realistic. In this case, the goal of the effects work is the same as effects work for live-action films. Effects for computer-generated movies are particularly exciting and challenging because you are engaged in a process of creating and conveying to the audience a new and different level of augmented realism.

Water effects shot from the flood sequence in *ANTZ* (©PDI & Dreamworks, LLC)

Digital Watermarking: What Will it Do for Me? And What it Won't!

Moderator
Jian Zhao
Fraunhofer CRCG
MediaSec Technologies LLC
jzhao@crcg.edu

Panelists
Eckhard Koch
MediaSec Technologies LLC

Joe O'Ruanaidh
Siemens Corporate Research

Minerva M. Yeung
Intel Corporation

Digital watermarking has emerged as an enabling technology for protecting digital intellectual property rights. The technology is completing its trial phase and becoming mature in diverse markets, from electronic commerce and digital broadcasting to DVD. Several cross-industry organizations are developing standards that will support wide deployment of this technology.

Digital watermarking embeds secure invisible or inaudible labels in multimedia data (such as images, audio, text, video, 3D graphics) for identifying copyright-related information such as origin, ownership, use-control, integrity, or destinations. The digital watermark is integrated with the multimedia and tightly bound with the quality of the content. The benefits of digital watermarking for content protection are twofold: it provides evidence of illicit copying after the event, and it discourages such misuse in advance.

Like any other emerging technology, digital watermarking has raised a number of questions both in research and business: Why is digital watermarking necessary? How does the technology work in the real world? What are the strengths and weaknesses of the technology? What data can be watermarked? How do you measure the security and robustness of digital watermarks?

This panel provides an opportunity to hear from some researchers as well as some startup entrepreneurs in the digital watermarking area. Panelists explain how to apply this technology in real-world applications and discuss where digital watermarking is likely to go in the future.

Eckhard Koch

Even though digital watermarking is still in its infancy and needs more research, the technology is already in commercial use. The main reason for this early adoption is that customers can't wait. Their bitter experiences with data pirating make them eager for protection.

Most customers for this technology are OEMs, system integrators, hardware manufacturers, and software vendors. They could add considerable value to their products and systems by licensing watermarking technology. The straightforward model for the watermark business is royalty-based licensing. An alternative model for large-scale mass-market products is to provide counter-based, off-the-shelf watermark packages. End users are charged by the number of watermarks locked in a software or hardware counter.

The major technical challenge is to develop a foolproof protection scheme while at the same time keeping the watermarks imperceptible. Absolute robustness is impossible, but digital watermarking can be useful as long as the process of tampering with and removing the watermark is costly and time-consuming.

Potential customers are pushing the digital watermark technology to enter the market faster than other new technologies. In fact, visionaries have adopted digital watermark products in some application domains where state-of-the-art technologies provide sufficient protection. For example, in the field of secure and invisible communications, where steganography has been accepted for centuries, watermark technology is considered ready for real-world applications.

Digital

Digital Watermarking: What Will It Do for Me? And What it Won't!

Several cross-industry organizations such as CPTWG (the Copy Protection Technical Working Group) and DAVIC (the Digital Audio-Visual Council) are working hard to develop digital watermarking standards that will allow wide deployment and acceptance of this technology. Other visible initiatives are underway at the International Federation of Phonographic Industry (IFPI) and the Secure Digital Music Initiative (SDMI) for audio watermarking, and at the Digital Audio-Visual Council (DAVIC) for watermarking in e-commerce.

In the foreseeable future, in my personal opinion, it's unlikely that the digital watermark will be standardized for all markets. Instead, standardization will likely happen in different sectors such as physical media (DVD, CDs), online media delivery, broadcasting, and document authentication in the imaging industry.

Joe O'Ruanaidh

Watermarks. The term evokes visions of shady characters secretly beavering away in dark basements surrounded by forged $100 bills drying on clothes lines.

In a digital media context, away from the traditional world of inks and paper, the same old problem remains but it relates not just to forgery but also to outright theft, because one digital copy can spawn millions of others with the single click of a mouse button. It is hardly surprising that the notion of a digital watermark has stimulated avid interest amongst artists and publishers alike.

It is commonly recognized that digital watermarks must be as robust as the media in which they are embedded. For example, a rotated, cropped, and rescanned watermarked image should still be a watermarked image. However, this robustness requirement directly conflicts with the need for a digital watermark to be unobtrusive. The most effective techniques used to embed watermarks are the result of a combination of secret key-based techniques used for military communication and simple models of the human visual system.

The most familiar application for digital watermarks is for copyright protection and protection of intellectual property. On its own, a watermark does not provide any legal proof of ownership. In other words, the use of a given digital watermark to protect intellectual property must be registered with a trusted third party to be of any value. Any technique for embedding robust digital watermarks must be compatible with methods for registering copyright.

A watermark's resistance to intentional and unintentional degradation has been the main subject of interest in the watermarking community. The main challenges are geometric transformations such as change of proportion or simple rescaling. One watermark removal technique that is supplied on the Internet simply shifts a corner of the image. Lossy image compression such as JPEG and filtering are more easily overcome and, generally speaking, watermarks have evolved into very resistant forms. The results are impressive in the laboratory, but will they really work in the real world?

Watermarking

Minerva M. Yeung

The Internet has been growing very rapidly, and the bandwidth available to users has increased as the Internet has grown. Digital subscriber lines and cable modems are becoming widely available at affordable prices. As transmission rates increase, the quantity and quality of available digital content (in the form of images, audio, video, graphics, and 3D models) will increase. However, this raises a major problem for content providers and owners: protection of their material. They are concerned about copyright protection and other forms of abuse of their digital content. On top of that, digital media content, coupled with Internet distribution, is subject to instantaneous mass replication and distribution, resulting in severe loss of revenues or royalty payments.

Many of those involved in cross-disciplinary research on effective content protection technology for digital media believe that technology can play a major role in providing the infrastructure for content protection and distribution. But they are gradually realizing that a comprehensive digital content-protection infrastructure requires more than data encryption and embraces technical innovations plus in-depth study of the functionality, thread models, limitations, and applications across many disciplines, ranging from cryptography, computer science, signal processing, software and hardware architecture, public policy, and law.

This panel focuses on end-to-end protection of multimedia content in Internet-related applications, the role of digital watermarking in the media content protection infrastructure, potential benefits and limitations, possible attacks and remedies, and the current status of watermarking research in the computer graphics field.

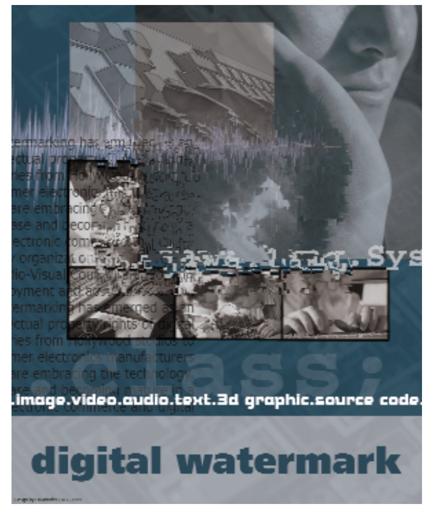

Jian Zhao

High cost and user reluctance to establish the "use-control" trusted model encourage exploration of alternative solutions for content protection. Although the original motivation for digital watermarks was for copyright protection, this technology has found a multitude of potential applications not originally envisioned by pioneers in the field.

Digital watermarks can be used for creation of hidden labels and annotations in medical applications, in cartography, and for multimedia (video and audio) indexing and content-based retrieval applications. In a medical scenario, watermarks might be used for unique identification of patient records. Patient records could be embedded directly into the image data for each patient, which would both speed up access to records and prevent potentially harmful errors such as a mismatching of records and patients.

For proof of authenticity, digital watermarks are of particular interest in electronic commerce and distribution of multimedia content to end users. The surfaces of ID cards, credit cards, and ATM cards could be watermarked. So could bank notes, personal checks, and other bank documents. Scanners, printers, and photocopiers could refuse to operate if they find a watermark that specifies permissions to manipulate the document and does not authorize scanning, printing, or copying.

As a robust covert communication channel, digital watermark technology has a wide range of applications in the defense and intelligence sectors, where traditional steganography has been employed for centuries. Also, digital watermarks may find potential markets in countries where strong cryptography is not permitted.

Moderator
Steven Feiner
Columbia University
feiner@cs.columbia.edu

Panelists
Henry Fuchs
University of North Carolina at Chapel Hill

Takeo Kanade
Carnegie Mellon University

Gudrun Klinker
Technische Universität Muenchen

Paul Milgram
University of Toronto

Hideyuki Tamura
Mixed Reality Systems Laboratory Inc.

Mixed Reality: Where Real and Virtual Worlds Meet

Ever since Ivan Sutherland's development of the first head-tracked, see-through, head-worn graphics display, researchers have been exploring the mixture of real and virtual objects. On one end of the spectrum is the real world itself (seen, heard, and felt without any virtual intervention). On the other end is the fully synthesized virtual world (theoretically a replacement for the real world, experienced through computer displays).

What at first glance are polar opposites, however, are not as far apart as they appear. Increasingly, our experience of the "unenhanced" real world is enriched by the sight and sound of computer displays on the desk, wall, or palm. And the virtual world presented by a current "fully immersive" VR system typically relies on the user standing or sitting on a part of the real world, and often makes substantial use of real-world textures.

This panel addresses some of the many ways in which virtual and real worlds are being combined in user interfaces to create "mixed reality" (MR). Topics range from augmented reality (AR), in which additional material is added to the user's experience of the real environment, to augmented virtuality (AV), in which real material is added to the user's experience of a virtual environment. Panelists situate their work along this "reality-virtuality" continuum.[1]

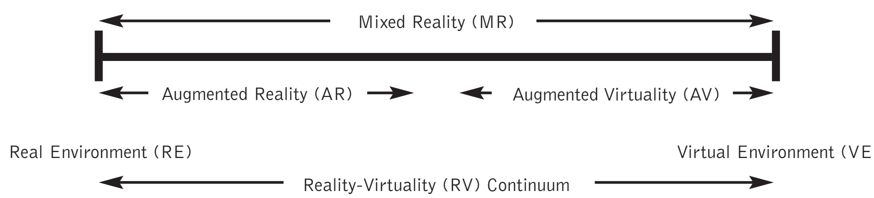

Steven Feiner

At Columbia, we are interested in how augmented reality and wearable computing can be combined. Our ultimate goal is to create a mobile mixed reality that can support ordinary users in their interactions with the world.

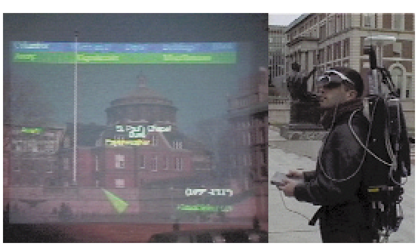

Outdoor campus tour overlays labels on view of surrounding world (left) seen by user wearing see-through head-worn display. (Columbia University)

My discussion stresses two themes:

1. Presenting information about a real environment that is integrated into the 3D space of that environment.

2. Combining multiple display and interaction technologies to take advantage of their complementary capabilities.

We refer to our experimental applications as "hybrid user interfaces," because they synergistically mix different display and interaction technologies, including 2D and 3D, see-through and opaque, monoscopic and stereoscopic, large and small, and mobile and stationary.

One system serves as a personal "touring machine" that assists the user in exploring our campus, overlaying relevant information on objects of interest. As users move about, they are tracked through GPS and inertial trackers, with information presented on see-through head-worn and opaque hand-held displays.

Another system explores collaboration within a shared 3D virtual space. The shared space is a virtual "ether" that envelops users, interaction devices, and displays that range from head-worn to hand-held to wall-sized.

Reference
1. P. Milgram and F. Kishino. *A Taxonomy of Mixed Reality Visual Displays.* IEICE Trans. on Information and Systems, vol. E77-D, no. 12, December, 1994, pp. 1321-1329.

Real

Henry Fuchs

Two of the applications being pursued at UNC Chapel Hill, immersive telecollaboration and "mixed reality" visualization assistance for surgical procedures, impose severe requirements that cannot fully be satisfied by today's technology. For telecollaboration, we wish to interact with distant collaborators as naturally as if they were in the same room with us; for surgical visualization, we wish to see real-time and pre-op imagery as though we possessed Superman's X-ray vision.

Turning such ideas into real systems requires careful balance between imperfect system components. For example, making distant collaborators appear as if they were next to us requires head-mounted displays or large flat-panel displays or high-resolution projectors, each of which is well beyond the state of the art. Enabling a physician to see imagery within a patient requires displays to combine real and virtual imagery to match with high-resolution, high-precision displays that are similarly beyond the state of the art.

For telecollaboration, we are pursuing image presentation via a multitude of overlapping front-surface projectors, calibrated by multiple digital cameras. For surgical assistance, we are developing head-mounted displays augmented by a pair of miniature video cameras, one in front of each eye.

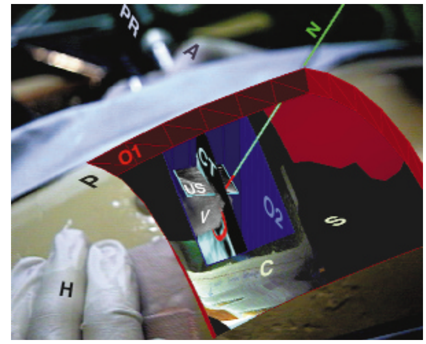

Live video-see-through augmented reality view shown in stereoscopic head-mounted display of experimental UNC system designed to assist surgeons with image-guided interventions. (University of North Carolina at Chapel Hill)

Takeo Kanade

Mixed reality is more than mixing real and synthetic images. It can include modeling, manipulating, altering, and editing reality. In the Virtualized Reality project at Carnegie Mellon University, we have been developing a 4D digitization technique and its facility that we call the "3D Room."

On its walls and ceiling, the 3D Room has 50 cameras that capture real events, such as sports or games, that occur inside the room. From these video sequences, our 4D digitization technique can create a time-varying 3D model of an event, in its entirety and in real time. Once a model is created, we can see the event from any angle for omnipresence, edit the model for event archival, and compute on the model for event simulation. Example output from the 3D Room demonstrates what more can be done with 4D digitization beyond today's image rendering and image overlay.

An input sequence is digitized into a Virtualized Reality 4D event model. (Carnegie Mellon University)

Gudrun Klinker

Current augmented reality research fans out into many different activities that are essential to eventually generating truly immersive AR experience. But the current state of technology cannot yet provide simultaneous support for an optimal solution to all aspects of AR. Today's AR systems have to balance a wealth of trade-offs between striving for high quality and for physically correct presentations on the one hand, and making shortcuts and simplifications on the other hand, in order to achieve real-time response.

In our work, we have selected two different positions among many possible trade-offs, one demonstrating the real-time immersive impression that can be generated with today's technology, and the other forecasting what quality might be achievable with continuously increasing processing power and data bandwidth.

Topics include development trade-offs and applications in design, construction, assembly, and maintenance of machines or buildings, as well as augmented board games such as Augmented Tic Tac Toe.

Augmented reality X-ray view inside a wall for building maintenance personnel. (Courtesy of Fraunhofer Institute for Computer Graphics)

Paul Milgram

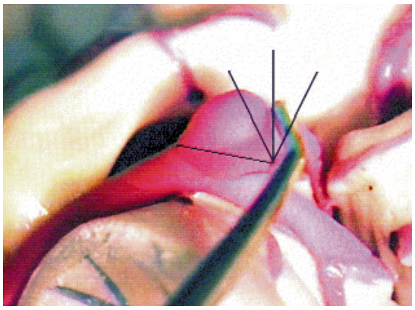

Virtual tape measure applied to intraoperative measurement of aneurysm diameter at the Ergonomics in Teleoperation and Control Laboratory. (University of Toronto)

The fundamental issue in mixed-reality research derives from perceptual confusions that arise when one attempts to align stereographic images with real objects in a real-world, unstructured (and thus unmodeled), stereoscopic video image. Such problems arise, for example, through use of our "virtual tape measure," an AR application that allows one to make 3D measurements of distances and dimensions within remote (unmodeled) stereoscopic video images. An extension of this, the ARTEMIS project, has been demonstrated as an effective means of telemanipulation over low-bandwidth communication channels, such as the Internet, and is currently being expanded for remote mining applications. Current research aims at identifying the fundamental factors that affect perception, and thus alignment accuracy, in such applications.

Our work on remote excavation makes use of a combination of video and laser range images, with a modeled excavator superimposed to facilitate human-mediated control. In this context, our research centers assist the human operator in trading off the relative advantages of exocentric versus egocentric viewing and ego-referenced versus world-referenced control. In a related project, we are applying these concepts to development of display enhancements for endoscopic surgeons.

Hideyuki Tamura

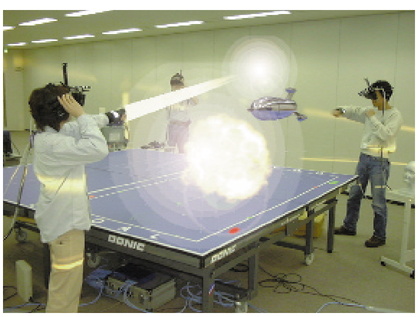

RV-Border Guards: A multi-user AR entertainment system. (Mixed Reality Systems Laboratory Inc.)

Following the concept of MR proposed by Paul Milgram, we started the Key-Technology Research Project on Mixed Reality Systems in Japan in January 1997. By adopting the relatively broad concept of MR, our goal is to develop technology that seamlessly merges the real and virtual worlds.

Different levels of mixture are possible for different application purposes. Thus, merging or fusion of the real and virtual worlds should not be considered as augmentation, which makes one world primary and the other secondary, but rather as a mixture.

The steps toward achieving a seamless MR space and the feasibility of MR technology could be clarified through building pragmatic MR systems. We have been developing several MR systems, such as a collaborative AR system (AR^2 Hockey, installed at SIGGRAPH 98), a multiuser AR game (RV-Border Guards), a visual simulation system with MR (MR Living Room), a cybershopping system (CyberMirage), and an image-based walk-through system (Cybercity Walker). We now truly realize that MR technology can be utilized in a wide range of application fields, including entertainment, education, and industry.

Virtual

Special Sessions ▶

Panelists
Rob Coleman, Ned Gorman,
Christian Rouet, Scott Squires, John Knoll
Industrial Light & Magic

Star Wars Episode 1: The Phantom Menace

The story behind the digital imagery of *Star Wars Episode 1: The Phantom Menace.* Five key members of the Industrial Light & Magic creative team discuss the production process, from conceptual R&D strategies to the final digital render, and present details and behind-the-scenes stories about one of the most anticipated films of all time.

Star Wars

Organizers
Don Brutzman, Naval Postgraduate School
Timothy Childs, Oz

Web 3D RoundUP

Web3D RoundUP is a high-speed shootout where the world's leading Web3D content developers and toolmakers demonstrate the latest Web3D technology and applications in a fast-paced, exciting format for a merciless, cheering audience. Over 30 demo-ees have only 60 seconds to five minutes to win over the audience, all of whom have noisemakers to voice their approval or disapproval. To ensure the event stays on track, the first few rows of attendees are armed and ready with Nerf presents for those who dare to extend their presentations.

Web3D RoundUP's goal is to present the best the Web3D world has to offer in an informative and entertaining format to one of the most technically discriminating audiences. Originally inspired by Rachel and Loren Carpenter's audience participation Electronic Theater debut at SIGGRAPH 91 in Las Vegas, the audience feedback devices are critical to the ongoing success of this event. The interacting audience actually becomes a main part of the event.

To more accurately reflect where the future of the Web3D technology and marketplace is heading, SIGGRAPH 98 broadened the original VRML Demo Special Interest Group, which became Web3D RoundUP. Now in its fourth year (and seventh show), Web3D RoundUP continues to push the visual-computing envelope. It has led the way for the adoption of non-VRML technologies within the Web3D industry by creating a venue where all things Web3D can be shown. Consider the Web3D RoundUP as the perfect instrument for carving out a time-slice showing the state of the art of the Web3D industry.

For more information, RoundUP Results, and/or to plan your next entry, please visit www.web3droundup.org

Web 3D

Organizers
Donald Levy, Sande Scoredos
Sony Pictures Imageworks

A Visit With an Animation Legend

New technology is a powerful instrument in the hands of talented artists and engineers who are free to dream with almost unbridled imagination. As the tools advance, so, too, the possibilities. This inspiring evening brings together the spirit of creativity and the power of insight and intuition that evolves over a lifetime of experience.

While technology has propelled art into the third and fourth dimensions, it is vital to remember the brush strokes behind the key strokes. This very special session is a rare opportunity to meet one of the founding fathers of animation.

Legend

Organizers
Andrew Glassner, Turner Whitted
Microsoft Research

Fiction 2000: Technology, Tradition, and the Essence of Story

Fish need water, songs need a beat, and stories need a medium. What's in a medium? To what extent does a medium influence, or even define, the fictional stories that can be told through it? Some stories work best around a campfire on a dark, spooky evening, while others require costumes, sound effects, or visual effects.

If McLuhan was right, and the medium is the message, then what sorts of stories will we tell each other using the new medium of networked computers? To some people, the Internet and home computer seem to offer an entirely new medium for fiction, unlike any we've seen before. They argue that the classical forms of fiction were limited by technology, and audiences are eager to participate as characters and decision-makers in narrative stories. Others disagree, and argue that traditional fictive forms must be respected as an art form matched to the human psyche, even as they are adapted and changed for the new technologies. The former group dreams of interactive story environments where readers are part of the story; the latter dreams of novel technologically enabled metaphors and genres for tightly-authored works of fiction. "Readers want to help shape the story," argues the first group. "An author's control is essential for a good story," argues the second.

Could they both be right? If so, can these viewpoints possibly be reconciled into one vision? Could the results ever approach the complexity, depth, and power of the modern novel? In this panel, we'll look at both points of view, and several in between, as we imagine the shape of fiction on the Internet over the next 10-15 years.

Fiction