



SIGGRAPH

Conference Abstracts and Applications

Computer Graphics Annual Conference Series, 1999
A Publication of ACM SIGGRAPH



SIGGRAPH

Electronic Schoolhouse
Educators Program | sigKIDS | Community Outreach

Jodi Giroux
The Allen-Stevenson School

Anne Richardson
StarMedia

Jill Smolin
Cinesite Visual Effects

Panels

Jeff Jortner
Sandia National Laboratories

Conference Abstracts and Applications

Emerging Technologies: The Millennium Motel

Kathryn Saunders
Royal Ontario Museum

Sketches & Applications

Richard Kidd
Cinesite Visual Effects

Computer Graphics Annual Conference Series, 1999
A Publication of ACM SIGGRAPH

Conference Abstracts and Applications

COMPUTER GRAPHICS

Annual Conference Series, 1999

The Association for Computing Machinery, Inc.
1515 Broadway
New York, New York 10036 USA

ISBN 1-58113-103-8

ISSN 1098-6138

ACM Order No. 435993

Additional copies may be ordered pre-paid from:

ACM Order Department
P.O. Box 12114
Church Street Station
New York, New York 10257 USA

Or, for information on accepted european currencies
and exchange rates, contact:

ACM European Service Center
108 Cowley Road
Oxford OX4 1JF
United Kingdom
+44.1.865.382338
+44.1.865.381338 fax

Credit card orders from U.S. and Canada:

800.342.6626

Credit card orders from the New York
metropolitan area and outside the U.S.:

+1.212.626.0500

Single copy orders placed by fax:

+1.212.944.1318

Credit card orders may also be placed by mail.

Electronic mail inquiries may be directed to:
orders@acm.org

Please include your ACM member number and
the ACM order number with your order.

Copyright © 1999 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page or initial screen of the document. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM Inc., fax +1.212.869.0481, or permissions@acm.org.

Electronic Schoolhouse

Educators Program | sigKIDS | Community Outreach

4

Panels

114

Emerging Technologies:
The Millennium Motel

160


Sketches & Applications

188

Everyone is an educator and
everyone is a student.

Electronic Schoolhouse:

Educators Program | sigKIDS | Community Outreach



Jodi Giroux

The Allen-Stevenson School

Anne Richardson


StarMedia

Jill Smolin

Cinesite Visual Effects

Classroom **c**
 Playground **p**
 Workshop **w**

Contents



- | | | | |
|----|--|-----|---|
| 6 | Introduction | 68 | Integrating Art and Technology in a Statewide Curriculum c |
| 7 | Committee/Jury/Reviewers | 70 | The Integration of Graphics, Video, Science, and Communication Technologies c |
| 8 | The 4D Virtual Museum of the City of Bologna, Italy c | 71 | The Interactive Learning Environment p w |
| 12 | A Creative Journey c, p | 73 | Introduction to 3D Concepts for Teachers w |
| 14 | Art and Technology: Electronic Resources from the Getty c | 75 | Introduction to Digital Effects c |
| 17 | Art Before Technology or Technology Before Art? That is the Question w | 77 | Math - What's the Use? c |
| 19 | The Atmosphere: Incorporating Interactive Multimedia into the Classroom c | 81 | Math and Computer-Generated Effects: Tools of the Trade c |
| 20 | City@Peace p | 82 | Minotaur: A Tactile Archaeology Game for Kids p |
| 21 | ColorWeb p | 83 | Museums and Computer Games c p |
| 22 | Computer Camp: For Girls Only! c | 84 | Organizing Summer Computer Graphics Camps c |
| 25 | Creating 2D Animation c | 85 | People in the Past: The Ancient Puebloan Farmers of Southwest Colorado p |
| 29 | Creating 3D Animation c | 86 | Proposal Writing 101: Ensuring Your Submission is Understood c |
| 31 | The Creation of the Nu.M.E. Project p w | 90 | The Round Earth Project: Collaborative VR for Elementary School Kids c |
| 33 | Creative Programming: Merging the Artist with the Computer Programmer c p | 94 | The SIGGRAFFITI Wall: Multi-Input Painting p |
| 34 | Developing Creativity: A Curriculum Based on the Use of Computer Graphics Technology c p w | 95 | SP3D and The Lighthouse: Explorations in 3D Internet Learning c p |
| 38 | Digital Design Education at UCLA c p | 97 | Supporting Online Collaborative Communities p w |
| 39 | Drawing & Learning c | 99 | The Teacher's Mid-Life Crisis: Moore's Stairmaster of the Fittest c |
| 41 | Education Delivered Through Storytelling: Using Virtual Reality as an Educational Tool p | 101 | Teaching and Creating Animatics w |
| 42 | Educators Workshop in 3D Computer Graphics w | 102 | ThinkQuest: Students and Teachers Exploring a Global Web-Based Education Project c p |
| 43 | Exploratories: An Educational Strategy for the 21st Century c | 103 | Virtual Harlem c p |
| 46 | FELIX 3D Display p | 104 | Virtual Science Laboratory p |
| 48 | Figures of Speech p | 105 | Visual Effects through Adaptive Technologies w |
| 49 | The Future in Computer Graphics Education c | 106 | Walking the Tightrope: Balancing Digital and Traditional Skills In Undergraduate Education c |
| 51 | Get a Job! A Recruiter Tells You What You Need to Know c | 109 | Web Pages, Interactive Interfaces, and Worm Holes: The Next Generation of User Interface Designers c |
| 54 | Going Farther in Less Time: Responding to Change in Introductory Graphics Courses c | 110 | When Children Draw in 3D c p w |
| 56 | Hands-On Animation w | 111 | Why is the Mona Lisa Smiling? c p w |
| 58 | Hands-On Universe: Teaching Astronomy with Java-Based Image Processing Tools p w | | |
| 59 | High-End Interactive Media in the Museum c | | |
| 63 | How to Marry an Éclair: Anatomy of an Animated Tale p | | |
| 64 | Incorporating Principles and Examples from Art/Design and Film/Video into a CS Computer Graphics Course c | | |



Educator

Everyone is an educator and everyone is a student.

This simple tenet is the foundation and inspiration that is the Electronic Schoolhouse. It was the visualization of this principle that led its three chairs to amalgamate three programs (Educators Program, sigKIDS, Community Outreach) to create a unique wonderland of integrated education, where knowledge sparks in every corner, and inspiration ignites during every conversation.



Education at the end of the 20th century is as multidimensional as it is challenging. The one-room schoolhouse is more evident than ever, yet its dimensions are infinite. We are a fabric of integrated circuits and interrelated topics. Students collaborate with peers halfway around the world, creating projects and establishing a lasting presence where they may not even share a common language. Educators retrieve and use tools from countries they have never visited; they scour library stacks of universities they will never see, exchange ideas with professionals they will never meet, and augment their curricula by perusing classes they will never attend. We learn tools as soon as they're built. We teach those tools as quickly as we learn them. And as soon as we learn them, we encounter students who have already constructed more tools to augment those we have barely learned.

We can choose to be overwhelmed or inspired.

For while we can be easily distracted by the technology, we should be equally compelled to remember our fundamentals, the basics that are hard-coded into every theory and provide the cornerstones of every dazzling technique. As educators and students, artists and scientists, it is vital to remember that enlightenment occurs in the smallest interactions, and transformation from the most unlikely places.



This Electronic Schoolhouse is such a place. We imagined an integrated experience where teachers could present papers in a Classroom and then walk their students through hands-on Workshops on the same topics. We envisioned an environment where students of all ages could learn the process of building a project and then explore the creation on their own in a Playground of interactive installations. We pictured a venue where professionals involved in various computer graphics industries could offer teachers and students interaction they might not be able to otherwise access. We conceived of a Library where a plethora of valuable resources, references, and publications could be borrowed indefinitely, and casual conversation could lead to collaboration and inspiration.

Through our incredible contributors, we believe we have realized our vision and we welcome you to the SIGGRAPH 99 Electronic Schoolhouse.

Student

Co-Chairs

Jodi Giroux
The Allen-Stevenson School

Anne Richardson
StarMedia

Jill Smolin
Cinesite Visual Effects

Electronic Schoolhouse Subcommittee

Laurie Burruss
Pasadena City College

David Cruz
StarMedia

Paul deBonis
Los Angeles Unified School District

Etta Dileo
The Write Approach

Eric Huelsman
Friedman 3D

Bill LaBarge
Student, Rochester Institute of Technology

Joe Lohmar
Digital Domain

Kathleen Milnes
Entertainment Industry Development Corporation

Randii Oliver
Raytheon

Christa Santiago
StarMedia

Loretta Schnurman
Orion Systems Group, Inc.

Evelyn Seubert
Workforce LA

Kristen Stratton
Warner Brothers

Richard Taylor
California State University, Long Beach

Jason Thomas
UCLA

Electronic Schoolhouse Jury

Jodi Giroux
The Allen-Stevenson School

Valerie Miller
SIGGRAPH 2000 Educators Program Chair
Georgia State University

Adele Newton
SIGGRAPH 2000 Community Outreach Chair
Newton Associates

Anne Richardson
StarMedia

Jill Smolin
Cinesite Visual Effects

Electronic Schoolhouse Reviewers

Doug Acheson
Eglin Ayson
Marc Barr, SIGGRAPH 2000 sigKIDS Chair

Debi Barrett-Hayes
Ted Brown

Glenn Campbell
Huguette Chesnais

Tim Comolli
Steve Cunningham
Dena Elisabeth Eber

Frank Gladstone
Scott Grissom

Mk Haley
Tim Harrington

Kristy Higby
Pam Hogarth

Karl Hook
Laurie Howard

Debra Howes
Paras Kaul

Nancy Krebsbach
John MacIntosh

Francis Marchese
James Mohler
J. Michael Moshell

Jacki Morie
Mark Ollila


Jonah Peretti
Maria Roussou

Steve Schain
Evelyn Seubert

Chris Stapleton
Karen Sullivan

Ami Sun
Scott Wilson
Rosalee Wolfe

Committee



Moderator/Organizer
Francesca Bocchi
University of Bologna

Panelists
Maria Elena Bonfigli
Manuela Ghizzoni
Rosa Smurra
University of Bologna

Fernando Lugli
Centro Ricerche S.C.R.L.

Bologna is a major Italian city. It is the largest town and the administrative center of the Emilia-Romagna region, and is located more or less midway between Florence and Venice, in the Po Valley, to the north of the Apennines, the mountain chain that separates the northern part of Italy from the peninsular part, at the northern end of several valleys linking Emilia-Romagna and Tuscany. Today, Bologna has a population of about 400,000, and it has a long historical and cultural tradition. It is known the world over for having the oldest university in Europe (dating back to about 1088) and for its characteristic architectural and urban structure. In the city center, nearly all the streets are built with arcades, known as portici, a continuous series of arches along both sides of the street forming an integral part of the urban structure. Even today, there are more than 25 miles of these arcades in the city center.

It is extremely difficult to construct a museum illustrating the history of the city inside a building, above all when there is a need to show the dynamic phases of history. New developments, periods of decline, new settlements, and changing buildings are all elements that are difficult to present to the public in a satisfactory manner. The New Electronic Museum is using computer graphic techniques to reconstruct the transformations of Bologna, so that the public can see the changes that took place, and to present the results of historical research with a high level of accuracy in terms of cartographic and architectural reconstruction.

Information technology and the information superhighway enable us to move on to a new frontier of research into urban history, making it possible to carry out a three-dimensional electronic reconstruction of the urban habitat and its historical transformations. The result will be a four-dimensional city, with the three spatial dimensions plus the temporal one.

This project combines historical research and new technology to produce results that traditional research methods could never achieve. It is therefore necessary to introduce new approaches to our work, both with regard to the methods used in historical research, and to the transfer of that research to a virtual reality environment. It must be stressed that the method illustrated is not valid solely for the city of Bologna, but rather the city of Bologna is the place where experimentation with this method is taking place. The same method is valid for all cities, whether they be ancient or modern, European, African, Asian, or American.

The Fundamental Role of Historical Research in the Implementation Phase

The first phase consists of historical research based on analysis of documents and identification of the cartographic sources and other useful pictorial sources. This phase is particularly complex with regard to the period (ancient and Medieval) for which no pictorial sources have been handed down to us. Precise visual representations of the buildings and urban structure no longer exist. The only available sources of information for those early periods are the records of the notaries public, government acts, property lists, and tax records.

Even the oldest pictorial records (from the Renaissance onwards) are not always easy to transpose into geometric and virtual forms. Once again, these records must be studied in depth before they can be used. It is necessary to work out why they were produced and what they were intended to illustrate in order to understand the contents and interpret them.

Autocad Modeling

Once the analysis of the historical sources (documentary, narrative, pictorial) has produced satisfactory results, the next step is the modeling phase, implemented by experts in the construction industry (engineers and architects) who have to formulate geometrical models based on the data taken from the documentary sources and combine models and materials used in the various historical periods with the building and urban components that were in existence at that time.

This is a critical phase in the reconstruction work and is preceded by a complex data-processing operation. We intend the model under construction to have a scientific rigor that distinguishes the New Electronic Museum from other products currently available on the market. In order to achieve this level of reliability, it is necessary to start from an accurate two-dimensional map, identifying the precise geographical location (latitude, longitude, altitude) of the buildings.

Four-Dimensional Navigation: the 4D Interface System


The results of the research and the modeling are entered into a virtual environment governing the "time machine."

The Use of Written Sources for Reconstruction of the City in Virtual Reality

When we reconstruct the city in a virtual reality environment, we are not only showing something that is no longer accessible. We are also examining the past by analyzing each part, showing how each element is linked to the others. The sources available to us provide a great deal of information relating directly to physical structures, such as streets, buildings and land, but they also provide much indirect information about various aspects of life in the city.

The records known as the Libri terminorum are, without a doubt, the documents that provide the most significant information about the urban structure in spatial terms. They were drawn up by teams of land surveyors and by officials appointed by the city administration on a number of occasions (1245, 1286, 1294) to lay down a series of ground markers or termini to mark off public areas in the squares, in the streets, and along the city walls from private property. They were laid down in positions that were easily recognizable, such as the corners of buildings and arcades or portici. The Libri terminorum not only provide a list of these ground markers but also a great deal of other information about the buildings. The surveys took into account the upper floors of the buildings, with overhanging beams and guttering: if there were any overhanging features, questions were asked about whether they were built over private ground.

For the purposes of our research project, the careful attention to detail of the land surveyors is of great importance. The precise measurements between one ground marker and the next enables us to see the exact relation between the various elements taken as reference points for the identification of the ground markers. These documents also enable us to locate many buildings on the city map: this degree of resolution is made possible by the care with which the various householders were listed together with their property.



The database designed to manage this documentation needs to show the buildings as the result of social, historical, and property relations over a long period of time. Most of the information in the database concerns people, structures, and places, but the contents are not simply a list of such items. We need to bring together the various components that contribute to urban development so that each person is linked to the others who are in some way related. In the same way, the buildings need to be linked up with their owners and inhabitants. The mass of data put together in this way produces the information that is the basis for the production of hypertexts and for modeling the structures that are recreated in the virtual reality environment of the New Electronic Museum.

The Map of Bologna in Perspective (1575)

For the Holy Year of 1575, Pope Gregory XIII ordered a room in his private apartments in the Vatican to be decorated with a large pictorial map (seven meters across and five meters high, scale 1:360) of Bologna, the city of his birth. It appears that this fresco served a strategic purpose for control and management of the city and the surrounding area. To be able to cast an eye over a pictorial map of the city was a way of possessing it, in material and not just symbolic terms.

This fine pictorial map of Bologna, in which topographic detail is enhanced by landscape painting, consists of a bird's-eye-view perspective in which the hypothetical viewpoint is located high up in the sky over the northeastern sector of the city. Until recently, this map has been used to support various hypotheses about the urban development of Bologna, but now software applications enable us to move on to a qualitatively different kind of analysis of the sources.

The fresco in the Vatican is being used as the main source to produce a three-dimensional model of Bologna in the second half of the sixteenth century. To this end, there is a need to ascertain the reliability of the fresco, using scientific criteria, in order to be sure that this was a pictorial account of the city as it really was and not just a fresco based on the painter's imagination. Because of the complex changes that have taken place in the urban layout, such an analysis cannot be based simply on a comparison between the elements reproduced in the pictorial map (streets, buildings, squares) and the surviving structures. It is necessary to carry out a careful analysis of the map.

An initial analysis shows that the map has a high degree of reliability: there is a close match between the street frontages of the buildings and the measurements of the building plots in the land registry records (1833). These comparisons lead us to believe that the Vatican artist worked from a plan made in the same period that has unfortunately not been handed down to us, showing the development of the city center on paper.

Historians are able to use the fresco of Bologna as a reliable historical source, free from inventive work and flights of imagination by the artist. It is important to underline that, for the purposes of our project, the map in the Vatican is of great interest for the reconstruction of sectors and individual buildings that have undergone major refurbishment or reconstruction.

Moreover, thanks to the map's scale, it is possible to carry out a close reading of the buildings, including features of individual buildings, both civil and religious.

From Historical Sources to CAD Modeling

The city itself is a complex form. It is not an entity that perfectly conforms to elementary and standard geometric forms. It is characterized by irregular, geometric forms found in its external confines and in its complex division of streets and housing lots. In these situations, there does not exist a street that is perfectly rectilinear or a tower that is exactly vertical. Also, the dimensions and forms of habitation spaces are not perfectly subjugated in one, simple reticular form. This diversity of form provides information about the genesis of the city itself, offering hints as to how the city appeared in other times and how it expanded.

The precision of topographic operations when dealing with the necessity of reducing scale, can only be solved today with digital cartography based on CAD generating systems, which allows the possibility of overcoming the scale problem in graphic representation. The computer does not memorize the graphic nor the mathematical characteristics. Instead, it allows each mathematical element to be associated with information concerning the characteristics of the objects themselves.

In this way, the computer initiates the algorithm of representation, or in other words, "produces" the image instead of merely reproducing it, which allows the output of the machine to regenerate every time from zero.

The 4D Interface System

The interface of the prototype version of Nu.M.E., the New Electronic Museum, is designed and implemented in order:

- to provide different types of users with simple and efficient navigation tools to look back over the historical and urban development of the city from the end of the first millennium to the present day.
- to show visitors the fundamental role that historic research plays in the project.

The New Electronic Museum Interface consists of a VRML browser for the 3D display and specific 4D navigation tools in order to improve management of both spatial navigation and temporal navigation.

The spatial navigation orientation tools consist of:

- A 2D orientation map of the area of Bologna reproduced in the New Electronic Museum that allows visitors to visualize their position in the virtual world with a red rectangle and the direction for viewing the city with a green rectangle. The map can be shown or hidden by means of a button on the console.
- A virtual terrain model that allows the user to visualize each virtual building considering the altitude parameter in order to measure distances and areas and to carry out simulations related to the relief of the terrain.


The temporal navigation orientation tools consist of:

- A time-bar that includes a text field showing the year on display and a bar with a cursor showing the time line. By selecting a year, it is possible to visit the entire reconstructed city.
- The sound environment by which each century is associated with a different soundtrack, that allows visitors to identify, at an intuitive and perceptive level, the period in which they are visiting the city.

Finally, the New Electronic Museum interface needs to make clear some concepts related to the historic research:

- The need to reverse the chronological order of the research itself. The visitor begins to tour the city in the present and then travels backward in time, watching existing buildings disappear into the ground or change appearance, after which the buildings that no longer exist then pop up.
- The principle that the New Electronic Museum shows only as much as the historical sources will justify. Each building is therefore accompanied by a hypertext link with a description and references to the historic sources on which the virtual reconstruction is based.

The New Electronic Museum of Bologna (www.cineca.it/visit/nume/) is supported by the National Research Council (CNR), the University of Bologna, Centro Ricerche S.C.R.L., and CINECA, the supercomputing center for a consortium of Italian universities.



A Creative Journey

Moderator
Victor Raphael
Artist

Panelists
Bob Goldstein
Digital Consultant

Jane Raphael
Wonderland Avenue School

"Victor Raphael @ ZZYX A Creative Journey" was created by Victor Raphael with the support of ZZYX Visual Systems. The CD-ROM is a virtual exhibition of Raphael's "Space Field" series that can be viewed anywhere on a personal computer. It has been accepted into the collections of The Museum of Modern Art, New York, the Bibliotheque Nationale de France, and the Skirball Museum, Los Angeles. Beyond its artistic merits, the CD-ROM is an educational tool. Teacher Jane Raphael has used the CD-ROM in her multi-age primary school classroom.

Jane Raphael explains: "As we prepare to step into the 21st Century, educators continue to ask, 'What's worth teaching?' and 'How will learning best occur?' Information and knowledge are expanding at such an accelerated rate that what our children need to know is very different from what was needed a generation ago. 'A Creative Journey' can help educators who are asking these questions. It renders the universe a tangible miracle to students. It brings the question 'What is our place in the universe?' into focus as a point of student inquiry."

The artwork, based on NASA images from outer space, and the audio, based on plasma sounds from our solar system, send students on a personal journey into space. As they navigate through the artwork, students use technology as a vehicle for exploration. They see images of the planets, stars, and galaxies – imagery that is both awe-inspiring and provocative. One student saw rocks banging into a spacecraft, and the next moment, she was using the cursor to trace the tendrils of a spiral nebula.

As the connection between the arts and learning becomes clear and evident, "A Creative Journey" is a tool to bring art into the classroom. Jane Raphael continues: "Viewing the artwork is the common experience that allows us to respond, interpret meaning, and make critical judgments, as individuals in a group context. When students view Victor's artwork, I raise such questions as: 'What do you see?', 'What is it about?' and 'How do you know?' The students' descriptions of what they see ranges from twisters and rainbows, to the moon, earth and stars. This talk about art seamlessly transitions into talk about science. It is in this transition that I am able to access what the children already know about the universe. The unanswered questions they bring shape future investigations into this unit of study. Children's ability to critically focus on an artwork is developed when they need to find the point of symmetry in order to interact and navigate the artwork. Using the art as a visual text to be 'read' for specific information gives support and confidence to those students whose strengths lie in visual perception."


Once students have explored the artwork and raised scientific questions about our solar system, they are able to explore Victor Raphael's studio interactively, and through the movie section, observe the artist and his creative process. The CD-ROM also explores collaborations with computer artists and technicians. As Bob Goldstein relates: "We shot video of Victor in his studio creating the 'Space Field' series and explaining the history and methods of his work. We made videos of all the technicians – digital camera operator, computer artist, programmer, QTVR artists, and Iris printer operator – to demystify the process and show other artists and students how art, science, and technology come together. Few people have seen an Iris printer, and fewer still know how it operates. As the Iris is the leading device in the fine art printing world, we thought it important to feature it."

Goldstein describes how the CD-ROM project began: "I began working with Victor Raphael, who came to me to help him produce fine art Iris prints from his Polaroid work. As Victor and I got into the process, I became fascinated with the way his images were transformed when they were digitized and viewed on a computer monitor. I was also interested in the process Victor was personally going through, seeing his art in a new way and beginning to work in the digital world. I thought that Victor represented a new breed of artist making the transition from traditional to digital art,

and I proposed we document this transformation. Because we were using a myriad of sophisticated digital devices and highly talented technicians, I thought we could blend the art and science of digital imaging into an interesting 'creative journey' that would document not only an artistic process, but the entire late-20th Century digital production process as well."

Victor Raphael has a strong emotional attachment to his subject matter: "When I was a kid growing up in the 1950s, I was excited about the prospect of being an astronaut and going out into space... I think that working with these space images was a way for me to put myself out there and create my own fantasy about space. Because we live in a time where the common experience of viewing the heavens has been minimized by urban light pollution, our primal desire to connect to something larger than ourselves goes unsatisfied. Especially with young learners, the CD-ROM provides an easy entry point for students who are interested in making sense of the world."

In the opening artist statement that begins the CD-ROM, Raphael continues: "All of these images of space are new and have come into being in our lifetime... I'm depicting space as a metaphor for a micro-macro relationship, both of which offer a glimpse into infinity... I want people to slow down when they look at these images because there is something incredible about our place in the universe, when we can think about ourselves as a very small planet in perhaps endless and vast expanses of space." For students of all ages, "A Creative Journey" integrates art, science, and technology to provide an initial engagement in a compelling unit of study.



Journey

Presenters

Ria Bagaybagayan
Candace M. Borland
Naree Wongse-Sanit
Getty Education Institute for the Arts

Anne Marie Schaaf
J. Paul Getty Museum

Art and Technology: Electronic Resources from the Getty

ArtsEdNet: A Web Site in Progress (www.artsednet.getty.edu)

ArtsEdNet began in 1995 with 250 Web pages. Today the Web site includes over 2,000 pages. In this paper, The ArtsEdNet team details the developmental process that the site has undergone in becoming an expansive and extensive Web resource.

Devoted to comprehensive arts education, ArtsEdNet offers a wide variety of curriculum resources including lesson plans and online images from ancient to contemporary times and representing cultures from around the world. ArtsEdNet also provides online professional development materials and an online community called ArtsEdNet Talk, an email discussion group with over 1,000 participants.

At its inception, ArtsEdNet consisted of materials that originated as print resources adapted for the online environment. Gradually, the team developed curriculum resources exclusively for the Web, utilizing the interactive aspects of the Internet. The four-year research and development period for ArtsEdNet is not yet complete; the demands of our audience and of this medium are constantly changing and so must the site.

Keeping in mind the audience that a Web site serves is crucial to its development. Over the years, we've devised guidelines for writing and presenting information on the Web in a manner that is useful to teachers.

Exploring the Artworlds of Los Angeles: Worlds of Art

This innovative Web curriculum resource provides lessons that tap Los Angeles' many artists, museums, community art programs, and public art. Teachers can bring LA's worlds of art into the classroom by combining use of the Internet with an interdisciplinary approach. Teachers outside LA can use these lessons to build connections between art learning and the art worlds of their own communities.

Worlds of Art comprises several lesson units designed to work together. The core unit, Understanding Artworlds, consists of four lesson plans that help students broaden their understanding of art and culture. It also includes components on exploring LA art worlds on the Internet and understanding the importance of art in other cultures. All of the lessons included were developed to correlate with the California Frameworks for various subject areas addressed.

The other lesson units, also known as multicultural units, allow students and teachers to delve more deeply into particular art worlds in Los Angeles. Three are now available: Mexican American Murals (November 1998), African American Artists (January 1999), and Navajo Art (April 1999). Each unit centers on a group of artworks found in the Los Angeles area and provides online images accompanied by multiple core and supplemental lessons, as well as worksheets and other handouts. Students and teachers can explore these art worlds, learn more about the artworks and the themes they address, and try making art that reflects what they have learned about each particular art world.

Worlds of Art is based on an approach to teaching and learning that is thematic and inquiry-based. It is part of ArtsEdNet (www.artsednet.getty.edu), the Web site of the Getty Education Institute. ArtsEdNet aims to provide K-12 teachers around the nation with quality curriculum materials based on art from ancient to contemporary times and from Western and non-Western cultures. It has been available on the Web since September 1995 and receives an average of 10,000 visitor sessions each week.



ArtsEdNet home page, the Getty Education Institute for the Arts' Web site.
www.artsednet.getty.edu

ArtsEdNet

Presenting and Managing Electronic Visual Resources: The Getty Experience

The Getty Education Institute and the J. Paul Getty Museum collaborated to produce online exhibitions utilizing the unique interactive qualities of the Internet. This collaboration focused on two of the J. Paul Getty Museum's opening exhibitions, "Making Architecture: The Getty Center from Concept through Construction" and "Beyond Beauty: Antiquities as Evidence," both of which have since closed in the galleries. Portions of these two exhibitions are still available in an online version on ArtsEdNet (www.artsednet.getty.edu), with added educational components.

Through this collaboration, the Getty has used digital images to virtually recreate elements of a physical exhibition. "Looking at the Art of Ancient Greece and Rome" on ArtsEdNet uses artwork and information from the exhibition "Beyond Beauty." The physical exhibition consisted mostly of three-dimensional sculptures. One concern was to attempt to translate these three-dimensional works onto a two-dimensional page. At the same time, because the main audience of teachers is often not able to utilize more advanced plug-ins such as ShockWave due to hardware limitations, we needed to find a simple way to accomplish this. We created "rotating" images by using an animated GIF that connects a quick sequence of several static images. For example, a statue of Aphrodite shows all perspectives (front, side, back, and other side) of the goddess in succession, giving the effect of movement around the statue. Despite the flat screen, viewers get a sense of the object's three-dimensionality.

"Trajan's Rome: the Man, the City and the Empire" on ArtsEdNet also utilizes material from the exhibition "Beyond Beauty," this time creating a unique and innovative "virtual tour" of the Forum of Trajan. Collaboration among the Getty Education Institute, the J. Paul Getty Museum, and the University of California, Los Angeles, produced a virtual reality model of Trajan's Forum based on available archaeological evidence. The video virtual reality tour of the Forum of Trajan played continuously in the physical exhibition and is now available online, along with QuickTime movies and VR stills accompanied by interviews with curators, professors, and archaeologists. To complement this technological resource, ArtsEdNet added an online curriculum unit called "Trajan's Rome: The Man, The City, The Empire." Six lesson plans explore life and culture during Trajan's Rome.



Looking at Art of Ancient Greece and Rome: An Online Exhibition




Introduction

Welcome to ArtsEdNet's exclusive presentation of selected works from the [J. Paul Getty Museum](#). This virtual exhibition draws from art objects that appear in *Beyond Beauty: Antiquities as Evidence*, on display at the museum at the new [Getty Center](#) during its opening year.



Looking at Art of Ancient Greece and Rome: an online curriculum resource on ArtsEdNet



Workstation for Art Access, the J. Paul Getty Museum's interactive multimedia system, in an Art Information Room at the Museum

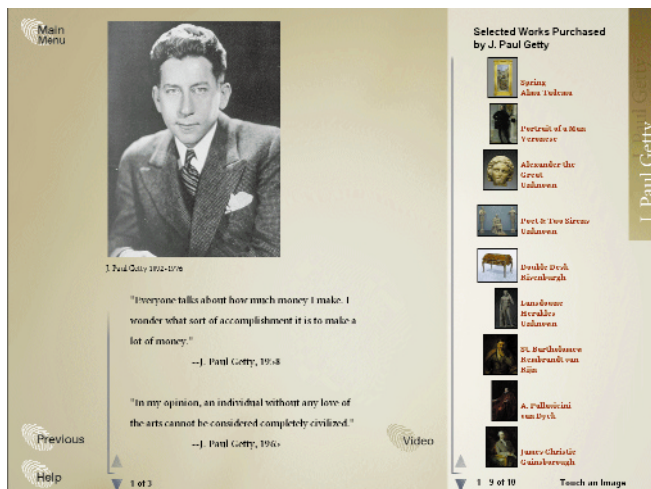
Within Art Information Rooms located adjacent to galleries with works of art, the J. Paul Getty Museum presents information about the museum's collections in an interactive multimedia system called Art Access. Using an underlying database structure, the touch-screen system allows visitors to acquire more information on works of art and artists. Art Access is a resource for the gallery experience, not a CD-ROM history of Western art. Why would the Getty Museum supplement real objects with art on the screen? Because a database system can provide great depth of intricately layered content, well beyond the brief words on a wall label. Curators cannot be there in person for every visitor, but this virtual tour can serve as a guide at any time.

The structure of Art Access is based on objects. Each object record contains multiple images, informative text with hyperlinks, and links to groups. Objects also link to their artists, who are indexed by specialty, nationality, and type. Artist records can contain portraits, life dates, text, audio, and video. A variety of indexes combine works of art from multiple departments in various ways: Time and Place; Subjects, which uses language to engage a layperson; Works of Art, an index by type of object that uses vocabulary for knowledgeable visitors; and Reference, an overall listing with a visual key. Additional elements enhance the object and artist content. Hyperlinks lead to definitions for glossary terms, and presentations offer additional information on topics within indexes. Viewers who seek more guidance can turn to introductory videos from the director, curatorial tours on video, and videos of manufacturing processes.

All elements of Art Access – text, images, audio, video, and index categories – are stored in a custom-made database and retrieved from servers when a user touches the screen. With quality as the first priority for onsite use, the museum has adopted various standards for its digital resources. Having made a significant investment in these resources, the museum chose a modular database, separating content from screen design and functionality, to permit multiple uses for this information. The museum's new collection management system, AMICO (Art Museum Image Consortium), the Museum's Web site, curricula on ArtsEdNet, and in-house slide presentations already use elements of Art Access content or may do so in the future.

The production process for a complex, truly live, interactive system based on a database model differs significantly from the usual one-shot, finished publication such as a CD-ROM. Editors must pull together creation, management, storage, and delivery of high-quality modular content, while concurrently handling both the continual flow of new content and review and adjustment of existing content. Since the system is non-linear, editors must juggle the multiple, evolving digital resources – text, images, video, audio, and animation – and maintain the complex links among them in a continuing campaign of editorial and technical management.

Screen for presentation on J. Paul Getty in Art Access, the J. Paul Getty Museum's interactive multimedia system.



Art Before Technology or Technology Before Art? That is the Question

Clifford Cohen
AnimAction, Inc.
cliffo@animaction.com

Workshop

In 1989, AnimAction was formed with a single purpose in mind: to create a unique and innovative environment for young people, where they could experience the spirit of collaboration, develop new skills, and exercise freedom of expression, all with the ultimate goal of articulating a powerful message through the medium of animated film. AnimAction has trained thousands of students on development and production of classical and computer animation in the United States, Canada, and Europe.

AnimAction has worked with children involved in programs at: The World Animation Celebration, The Los Angeles Unified School District, LA's BEST After School Enrichment Program, LA Children's Hospital, Glendale Unified School District, The Office of Criminal Justice Planning, YWCA, YMCA, Torrance Unified School District, Long Beach Unified School District, Sweetwater Union High School District, The San Fernando Arts Council, The Light-Bringer Project, The Gene Autry Museum, The Federation of Saskatchewan Indian Nations, Sioux Lookout Zone Hospital, The Solicitor General of Canada, The Teachers Advisory Council on Alcohol and Drug Education (TACADE), and The Mentor Foundation.


Introduction

The graphic arts and animation industries have grown in leaps and bounds since AnimAction was first started, particularly the animation industry. Today more than ever, the animation industry is crying out for artists to work in traditional style and computer-generated animation. Now that the Internet is available to the general population, our world's communication system is becoming easier to traverse, and the need for technicians is becoming greater. Indeed, the whole media production world is opening up, and there are more employment opportunities for a variety of careers within the rapidly expanding fields of animation and digital production.

However, there seems to be a great crevasse between the educational world and the industry. Should we teach our kids how to work, live, and breathe computers first, whilst not seriously taking into account a solid foundation in art and the use of the pencil and paper?

Of course, the animation industry looks primarily for experienced artists who have acquired these artistic skills for a number of reasons. They are more interested in working with people who can draw and understand the flow of a line. Animation studios train these artists with the necessary digital technology, making them ready for successful employment.

Art / Technology



Art Before Technology or Technology Before Art? That is the Question

Goals of the Program

- Introduce the classical animation process, stressing the importance of the basics and how a solid foundation in the art of animation will help in digital production.
- Demonstrate and involve all participants in a hands-on experiential workshop and give them the opportunity to take home their own products on video.
- Demystify classical animation production in the classroom (grades 5-12).

The Program

- Three hour workshop
- Ages 8-80 (presentations can be designed for a wide age range under the same roof)
- Participants are arranged in groups of up to four people, as small production teams. Each team works together throughout the workshop. From the very beginning, this introduces and builds the teamwork.
- From this launching point, we work through every stage of animation production with the goal of producing up to 10 seconds of animation. Each team takes on the challenge of producing a short film based on the theme: a whimsical look at a digital animation artist compared to a classical animation artist.

Each animation stage is covered:

- Importance of the story concept
- Character development
- Timing
- Storyboarding
- Production
- Color
- Filming

Methodology and Quotes

AnimAction's staff consists of artists and students of scriptwriting and animation. We're a very young-minded company made up of extremely curious minds. In this workshop, we demonstrate the endless possibilities of peer education and how it inspires us all as we work together experiencing each other's talent.

AnimAction staff artist Esdras Varagnolo, independent 3D animation student, and SIGGRAPH member, says: "Working with kids has improved my improvisation ability and constantly makes me look back at the basics for inspiration. Digital production is a tool used as an enhancement. The real strength is in the art!"

AnimAction staff artist Wade Bradford, English masters program, California State University, says: "I get to discover brand new styles and fresh story concepts working with kids."

AnimAction staff artist Brigitte DiSalvo, student at Santa Monica College says: "Working with kids has improved my timing and conceptual creative approach. It's also great practice, and somehow I seem to work faster now. A solid foundation in art is essential as the computer is only a tool. Creativity comes from within."



The Atmosphere: Incorporating Interactive Multimedia into the Classroom

MaryEllen Coleman
IBM Corporation
mea@us.ibm.com

Classroom

The Atmosphere is an interactive multimedia project that I did while I was a graduate student at Rensselaer Polytechnic Institute. The title comes from the Lutgens and Tarbuck textbook used in an introductory atmospheric science course at the State University of New York, Albany. This course is offered to non-science majors to fulfill a distribution requirement. Most of the students in the class find math and science very difficult and are unmotivated to study the material. So a multimedia application for this class serves two purposes: to explain difficult material and to motivate students to learn the material.

This project was a collaborative effort. A classmate worked on the project with me. She is more skilled at graphic design, while I am more skilled at programming. We also enlisted the help of two graduate teaching assistants in the Atmospheric Science Department, who were responsible for this class. They were our content experts. They identified topics that were especially difficult for students to grasp through lecture and readings, and gave us books and lecture notes that explained these concepts for us.

My partner and I wrote scripts based on this material and worked with the teaching assistants to ensure that our text was technically accurate. Then, we illustrated the scripts and motion through simple storyboards. These are just pencil and paper stick-figure diagrams that roughly show the illustrations that appear in the scene, and use arrows to indicate motion. Again, we met with the teaching assistants to review the storyboards to make sure we were still on track. Next, we found clip-art illustrations and built our scenes inside our tool, Asymetrix Toolbook. We used arrows to represent the atmosphere, and used Toolbook's hide and show commands to "move" the arrows to show the atmosphere's motion. In another scene, we used Toolbook's motion recorder to move an object from one point to another.

When we put it all together, we showed the application to the teaching assistants and their professor to get their feedback. We also asked co-workers who had never taken the class to try out our program and let us know what they thought about its navigation, layout, and content. Everyone was pleased with the program. Most everyone found the navigation intuitive. Many suggested that we add a "back" button, which we did. One reviewer suggested that we add a quiz at the end, in a future release.

We packaged the program on two diskettes and gave it to the teaching assistants to distribute to their classes as optional materials to supplement lectures. They reported that many students used the program and conveyed that it did help them better understand the concepts. The teaching assistants then began to use the program in class as part of the lecture. The professor did not have to wave his arms around any more to describe precipitation processes. Students could now watch precipitation in action.

Our process can be summarized in four major steps:

- Planning (10%) — audience analysis, anticipated use
- Design (40%) — script writing, storyboarding
- Development (40%) — drawing illustrations or finding images, videos, and sound, putting the pieces together in the authoring tool
- Testing (10%) — reviewing the final piece with anticipated users, gathering feedback, incorporating feedback into future releases

Animation is not a technique for film studios or design students only. However, it is time-consuming. We spent about 10 weeks working on The Atmosphere, part-time. The professor and teaching assistants wanted to incorporate multimedia into the classroom, but they didn't have time to develop the application while also teaching their classes. My partner and I had the time, but were not experts in atmospheric science. Because the teaching assistants were willing to spend a little time with us to review our work along the way, we were able to form a successful collaboration. After this project was completed, the professor realized the power of using multimedia in the classroom. He has asked his teaching assistants to find more video and animation to incorporate into lectures. A large volume of educational material is available free from the World Wide Web.

City@Peace, sponsored by the Community Mediation Program in Santa Barbara, California, brings together a diverse group of teenagers, age 13 to 19, from all socio-economic and ethnic backgrounds to promote peace and cross-cultural understanding through the performing arts. Assisted by professional artists trained as mediators, the teens explore conflicts in their lives and use improvisational theater techniques to write, produce, and perform original musical plays. Through collaboration with the professional artists (writers, actors, directors, visual artists, dancers, composers, musicians, poets), the teens grow in self-esteem and become productive (rather than destructive) members of their communities. They also invite an ever-expanding audience of their friends, families, extended families, and community supporters to enjoy their work.

This year, City@Peace adopted interactive, computer-generated sets to replace traditional props and backgrounds. The computer imagery shows physical locations where the action takes place (for example, in homes, on the streets, etc.) The image-processing capabilities are also used to show not only reality, but how the reality might feel. The German Expressionists called this "Theatre of the Mind."



The ColorWeb installation is a subset of Web-based materials created in the Exploratory Project at Brown University. An exploratory is a computer-based combination of an exploratorium and a laboratory that embodies an approach to learning by experimentation and investigation. It provides multifaceted, interactive microworlds that model objects, concepts, and phenomena, and that exhibit appropriate behaviors when interacting with students.

Our current implementation embeds Java applets in a hypermedia framework that (a) provides an immediate context for both the individual applet and for the larger conceptual frame of reference, and (b) enables students and teachers either to work within the framework provided or to place the applet in an environment of their own choosing.

The first set of exploratories on this site teaches basic concepts in additive and subtractive color mixing. These modules are being used in our introductory computer graphics course as classroom demonstrations, in assignments, and for self-directed learning. They are appropriate for a wide audience and assume no prior knowledge of color theory or computer graphics. The second set is designed to develop an intuitive feeling for the signal processing aspects of color perception and has been used by teachers around the world in their existing curricula. These exploratories are appropriate for a general technical audience.

Text for the color mixing applets is from Anne Morgan Spalter's recently released text *The Computer In the Visual Arts* (Addison-Wesley).

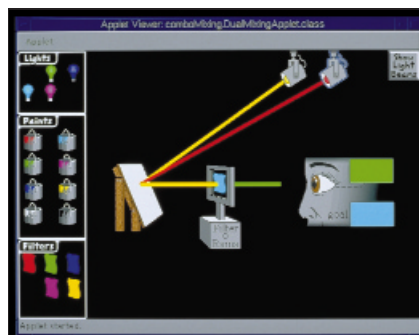
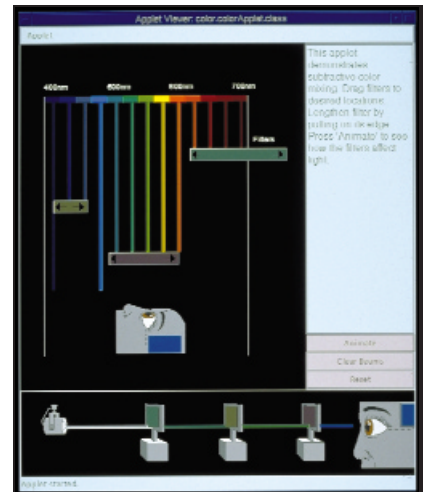
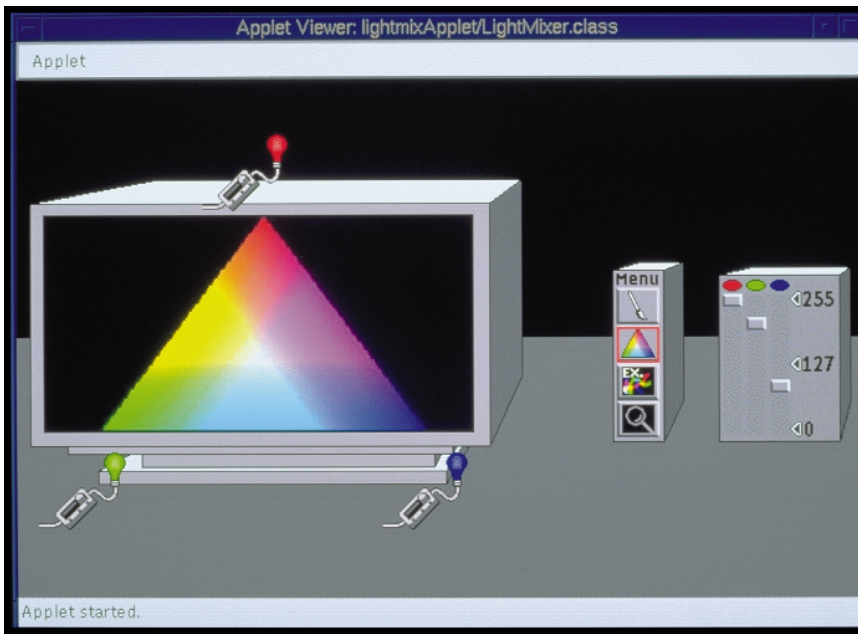
Text for the signal processing color applets is by John Hughes.

The signal processing color applet suite was developed by John Hughes, Adam Dopplet, and Jeff Beall.

Color Mixing Applets were brought to Java Platform 2 compliance by Robert Benjamin George.

Exploratory Project Leaders

- Andries van Dam
- Rosemary Simpson
- Anne Morgan Spalter
- Brown University



Computer Camp: For Girls Only!

Abstract:

A Web-based computer camp for girls only (7th - 10th grade) was offered during the summer of 1999 in the Philadelphia region. The campers learned HTML and JavaScript programming in order to understand basic computer concepts, while enjoying individual creative exploration and developing Web pages. The female orientation was intended to assist in thwarting the image of computer and math-related activities as male dominions. 2D and 3D graphic design, as well as interactivity and human-interface concepts, were emphasized. Programming topics included: language syntax, functions, block structure, sequential and iterative statement sequences, and variables. Computer fundamentals (such as word processing, and basic understanding of hardware components) were summarized. Students were assigned female mentors who communicated with them via email and chat rooms during the camp week. The goal was to have each girl create a personal Web page.

Description

Computer camps have been traditionally viewed as a haven for sallow-faced, antisocial, male teenagers who prefer to bask in the glow of a terminal rather than in the summer's sun. Often, this male orientation is enhanced by promotional materials that emphasize activities like "hacking" and "Dungeons and Dragons" in addition to traditional programming. Although some of the camps are better than others in attracting females, it is a matter of serious concern that they may be among the factors contributing to the early perception of computer programming as an inappropriate career choice for females. Even worse may be the adverse impact that the reduction in hands-on recreational computer time has on the ability of females to compete in higher-level programming courses as they proceed through the curriculum. These assertions find support in studies indicating that although males and females elect high school computer courses in equal numbers, only a third of the bachelor's degrees in computer science and computer engineering are issued to women. This disparity begins as early as the freshman year, when about half of the qualified males choose scientific majors, with only a sixth of the qualified females making similar selections.¹

In an effort to reverse this trend, a number of organizations have attempted early intervention programs. One of the more notable of these is the Math Options program, now in its ninth year at Pennsylvania State University. This project targets seventh-grade girls, who are brought to college campuses to attend a day of activities, panel sessions, and workshops led by women who have chosen math-related careers. There, the girls are given an opportunity to communicate one-on-one with the women about their professions and lives. Follow-up sessions and mentoring relationships are also provided. This project has met with tremendous positive approval, and serves many hundreds of girls each year.²

Having been involved with the Math Options program since its inception, I wanted to extend this concept within the computer camp scenario. One of the locations expressing interest in offering this program was Mercer County Community College (MCCC), a suburban campus in central New Jersey serving a diverse student population. Among the many non-credit offerings is a group of summer courses called Camp College for junior high and high school students. For a number of years, MCCC has offered various week-long computer camps entitled Computer Hackers Workshops, which emphasize programming in the Pascal and C languages. The project goal is development of a computer game. Interestingly, to the best of the program directors' recollections, not a single female had ever applied for admission to any computer camp session. Clearly these demographics indicated that something was amiss.

Camp

The first step, therefore, was to develop a theme that might better attract the large untapped market of young female computer users. The Internet and World Wide Web seemed to promise a viable educational scenario, since its multimedia presentations, timeliness, and easy accessibility attracts both girls and boys, yet its verbal orientation, particularly in chat sessions and via email, provides an especially strong female appeal. Also, the slight edge in verbal skills among girls of high school age possibly makes them better “surfers” than boys.

Penn State chose seventh grade as its target age group for the Math Options program because, at that point, the sex-based math and verbal differentiation is not yet as noticeable. Our camp would include students between 7th and 10th grades, since that was currently a successful age group for the male attendees. Since same-sex schooling has been shown to enhance participation and learning by females in male-dominated subject areas (such as mathematics and science), it was decided to offer the Web-based camp for girls only. This labeling in the promotional brochure was intended to further attract the attention of female students and their parents.

A few years earlier, I co-developed a Web-based “Introduction to Computer Programming” course at Drexel University, aimed at college freshmen.³ This course used HTML and the JavaScript language to teach such concepts as language syntax, functions, block structure, sequential and iterative statement sequences, and variables. Computer fundamentals (such as word processing, and basic understanding of hardware components) were summarized. Good human interface design and interactivity were emphasized, through the inclusion of tables, forms, and hyperlinks in Web pages. Students were provided with software tools for 2D graphics, and they were encouraged to explore use of original or customized imagery in their Web presentations. Some of the more advanced students surfed for animations and other code that they linked to or modified and included directly in their pages. The materials developed for this introductory course were simplified somewhat and used for the programming training portions of the computer camp.

Prior work for the graphics portion of the computer camp was based on my earlier collaborative effort involving the Franklin Institute Science Museum in Philadelphia and the Community School of Music and Arts in Mountain View, California. Here, cartoonist Mike Mosher instructed students to develop images that depicted their fears. These drawings were then enhanced with sound and 3D graphics in order to form a kiosk-style display. This successful exploration with evocative feelings should transfer well into the Web modality.⁴ The Franklin Institute, which serves nearly a million visitors each year, is currently involved with the National Science Partnership for Girl Scouts and Science Museums. Girl Scouts in regional clubs (from Maine to Virginia) have the opportunity to participate in educational programs at the museum, including overnight camp-ins and mentoring sessions by women in science-related careers. Illustration 1 shows a recent lecture/demonstration about Web programming to a group of Cadettes and their troop leaders.⁵

The Girls Only Computer Camp had as its goal that each participant develop her own original Web page, from scratch (using a non-WYSIWYG editor), around a personal theme. The campers began the week by learning how to use the browser and search engines in order to surf the Web for specific topics (in a scavenger-hunt fashion). The girls were directed to specifically female and kid-oriented sites and permitted to explore for items of their own interest (parental control software was used to insure some level of safety). The campers were given access to chat rooms, email, and digital whiteboards, in order to facilitate cooperative work within the group.

Computer Camp: For Girls Only!

References

1. Statistics from the Computing Research Association's Taulbee surveys.
2. Math Options can be reached by contacting: Alice Sayles, Penn State Abington, 1600 Woodland Road, Abington, Pennsylvania USA 19001-3990, aws5@psu.edu, www.abington.psu.edu/MathOptions
3. Report presented at ACM SIGCSE '98, see proceedings.
4. For "Fears" information and photos see: www.ylem.org/artists/mmosher/fears.html
5. Information on the National Science Partnership for Girl Scouts and Science Museums can be found at: sln.fi.edu/tfi/programs/nsp.html

Each camper was assigned an adult female mentor, similar to the Math Options scenario, who greeted her mentorees electronically and remained available for email consultation throughout the week (and possibly also after the camp was over). Various hardware and software tools (including 2D and 3D graphics, digital still pictures, animated videos, and audio) were explored for development of multimedia presentations. While the students began to storyboard their concepts for their Web pages, they were provided with small programming tasks in HTML and JavaScript that taught them the language basics. At the end of the camp week, the Web pages were mounted, and the students and their mentors had an opportunity to look at and interact with each others' work.

The collected set of pages was posted on the Web. Results of the program to date are reviewed in a Classroom presentation, including the course syllabus, examples of training materials, and student work.



Senior Girl Scouts and their troop leaders at the Franklin Institute Science Museum attending a lecture/demonstration by Rebecca Mercuri on careers in Website development.

Photo by Ed Wagner, Franklin Institute.

Digital Compositing and 2D Animation

Compositing and 2D animation is an exciting field for digital artists, effects editors, and designers. There are many opportunities in both the broadcast video and feature film markets for visual effects artists. Today more than ever, there is an increasing need for people with that special blend of creative and technical skills. Not only does one need to understand the basic “tools of the trade,” which can include a variety of software, one must also be aware of the visual nature of the process.

Remember, the bottom line is that all the technical considerations are unimportant when confronted with the question: “Does it look right?” Obviously this is a subjective judgment, and a good compositor able to make these decisions will always be in high demand.

What is a Composite?

Compositing is simultaneous multi-layering and design for moving pictures. It is the digital blending of a background plate and one or more foreground elements to create one seamless, well-integrated film image. Modern designs often use many techniques together, such as painting, retouching, rotoscoping, keying/matting, digital effects, and color correction as well as multi-layering to create complex animations and opticals for promotions, title sequences, and commercials as well as program content. Besides the creative element, there are other important applications for compositing such as image repair, matte painting, and wire removal – especially in motion pictures. The artist and the equipment can be crucial, especially where seamless results are demanded.

Digital Compositing Tools

The tools used in the process of creating 2D animation and compositing include:

- Paint programs
- Color correction utilities
- Warping and morphing tools
- Matte-extraction software
- General-purpose compositing packages

Basic Terms

Elements

Elements are your images. You may also commonly hear elements referred to as “layers.” A subset of elements, called “plates,” usually refers to original scanned sequences. Intermediate elements generated during creation of a composite are not referred to as plates, but as “precomps.”

Mattes

A matte is an image designed to control the transparency and opacity of another image. Mattes are used during compositing when we only wish a portion of one image to be included in the output image. You may also hear the term “mask” used when referring to mattes, and in general the two terms are usually interchangeable. “Mask” is more common when specifying an image that is used to control or limit a color correction (or some other form of image-processing) on another image.

Animation

Matte Generation

There are many different types of mattes, and there are many different methods used to generate mattes for compositing. This process, particularly when automated, is also referred to as "matte extraction" or "pulling a matte."

There are two approaches we can take to generate these mattes. The first is hand-drawing a matte for the object in question over every frame of our sequence. This is still occasionally done, but only after all other options have been exhausted, since the process of hand-drawing a matte for every frame of a sequence is time-consuming and error-prone.

One slightly less brute-force method of generating a traveling matte involves the use of splines to "rotoscope" basic outlined shapes for an object. The rotoSCOPE artist can specify certain key shapes, and the software will smoothly interpolate any in-between frames. Unfortunately, objects that move or change shape a great deal may end up needing a key shape defined for every frame anyway!

There must be a better way, you say, and there is. Instead, we rely more on procedural techniques, where some initial parameters are determined that are capable of extracting a matte, and then software is allowed to apply these same parameters over a sequence of images. Some of these methods include The Color Difference Method, which involves photographing the subject in a manner that greatly simplifies the extraction of a matte.

The Color Difference Method requires shooting the foreground object we wish to isolate in front of a uniformly colored backdrop. Any color for the backdrop (or "backing") may be used, as long as the foreground is essentially devoid of this color. The term "color difference" refers to the difference in color between the foreground and the colored backing.

The most common backing colors are blue and green, and the choice of which to use is generally determined on the basis of the subject's colors. If someone says that the subject must be wearing a blue shirt, then typically a green-screen shoot would be dictated. Generally, tests are done if the choice is not obvious. The process of extracting a matte with this method is known as "keying," and the extracted matte can be referred to as the "key." The software used to key something from its background varies in complexity and capabilities.

There is also a method of creating a matte known as "difference matting," where a frame of the scene without the subject is subtracted from a frame with the subject. In theory, all you are left with is the subject. In practice, slight lighting differences, shadows, and grain make the difference between the two images unpredictable. Difference matting is sometimes the only solution available, and it is actually a very useful first-pass that can then be cleaned-up by hand.

Even the best tools can have problems with certain images, and real-world situations often deliver to the compositor plates that

are less than perfect in terms of evenness, graininess, and absence of objects that are not intended to be seen in the final composite. To help deal with these, one almost always creates garbage-mattes around the subject. These are loose-fitting shapes designed to quickly remove problem areas from the scene.

Several of the techniques mentioned will often be used in conjunction with one another, combined until as flawless a matte as possible results.

Image Tracking and Stabilization

When photographing an element to be used for visual-effects work, one sometimes has the ability to specify that the camera be unmoving, or locked-off, for the duration of the shot. However, it is often not possible, or even desirable, to do this. Multiple shots without camera moves can become boring and lifeless. In situations where the need arises to composite elements that were shot without identical camera moves, one must resort to tracking. Tracking is the process of selecting a particular region of an image and determining that region's movement over time (on a sequence of images). The data are stored as a series of moves or positions and then applied to one or more images.

There are a variety of situations where tracking can be used. One reason would be to "stabilize" the sequence you are working on. Another reason would be the need to synchronize the movement of an object you are adding to the scene, with something already in the scene. (The object in the scene may be moving, or the camera may just be moving relative to the object).

Preparation of Elements

The best composites are those whose elements were planned and photographed with the explicit intention of creating a composited image. While this may seem like an obvious, elementary statement, you'd be surprised at how often it is ignored in the real world. In just a bit, we'll look at what can be done to help fix improperly shot plates, but first, let's look at some things that can be done to make everyone's lives easier.

Matched Lighting

Whether you are planning to integrate elements shot with blue-screen or synthetic CG images, or are just soft-splitting two plates together, it is critical that the different pieces look like they were lit with the same lights. Lights should hit the objects in the scene from the same angle, have the same apparent intensity, be of the same color, etc.

Matched Cameras

Almost as important as the lighting is the synchronization of the camera for all elements. Be aware of the camera's positioning relative to the subject and the height of the camera from the ground plane, and ensure that the same size lens is used when photographing all plates. If your camera moves throughout the shot, either plan on shooting it with a mechanical motion-

control move (a device that allows the camera to repeatedly execute the exact same move) or be prepared to do a lot of post-processing tracking to try and duplicate the move.

Film Stock

If possible, shoot all elements with the same type and speed of film (and expose and develop them similarly). Different film stocks have widely different grain characteristics, and the discrepancy can be obvious.

Element Repair

Fixing Elements

The usual situation in this business is that by the time you receive the elements for your composite, there is no longer an opportunity to re-shoot anything to correct faulty plates. At this point, your only option is to manually fix the problems. The following section lists potential and common pitfalls and gives a few solutions. Remember, every shot has its own problems, and the true test of a good compositor is the ability to come up with efficient and creative solutions to these issues.

Wire Removal

It's very common for visual effects to be mixed with practical effects that require things like gag-wires, harnesses, ropes, and other mechanical devices. It is not always possible to fully hide these items from the camera, so digital effects may be needed to remove them. There is limited commercially available software for wire removal, and the general consensus is that this software works well on the easier situations, but the more difficult shots are still going to be done at least partially by hand.

Plate Instability

As mentioned earlier, tracking software should be used, either to stabilize the plate or to "bounce" the new elements to match the background.

Mismatched Lighting, Cameras, Action

Trying to tie together images whose lighting doesn't match can be one of the most frustrating tasks a compositor can undertake. Troublesome highlights on objects can often be decreased via specific color-corrections and masking. When your A and B plates were accidentally shot from wildly different positions or with different length lenses, you may need to compensate by moving and scaling the various elements relative to each other. Even the timing of the action may not initially work correctly. Not only can this be dealt with by "slipping" the synchronization between the plates (so that frame 1 of element A is combined with frame 30 of element B, for example), but you can also adjust the speed of action by dropping, duplicating, or averaging frames.

Hand Painting

Ultimately, for certain problems, there may be no other solution but to hand-paint the offending area, possibly on every frame. This is usually somewhat of a last resort, but the fact remains that it is a totally valid and acceptable solution. Sometimes you may be able to hand-paint a fix on a single frame and then use tracking and warping tools to apply this fix to the rest of the sequence.

The Final Touches

The more time you spend compositing, the more you'll learn about what things are important in order to fully integrate the elements. There are a number of techniques (and several tricks) that can be used to trigger the visual cues the eye is accustomed to seeing. There are also a lot of common problems that can be easily addressed if they are identified. First and foremost, a good compositor should understand how a real camera behaves. You will often need to mimic artifacts and characteristics of shutter, film, and lens. Many of the items mentioned below are related to this issue.

Blur Levels

It is rare that every element in a scene is in sharp focus. Depth of field dictates that objects farther or closer than the focus point will grow more and more unfocused. Determine the distance your element should be from that focus point and blur accordingly. This focus relationship can change over time (rack focus). Be sure to match the animation timing as closely as possible.

Also, keep in mind that a moving object, when recorded on film or video, will have "motion blur." This is due to the distance the object moves while the film is exposed (or the video-camera is recording). If you wish to place a moving object into a scene, and that movement is something that you created (either as a 3D element or with a 2D transformation), you should plan on motion-blurring your element. Most high-end systems allow motion-blur to be added to a 2D move or rendered with the 3D element.

Lens Flares

In the real world, when a bright light source shines directly into a lens, you will get a flare artifact. It is often desirable to duplicate this flare when creating an image that has bright light sources that were not present in the original elements. Be careful with this effect, as it has become over-used in much of the CG imagery being created these days.

Film Grain

Another extremely common mistake with adding pure CG elements to a scene is to ignore the amount of grain that the rest of the plate has. The CGI element will consequently appear far too "clean." In general it's much easier to add grain to an element than it is to remove it, so for this reason effects photographers try to use the least-grainy film they can get away with.

Interactive Lighting

If your background scene has flickering or inconstant lighting, (and your foreground plate doesn't), you will need to do your best to match the background. In some cases it may be as simple as adding a fluctuating brightness effect that is synchronized to the background. In other situations, you may need to have articulated mattes controlling the light so that it only falls on certain areas.

Edges

Examine the quality of the edges of objects that are already in the scene and try to match them as closely as possible. Edges can be sharper or softer depending on the amount of backlight an object is getting or how out-of-focus it is. Be particularly aware of how edges are behaving over time. What looks acceptable on a still frame may "chatter" in the moving sequence.

Shadows

A common novice mistake is to forget the fact that an object should cast a shadow (or several shadows). Many methods of extracting an object from its background do not allow the object's shadow to be brought along. In this case, you'll need to create a shadow yourself. It is often acceptable to simply flop the matte of the foreground element and use it to darken a section of the background. Remember to match the rest of the shadows in the scene, in terms of size, sharpness, and density.

Black Levels

Black levels in the elements you add should match the background plate. This is probably the most often-violated compositing rule, particularly when integrating CGI imagery with live-action. Rarely are the darkest parts of a scene pure black, and if you want to add something new to that scene, you should make absolutely certain that the darkest parts of your new element aren't any darker. An excellent trick for ensuring that your element's levels are matched to the background is to adjust the monitor brightness to high and low extremes. This will help to bring areas that the human eye is less sensitive to into a range where differences become more obvious.

Camera Movements

The nature of effects photography often requires that the elements of a composite all be shot with an unmoving, locked-off camera. While this makes for easier composites, it doesn't really make for interesting (or believable) cinematography. Fortunately, it is possible for the compositor to add camera moves after the plates are shot. This move may be as simple as a nearly imperceptible camera shake, or as complex as a tracked, 3D match-move.

Atmosphere

Often there may be fog, smoke, or haze in the scene you're wishing to integrate with. The farther away from the camera your element is supposed to be, the more atmosphere would need to be added. Again, examine other elements in the scene that are at the same distance and try to match their levels. If the atmosphere in a scene is very uniform, you can probably get away with just decreasing the contrast in your element and adjusting its color. With more distinctive mist or smoke, you may need to use a separate smoke element and add it in explicitly.


Rendering

Last, but not least, you have to render your animation. The computer will render each frame in turn, each comprised of various operations to achieve the desired effect. The rendering may take minutes, hours, or days depending on the effect's length and complexity and, most importantly, the speed of your computer.

Parting Thoughts

Finally, remember that every person who views an image has a little expert he carries around with him. This expert, the subconscious, has spent a lifetime learning what looks "real." Even if, consciously, the viewer is unable to say why a composite looks wrong, the subconscious can notice warning signs of artificiality. Beware.

Special thanks to Ron M. Brinkmann for allowing me to excerpt his Digital Compositing SIGGRAPH paper.



3D computer animation has become wildly pervasive throughout our media, ranging from Saturday morning cartoons to high-end feature film creatures and effects. What may not be obvious, though, is just how these effects were created. For those seeking an awareness of the field, this presentation covers the basics of the 3D artist's production tasks.

Overview

1. Cinematic special effects can be seen as the creation of fictional events.
2. We are experiencing a shift from manual (practical) effects work to electronic digital techniques.
3. The most successful work is a result of combining the two approaches, capitalizing on the strengths of each.
4. 3D software applications are primarily used for creation of character elements, props, effects, or environments.
5. 2D software applications are primarily used for assembly and enhancement of individual elements, both digital and practical.
6. 3D software generally appeals to "object-builder" personality types, whereas 2D software appeals to "image-maker" personality types.
7. The key to mastering this medium is comfort and understanding of computer technology. Presently this is a heavily technical enterprise.
8. Always remember: the act of drawing is primarily visual and secondarily technical. Digital media are no different, but they demand more technical discipline.

Sample 3D Production Workflow

1. *The Modeling Phase*
 - A. Modeling typically accounts for 20 percent of a project's schedule.
 - B. Modeling requires the skills of a sculptor or architect.
 - C. Modeling involves creation of wireframe 3D geometry.
 - D. The geometry defined is typically NURBS (freeform deformable surfaces) or polygons (point meshes).
 - E. 3D applications can import EPS or CAD file types (vector-format data).
 - F. One can build from sketches, measured drawings, video frame captures, or 3D scans.
 - G. Modeling requires a good feel for 3D spatial visualization.
 - H. Character modeling is one area of specialization for artists.
 - I. Hard models, such as vehicles and props, are another area of specialization.
 - J. Efficiency of modeling is very important, to allow ease of production later on.
2. *The Lighting Phase*
 - A. Lighting typically accounts for 20 percent of a project's schedule.
 - B. Lighting requires the skills of a photographer, director of photography, or gaffer.
 - C. Lighting involves placement and adjustment of various light sources.
 - D. CG light types can be spotlights, point lights, ambient lights, or directional lights.
 - E. CG lighting is presently a crude simulation of real-world lighting complexity and subtlety.
 - F. An emerging more realistic lighting paradigm is that of radiosity calculations.
 - G. Lighting can have a multitude of approaches, such as theatrical stage, studio, architectural, cinematic, product, and character.
 - H. Character lighting is an area of specialization for artists.
 - I. Hard model lighting, such as vehicles and props, is another area of specialization.
 - J. Efficiency of lighting is very important, to allow ease of production later on.



Recommended Bibliography

Film Architecture: Set Designs from Metropolis to Blade Runner, Dietrich Neumann, Prestel

Digital Filmmaking, Thomas Ohanian, Michael Phillips, Focal Press

The New City, Lebbeus Woods, Simon & Schuster

American Cinematographer Manual, American Society of Cinematographers, The ASC Press

Cinematography: a Guide for Film Makers and Film Teachers, J. Kris Malkiewicz, Van Nostrand Reinhold

Special Effects Cinematography, Raymond Fielding, Focal Press

The Language of Visual Effects, Michael J. McAlister, Lone Eagle Publishing

3. The Texturing Phase

- A. Texturing typically accounts for 20 percent of a project's schedule.
- B. Texturing requires the skills of a painter, illustrator, or photographer.
- C. Texturing involves creation and application of images or surfaces onto the modeled geometry.
- D. Texturing is intrinsically related to the lighting of surfaces.
- E. Texturing can be designed or computed procedurally.
- F. Texturing is the key to establishing scale, age, and weathering of surfaces.
- G. Texturing can replace the modeling of 2D surface detail.
- H. Character texturing is an area of specialization for artists.
- I. Hard model texturing, such as vehicles and props, is another area of specialization.
- J. Texturing relies heavily on Photoshop or similar applications.

4. The Animation Phase

- A. Animation typically accounts for 20 percent of a project's schedule.
- B. Animation requires the skills of a classic animator or choreographer.
- C. Animation involves establishing the motion and performance of geometry, effects, lighting, and camera moves.
- D. Animation deals with working with abstract motion curves within graphs.
- E. Animation is a lifelong art and education, requiring many years of training.
- F. Character animation is an area of specialization for artists.
- G. Effects animation, such as explosions or dust effects, is another area of specialization.
- H. Camera animation is yet another area of specialization.

5. The Rendering and Output Phase

- A. Rendering computation typically accounts for 20 percent of a project's schedule.
- B. Rendering computation places a very high demand on hardware, such as CPU cycles.
- C. Rendering computation requires distribution of frame processing over a large network of machines, with most rendering done at night.
- D. Rendering computation requires massive amounts of hard-drive storage and RAM.
- E. Back-up of the resulting data is also a large task.
- F. Once frames are complete, they need to be put onto videotape or film.

6. Possible Hardware and Software Routes

A. SGI/UNIX Platform Route

Highest quality, professional level software. Has the ultimate flexibility and capability.

- High-end: 300mhz R12K Octane MX2 with Maya Infinity
- Mid-range: 195mhz R10K Octane SI with Maya Complete
- Entry level: 195mhz R5K O2 with Animator

B. Mac Platform Route

Easiest interface, good media capabilities. Common in graphics communities.

- High-end: 300mhz G3 PowerMac with Electric Image
- Mid-range: 200mhz G3 PowerMac with Lightwave
- Entry level: 200mhz G3 PowerMac with Strata3D

C. PC Platform Route

Good price/performance ratio. Becoming on par with Unix/SGI.

- High-end: Quad 533mhz P3 NT box with Softimage Extreme
- Mid-range: Dual 450 mhz P2 NT box with 3D Studio Max
- Entry level: Single 400mhz P2 box with Lightwave

See page 8 for a related classroom presentation

Playground
Workshop

The methodology of the Nu.M.E project, as presented in the Classroom Panel "The 4D Virtual Museum of the City of Bologna", is based on the virtual reconstruction of Medieval and Renaissance buildings and streets. This workshop presents the 4D navigation and graphic processing phase of the project.

In the new millennium, the world will be increasingly connected by high-bandwidth links that will allow the virtual visitor, seated in front of a monitor anywhere, to take the New Electronic Museum's electronic tour of the Medieval streets of Bologna. Visitors will be able to look up at the Garisenda tower, as Dante did, and be impressed by the gathering clouds above, go up to the top of the Asinelli tower, admire Piazza Maggiore as it appears today or as it appeared in the 13th century, see the urban settlement where the oldest university in the world was founded, and walk under the portici today or in the past. They will be able to choose any street or period, in the hope of being able, sooner or later, to see this extraordinary city in real life.

Implementation Phases and the Time Machine

The results of the research and the modeling work were entered into a virtual environment governing the "time machine." Modeling techniques made it possible to reconstruct buildings that are no longer in existence. The navigation system enabled views of historic elements that are no longer in existence and removal of items that were not yet in existence during a given period.

A "time machine" in the program's tool bar allows visitors to select a historical period for an electronic walk through the city of the past. Visitors can access views of Piazza di Porta Ravennana with the Asinelli and Garisenda towers and the surrounding streets, the Chapel of the Cross (fourth-18th centuries), three medieval towers that were demolished at the beginning of the 20th century, Piazza Maggiore, and the interior of the basilica of S. Petronio.

Coordination of the Working Environment

At this stage of the project's development, various groups are working in parallel, coordinated by means of a "shared environment" that enables them to exchange information among themselves. The Nu.M.E. Cooperative Open Environment is used to gather the data from the various laboratories and download it for processing.

The work presented in this workshop is only the tip of the iceberg of a great scientific enterprise with sophisticated software applications that the "visitor" cannot see, but which are the basis of the final product. Every mark, drawing, and movement is derived from in-depth historical research, since it is the quality of the historical research, over and above the software applications, that makes the difference between a multimedia product of a purely commercial kind, of which there are many on the market, and the New Electronic Museum project.


Four-Dimensional Navigation as a Summary of Historical Research.

In order to provide the visitor with the scientific information on which the product is based, every building, every urban sector, and every single item is accompanied by the sources on which the vector model is constructed and which make the three-dimensional navigation possible. These sources may be called up on the screen at any time during the electronic tour. One advantage of this electronic museum is that it is accessible at all levels of competence and knowledge, and visitors can make their own personal interpretation of the materials.

The project is both an instrument of scientific research and a new means of disseminating scientific knowledge (in this case, about urban history and cultural and conservation studies), and, at the same time, a means of stimulating interest in applications of technology that are likely to see further development.


Implementing the Nu.M.E. Project

The idea of making the Nu.M.E. or New Electronic Museum project accessible on the Internet (www.cineca.it/visit/nume/) arises from the need to facilitate interaction with historical researchers, to reach a greater number of users, and, of course, to publicize the project. For this reason, we chose to implement Nu.M.E. with VRML and Java that also make it possible to develop dynamic and interactive 3D environments in a portable, standardized, and platform-independent manner.



Briefly, the implementation process of the Nu.M.E. project consists of progressive introduction of 3D models of buildings in the virtual 4D city model. This process can be subdivided into four phases. During the translation phase, the CAD model of a building must be translated into VRML2 format, which requires modifying some characteristics of CAD solid modeling:

- To maintain only the surface information to be used in the VRML representation.
- To introduce a hierarchy among the different components that constitute a building, so that during the graphic processing phase it is possible to process them individually and to reproduce their temporal development, if required, separately from the rest of the building.



During the graphic processing phase, the VRML files are processed in order to make them photo-realistic. In order to improve performance in terms of download time and frames per second, we simplified construction of individual 3D architectural elements (roofs, portici, columns, etc.) with the use of photos as texture to be applied to them in a level-of-detail technique. It is necessary to strike a balance between a lot of details for maximum realism and quick drawing for maximum interactivity. Moreover, the Nu.M.E. project has a great variety of potential users (for example, historical researchers, urban planners, virtual tourists, etc.). Clearly, historical researchers need a more detailed 4D visualization than virtual tourists, who probably give preference to speed of access rather than to careful reproduction of too much detail. So we developed two models: high-resolution and low-resolution, characterized by the different resolutions in which 3D models and multimedia data are stored, but with the same interface. The high-resolution model may be used by accessing Nu.M.E through a LAN connection or a visualization in a local environment. The low-resolution model is accessed via a remote connection to the Internet.

During the spatial collocation phase, the VRML model of the building is placed in the area of the city recreated in the Nu.M.E. project, considering only the spatial coordinates, eventually moving and rotating it in order to place it in the correct position. Recently, we introduced the option of visualizing a 3D model of the area, in order to place the buildings also at the correct altitude. The temporal collocation phase consists of positioning the building in the temporal dimension. Each building is identified by: the VRML model created during the previous three phases, properties that vary with the passage of time, and the historical sources underlying the four-dimensional reconstruction of the building. From the implementation point of view, we associated each building to a set of different files:

- building.wrl: the VRML model created during the previous three phases.
- building.db: the text file that describes the properties of the building that vary over time (defined by the dates of construction, alteration, demolition, etc.).
- _building.wrl: this file allows the user to select the VRML file (building.wrl) to visualize a specific hypertext document (building.html).
- building.html: this file allows the Nu.M.E. visitor to examine each building individually, together with the historical sources that were used for its reconstruction. It consists of two HTML frames visualizing respectively:
 - fonti_building.html: a hypertext document that describes the historical significance of the building citing the sources used to collocate it on a time scale, visualizing its changes from 1200 to the present.
 - building_bis.wrl: a high-resolution VRML model of the building isolated from its context. In this model it is possible to add light-spots to enhance some components and specific viewpoints.

In each phase, a constant exchange of information is required among partners located on separate sites with different platforms and different computer skills. We therefore implemented the Nu.M.E. Cooperative Open Environment, which provides all the partners with remote access and updating of all the archive data, building by building and period by period, which is simpler than interacting directly with the entire virtual environment.

Recently, we have begun to design a new immersive physical access to the Nu.M.E. scenario in order to prepare a virtual theatre using technologies such as Reality Center SGI. There is reason to believe that the virtual theatre could be the natural development of our high-resolution model, and that it could provide a public venue for the New Electronic Museum to be seen by groups of visitors.

Creative Programming: Merging the Artist with the Computer Programmer

Mark Ollila
University of Gävle
molly@hig.se

Eva Carling
University of Gävle

*Classroom
Playground*

The Creative Programming program at the University of Gävle is unique in its approach to education and selection of students. It aims to give opportunities to highly creative and intelligent people to experiment and enhance their abilities in the world of digital media by combining their skills from other areas.

The area of multimedia systems and digital media itself is becoming increasingly important. There is a shortage of people with the required educational background who can design systems in this new medium, in research terms, as well as in visual development and implementation. The skills involved in creating software applications have expanded from traditional programming to include creative skills, such as graphic design and writing interactive narratives and storyboards. Research also plays a very important role in design of systems and implementation of new techniques yet to be utilized in current technologies. Universities with traditional computer science programs must now examine the need for these new multidisciplinary skills. The skills required to develop applications using the new medium are essentially varied depending on the end result. Several applications will require a programming, logical approach as a solution, whereas others will require a creative, visual solution to the problem.

The Creative Programming program looks at both the technical and aesthetic issues involved in current digital media. Education is in the form of lectures and practical sessions with skilled instructors from industry or academia. This keeps the students up to date with what is happening in industry, as well as the latest developments in research. The program helps students to combine existing skills with the technological skills needed in this area, enhancing their knowledge both technically and analytically.

Applications are examined by a jury of experts in the area of digital media, both from academia and from industry. A previous Creative Programming student is also included in the jury to provide a student's perspective. The applications include samples of artwork. The jury reduces all the initial applications to 30 possible candidates, who are then invited for interviews.

The interview generally involves an examination of the applicant's intellectual ability, creative ability, skill at handling a computer, creative thinking, and sociability. One element of the interview stage is the use of Raven's Advanced Progressive Matrices (APM) to indicate the analytical and logical ability of the applicant. It is one of the most well-examined tests of general intelligence. It is nonverbal, and it has very good psychometric properties, normalized against lots of populations with different cultural and educational backgrounds. The advanced version of the APM used in this process is designed for people of above-average intellectual capability. The test-retest reliability is over 0.93 for age 30 and under. The test is used to evaluate the logical, intellectual capacity of those very artistic applicants who have no experience in traditional computer science. As APM is a very good predictor of academic success, it can also be used to raise the confidence of the applicants.

Social and emotional competence are also judged in the interview, to make sure that an open and dynamic group process is promoted by the students accepted. Creation of this group dynamic is also the rationale for selecting a good mix of gender, age, and artistic and computational profiles into the group. In this way, the most crucial conditions for merging the artist with the programmer are met.

This new digital media program provides a place where programmers who have artistic skills and artists who have never approached a computer can learn and improve themselves. In this lab environment, different cultures experiment with new techniques and use each other as references both technically and aesthetically. Those who are familiar with computer programming educate the artists in some aspects and in return receive aesthetic criticism and tutorage.

The educational process concentrates on theoretical aspects of multimedia, design and cognition, computer graphics, modeling and animation, aesthetics, and media issues. The practical side generally involves reinforcing the theory but also acts as an avenue for the students to enhance their creativity. Industry involvement is also crucial, so the students often find that their work has a real-world purpose. Guest lectures from industry, some not directly associated with digital media but affected by its presence, allow for intellectual discussion groups to develop between creative students and the industry of which they so often have misconceptions.

Creative

Developing Creativity: A Curriculum Based on the Use of Computer Graphics Technology

This curriculum addresses the needs of people who have access to personal computers and are excited by the idea of harnessing their computers as creative tools. The objective is to empower students to explore and expand their creativity using the powerful computer graphics technology now available.

Through carefully structured assignments, students:

1. Master a powerful paint program used in combination with a pressure-sensitive graphics tablet.
2. Expand their artistic, creative, and perceptual boundaries.
3. Gain knowledge, skills, and confidence that can be applied at a personal or professional level.

Computer

Media

A number of exercises, such as those using scripting, take advantage of the unique capabilities of digital media. Other exercises can be applied to "traditional" or "natural" (non-digital) media as well as digital media. Direct tactile experience using "traditional" media helps students realize the full potential of the digital tools that emulate "natural" media.

Modality

This curriculum can be taught by an instructor in a computer classroom or online, or followed by an individual independent of an instructor. The structure presented here, and the delivery tools described (workbook, hand-outs, assignment instructions), can be adapted to suit each specific teaching situation.

My experience with all three modalities (classroom, online, and individual) has led me to value the benefit of direct, in-person, student-teacher and student-student interactions. An instructor can motivate, encourage, energize, cajole(!), challenge, and stimulate students, as well as resolve problems quickly. Students learn from each other, from seeing how others apply the tools and solve problems. Collaborative exercises teach students to let go of the need for total control, let go of treating their brush strokes and materials as "precious," and learn to work from other people's brush marks.

A classroom setting offers the opportunity for full use of the curriculum. When the curriculum is applied in an online class, or by an individual, the collaborative exercises are omitted.

Structure

The Big Picture

I start with the big picture: the map. Students are given an overview of the course and the objectives. I break down the curriculum into easily digestible parts so that no single part seems insurmountable.

The Tools

The primary digital tools are Painter software and a graphics tablet with pressure-sensitive stylus, on a Macintosh or Windows computer. There are two software options with the curriculum. The simpler option is Painter Classic, which is bundled with many Wacom tablets. Painter Classic is suitable for institutions and individuals with a limited budget who are not interested in a versatile and powerful professional-level graphics program.

For the professional user, the software is MetaCreations Painter 5.5 Web Edition (or the latest version of Painter available).

Mastering the Technology

The students' first step is getting to know and understand the tools. This is more than being able to recite a manual or point at things on the screen and know what their function is. It's about developing a deeper, intuitive sense of the way the tools work, how to approach a task, the way to make any job easy and fun and efficient. I place great emphasis on strategy; not just explaining what things are and where they are, but the series of questions students can ask themselves that naturally lead to an elegant solution. I emphasize building up good habits so that application of effective strategy becomes second nature, like the muscle-memory of a trained dancer. One of the hardest challenges in technically mastering a program like Painter is letting go of old habits and behavioral patterns.

By necessity, this section of the course is more technical and less creative. Class exercises for this section are based on experimentation and familiarization, diving in and trying out brushes and art materials, getting a feel for the variety of looks and effects at the user's disposal. This section is like getting your hands used to playing the piano keys without needing to look and think. It's the necessary hard work that needs to be put in before you can play beautiful music.

The focus of this curriculum is learning expressive creativity by hand using a (digital) paint brush. The lessons learned can be applied to other areas of computer graphics such as photographic manipulation, Web design, and animation.

Embracing Chaos

Students dive into their greatest fears: making a mess, using wild colors and wild brushes, experimenting, purposely creating chaos, using gestural movement and expression in their brush strokes (see the examples in Figure 1). Students are encouraged to enjoy the process and not worry about the end product. I want them to push all their creative blocks: the emotional resistance they may have to embracing chaos, the urge to create something "pretty," inhibitions against using the whole canvas, and the full dynamic range of the materials at their disposal. This stage involves letting go of the need to be in total control and knowing what the final result will look like. This is the moment of liberation, of being given permission to forget every "rule," every "should," every "can't." In this section, students develop vitality and energy. They experiment with brushes and materials unencumbered by the need to create something that "looks right" or resembles something.

Stillness, Connection, and Visualization

Having prepared a canvas for paint, the students take a step back from the world of technology, pixels, and programs. They select a subject to paint. They place the tablet on their laps and rest their hands, palms down, on the tablet. They close their eyes and take a deep breath. They let go of all the worries and tensions from the day. They listen. They become aware. Awareness of self is the starting point for observation.

Students open their eyes and look at their subject (this curriculum is based on drawing from the live subject). They observe carefully, being aware of what their reaction is to the subject, what stands out, what moves them. This is the moment when the composition begins in the students' minds. They begin to visualize the finished picture. Details are not important at this stage. It is more a matter of tuning in to their subject and visualizing the picture they are about to create, setting a visual goal in their minds.

Graphics

Developing Creativity: A Curriculum Based on the Use of Computer Graphics Technology

As they observe their subject, the students rest their finger tips on the tablet, following the subject's contours like invisible brushes. In this stage, students make the tactile hand-eye connection between the subject they are observing and the physical interface with the computer.

The Art of Seeing

Students are encouraged to observe with great care, to differentiate between what they "know" is there and what they actually see.

Visual Vocabulary

The exercises here build vocabulary that can communicate the impression of a 3D solid form on a 2D canvas. Every mark on the canvas affects and is affected by other marks on the canvas. It is the relationship among all these marks that collectively gives the impression of form and that determines the expressive power, beauty, and impact of the picture.

Throughout this section, students are encouraged to take risks, to push themselves and go for "overkill," to exaggerate. This course is a laboratory for experimentation. If you're going to make a gesture, make it big and bold.

Style

Students create paintings in the style of other artists, paying attention to the way different artists have applied their paint and achieved different effects (see the Picasso-style example in Figure 1). I encourage students to experiment, emulate others, and ultimately develop their own individualistic styles.

Transformation and Variation

Students experiment with transforming and recycling images, or portions of images. They learn to celebrate impermanence and change rather than be afraid of it.

Order from Chaos

This is where all the elements come together. Students learn to transform chaos into order through intuition, visualization, and observation. This is the final stage of creative realization, where contrasts are brought into harmonic balance. This is where the stillness, connection, and visualization exercises come to fruition. The circle is closed as the seeds sown earlier now make their influence felt (see Figure 2 for examples of student work).

Students complete the creative process by closing their eyes, breathing deeply, and returning to the peaceful, aware state they entered at the beginning of the curriculum.

Final Touches

Practical matters are covered such as:

- Determining canvas size
- Creating backgrounds
- Deciding when a painting is "finished"
- Frames and borders
- Preparing files for output, display, and storage
- Archive and documentation
- Output, mounting, and display issues
- Copyright protection



Figure 1
Student work (by Lesley). Clockwise from top left - Embracing Chaos, After Picasso, Self Portrait with Wild Colors, First Self Portrait.

Delivery Tools

The primary delivery tool of this curriculum will be a hard-copy workbook that contains the course description, background reading, and assignment instructions, supplemented by additional handouts that contain topical reading material. Teachers can adapt the workbook and handouts to suit their own teaching styles and circumstances. There is a substantial amount of material that could be covered, much more than there is usually time for in a regular class or workshop. Instructors will need to exercise their own judgments in selecting the portions of the course most suited to their needs and those of their students. Teachers may wish to take out specific exercises, change them slightly, or add their own. I see this as a living document, not as the final word.

Summing Up

This curriculum combines the teaching of computer graphics skills with a creatively challenging arts program. It addresses two questions:

1. How do we bring more creativity and art into our computer laboratory?
2. How do we bring computer technology into our fine art classes in a relevant and effective way?

This curriculum is not an attempt to merely duplicate a traditional art class using digital media. It brings together the mutual benefits and strengths of both worlds, the technological and the creative, in a manner that is exciting, stimulating, and educational.

About the Author

Internationally recognized award-winning artist and educator Jeremy Sutton is a graduate of Oxford University, author of *Fractal Design Painter Creative Techniques* (Hayden Books 1996) and *Total Painter* (Total Training video course 1998), and a former faculty member of the Academy of Art College and San Francisco State University. He teaches computer painting seminars and runs an online digital portraiture course.

Resources

See www.portrayals.com for a list of useful resources plus Painter tips and many examples of Jeremy Sutton's and his students' artwork.



Figure 2
Order from Chaos, Student work
The left-hand pictures are the initial images, and the right-hand pictures are the final results. Top two images by Lesley, bottom two images by Lois.

Technology

Moderator
Cameron McNall
University of California, Los Angeles

Panelists
Rebecca Allen
Mits Kataoka
Gail Swanlund
University of California, Los Angeles

Digital Design Education at UCLA

A panel of educators discusses the digitally oriented curriculum of the Design Department at the University of California, Los Angeles, and presents examples of student work.

The UCLA Design Department completely revised its program three years ago to incorporate digital technologies. Every attempt was made to maintain the core studies needed for a well-rounded designer while fully engaging the possibilities offered by new digital media and technologies. Preparation of the new curriculum involved several challenges, including:

- How can a liberal-arts based design education best prepare students for future work with newly emergent technologies and media?
- What is the best balance of theory and studio work, of conceptual to professional work, of hand-to-computer exercises?
- How does one teach with new technologies without spending too much time just learning software?
- Hardware and software acquisition
- Space planning
- How to reconcile the "computer room" mentality with the studio environment so common to art and design programs.

The panel illuminates these issues for educators who are interested in incorporating the new digital technologies into art and design programs. The presenters were all engaged in design education long before the emergence of the computer and are aware of the issues common to all programs facing the transition to digital technologies. They candidly discuss which aspects of the curriculum seem successful and which still might need refinement, and they offer predictions for the direction of digitally based art and design programs.



A number of London-based art and design institutions are in the process of developing a family of products which address the important subject of drawing. The focus of the project will be fundamentally and crucially concerned with the process of developing visual literacy in students.

In 1995-96, almost five percent of higher education students in the United Kingdom were in art and design. If related subjects with a clear interest in drawing such as architecture, engineering, and technology are included, the total rises to almost 16 percent.

Drawing is central to all that is produced within the broadest spectrum of art and design. It is the core around which the conceptual and intellectual development of students takes place. By the end of this three-year project, the proposed outcomes will be a set of fully tested, quality assured, inter-related interactive multi-media (CD-ROM, video, and Internet) products covering the important subject of drawing. The products will provide an extremely valuable teaching and learning tool for use throughout education. They will enable students to develop as independent learners while also providing a much needed, innovative and cost-effective teaching aid to support hard-pressed teaching staff. The University of Glasgow Evaluation Group will oversee the process of production, testing, and delivery.

The project will address this problem by producing a set of computer-based resources which will have wide applicability to the subject of art and design and related subjects. This applicability will also extend across a range of levels. The products will be readily integrated into the teaching and learning practices within higher education and will reduce the amount of costly lecturer time presently expended on this.

The consortium partners are: The London Institute; University of Brighton, Faculty of Art, Design and Humanities; Falmouth College of Arts; Ravensbourne College of Design and Communication; Surrey Institute of Art and Design; and University of Ulster, Faculty of Art and Design. The consortium is seeking funding of £300,000 over three years. The consortium represents a large part of the total art and design activity in the higher education sector. The subject expertise which is available to support the project is formidable. The consortium partners have significant experience in the management and delivery of complex publicly funded projects.

Introduction

The growth in higher education student numbers and in the diversity of their nature, e.g. mature students, has placed continuing strain on teaching resources. The skill base of students is more varied than ever before. It is increasingly difficult to cosset students through the development of the essential skills of drawing because of increased student numbers and obvious cost constraints, yet it is necessary to ensure that student expertise and skill in this particular area are well developed.

Drawing is central to all that is produced within the broadest spectrum of art and design. It is the core around which the conceptual and intellectual development of students takes place. Drawing allows individuals to learn to look, to record what they see, and is used to develop thought and ideas for artwork and for design, in both two and three dimensions. Furthermore it is a language capable of emotional and formal expression and communication with others, in fine art, design, and fashion, as well as in architecture and engineering. There is a rich history of drawing as part of the history of art and design as well as the history of architecture, and of technology, in relation to European and world cultures.

The working titles of the initial products will include: Ways of Looking; Perceptual and Conceptual Approaches to Drawing; Measured Drawing Systems; The Use of Perspective, Scale, and Tonality; Using the Computer as a Drawing Tool.


Learning

Aims

- To develop computer-based products which will be of significant value in providing cost effective support to students of art and design.
- To develop a family of related products which exploit the diversity of drawing.
- To enable partner institutions, either individually or in conjunction with each other, to develop an aspect of drawing.
- To produce products that will build into a comprehensive coverage of the subject of drawing and be of significant value to education.
- To subject product development to clearly defined standards covering platform, scripting, delivery, and dissemination.
- To adopt a regime of stringent editorial management, ensuring quality of content and consistency of presentation.

Objectives

- To produce computer-based products which will emphasize the development of observation, skill, and accuracy.
- To produce products which will emphasize the understanding of form and space.
- To produce products which will improve the ability to utilize software applications for 3D modeling and will enhance the teaching of formal drawing systems such as projection and perspective.
- To publish teaching materials through CD-ROMs and other media, marketed within and beyond the educational sector as valuable learning resources for a wide range of applications for which the understanding of drawing is intrinsic.



Education Delivered Through Storytelling: Using Virtual Reality as an Educational Tool

Charles Kesler
East Carolina University
keslerc@mail.ecu.edu

Playground

Diana M. Henshaw
Mike Dorsey
Alissa Chapman
David C. Balch
East Carolina University

Jeff Taylor
University of Lapland

Tom DeFanti
Dan Sandin
University of Illinois at Chicago

This use of virtual reality for delivery of education through a storytelling motif is presented by a collaborative group of educators. The project environment allows children to interact directly with a 28-foot curved virtual reality world, which at scheduled times is a huge interactive computer screen offering educational vignettes, and other times runs in demonstration mode. The choices being considered for availability to the students using the VR motif range among the following educational programs utilizing the Panoram Technologies image stitching and projection system:

- A visit to the North Carolina Outer Banks, sounds and estuaries, with information on the troubled water situation that threatens the wildlife nursery.
- A visit to the Civil War vessel SS Monitor, sunk off the North Carolina coast.
- A visit to the pirate ship Queen Ann's Revenge.
- A visit to a museum without walls housing the largest collection of African art outside the Smithsonian Institution.
- A visit to the history of North Carolina pottery.
- The folklore of Lapland.
- A wild ride on a snowmobile through the great northern woods of Lapland.
- A history of the Cold War and pursuit of the elimination of nuclear weapons, featuring past SIGGRAPH Computer Animation Festival works ("Nuke the Duke" and "Wag the Flag," for example).
- A story about the plight of endangered species in "Save the Animals."
- Health education

Storytelling

Workshop

"If Walt Disney were alive today,
where would he be?"

"Probably brushing up on his
computer skills!"

This excerpt from a recent article in the "Toronto Globe & Mail" sets the stage for this workshop. The reporter interviewed a number of educators as well as production managers and technical directors in the film industry and included many quotes, such as:

"Today's major animation studios and boutiques are insisting that employees be hybrid artists - individuals possessing both Bill Gates' technical proficiency and Vincent Van Gogh's artistic flare. It is a recent phenomenon that animators can not afford to ignore."

Fueled by Hollywood's demands, educational programs have begun to surface at a number of colleges and universities to help provide students with the skill sets necessary to meet these demands (both computer-related and technical). Many artists and would-be animators are now investing a lot of time and effort in seeking out educational programs that will allow them to learn and develop these skills.

But why wait until the college and university years to introduce this technology to hungry, creative minds? Why not introduce this technology at a much earlier stage so it can grow, develop, and mature prior to the time when one has to make career decisions?

A decade ago, general computer skills were introduced only at the college and university level, largely because of the high cost of computer hardware and software, and the relatively low budgets at the public school level. Today, with the availability of lower cost hardware and software, and in many cases donated software, kids have access to computer technology at a very early age. We now have young adults entering post-secondary programs who are fully prepared to study computer skills and applications at a much more advanced level. As a result, they have very bright futures ahead of them, in positions that are not only challenging and creative but very lucrative as well. These graduates are now designing the software of the future and creating award-winning effects in the movie industry.

Today, we recognize that this integration of computer technology into the classroom has been a great benefit, not only as a teaching tool but also as a doorway to creativity and exploration. Simply put, it's fun as well as challenging. And it is this combination of fun and challenge that keeps these students motivated and eager to learn and grow.

This workshop introduces educators to the world of high-end 3D computer graphics technology and demonstrates how easily it can be integrated into the curriculum and classroom at the public and high school levels, using real software that is recognized for its abilities to create photo-realistic models, award-winning special effects, and truly believable animation. It utilizes some of the most innovative and flexible software available in the market today: Houdini, from Side Effects Software. The technology is described in lay terms as participants work through a real-life model project. Full course notes on the guided exercise are provided for all attendees.

The Scenario

A designer has an idea for a spiral staircase for an architectural design presentation to a client. The designer is a bit sketchy on the design but he does know the rise, the number of treads necessary, and the degree of rotation that the staircase must travel through. He is looking to the computer designer for a few ideas. He needs to see a prototype, but he also wants a considerable amount of control over design changes.

Workshop participants build the spiral staircase using a number of mathematical and interactive manipulation techniques. Materials and appropriate lighting are applied for final presentation of the prototype to the designer.

What computer-based learning environment will students of all ages be immersed in five, 10, or 20 years from now? And how can we best prepare for learning and teaching in that environment? These questions may be partly answered by obtaining a deep understanding of Web-based (or whatever succeeds it) interactive learning that engages a broad range of human capabilities.

The Exploratories project, an interactive Web-based educational research project at Brown University, uses the introductory undergraduate computer graphics course as a testbed to address the following challenges:

- What goes into effective interactive Web-based learning modules?
- How can such tools be most effectively deployed?
- What categories of the learning environment can benefit most from this approach?

Exploratories

An exploratory is a computer-based combination of an exploratorium and a laboratory that embodies an approach to learning by experimentation and investigation, in the tradition of Kay¹ and Papert.² It provides multifaceted, interactive microworlds that model objects, phenomena, and concepts, and that exhibit appropriate behaviors when interacting with the student.

Our current implementation embeds Java applets in a hypermedia framework that (a) provides an immediate context for both the individual applet and for the larger conceptual frame of reference, and (b) enables students and teachers either to work within the framework provided or to place the applet in an environment of their own choosing.

For example, one set of exploratories, which develops an intuitive feeling for the signal-processing aspects of color vision,^{3,4} has been used by teachers around the world within their pre-existing curricula. Another set of exploratories, which teaches basic concepts of additive and subtractive color mixing,⁴ is being used in our computer graphics course as classroom demonstrations, in assignments, and for self-directed learning.

Project Context

The context of the Exploratory project both reflects and reinforces the flexibility of the project approach described above. Because the Brown Computer Science department has long-standing programs for undergraduate research and teaching assistantships, our resources include eager young developers who combine imagination, energy, and intelligence with a sound technical base. These students work under the guidance of faculty, staff, and older students who provide domain and development expertise. The constraints include the heterogeneity of student-developer experience and platforms, the multiple ways in which we use the exploratories, and the episodic (and sometimes unpredictable) nature of undergraduate availability.

The Design Strategy Handbook

The ultimate goal of the project is to understand the patterns, in Alexander's sense⁵ that underly effective Web-based interactive education and to transmit that understanding through an interactive, Web-based Design Strategy Handbook that is itself an example of what it teaches. This handbook will present the results of our research in the form of guidelines, rules-of-thumb, worksheets, and patterns to guide teachers and educational material developers through a process of analysis of their context, needs, and resources; development of strategies and plans; and finally deployment of effective, finely tuned materials.

References

1. Alan Kay and Adele Goldberg. Personal Dynamic Media. IEEE Computer, March 1977.
2. Seymour Papert. *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books, 1980.
3. Jeff Beall, Adam Doppelt, & John Hughes. *Developing an Interactive Illustration: Using Java and the Web to Make it Worthwhile*. In Proceedings of 3D and Multimedia on the Internet, WWW and Networks. April 1996.
4. Exploratory Color Theory ColorWeb: www.cs.brown.edu/exploratory/ColorWeb/
5. Christopher Alexander. *A Pattern Language*. Oxford University Press, 1977.
6. Virtual Laboratory: jersey.uoregon.edu/vlab/
7. Educational Fusion: lumina.lcs.mit.edu/EduFuse/
9. Daniel L. Gould, Rosemary M. Simpson, and Andries van Dam. *Granularity in the Design of Interactive Illustrations*. Proceedings of SigCSE '99. New Orleans, March 1999.
8. Reinhard Klein and L. Miguel Encarnacao. *A Web-based Framework for the Complete Integration of Teaching Concepts and Media in Computer Graphics Education*. In Proceedings of ED-MEDIA '97. Calgary, Canada, June 1997.

Handbook Topic Examples

We are currently mining our own experiences as well as those of other Web-based educational project designers^{6,7,8} for patterns in the areas of user interface design, pedagogy, interactivity, Web-based information structure design, and component-based software. Our initial analysis has focused on optimal applet granularities⁹ and on the process through which we develop user interface widgets.

Granularity


Granularity, by which we mean the conceptual scope of an exploratory, has significant impact on the rest of the design process; fine granularity usually leads to multiple, single-featured applets, whereas coarse granularity tends to produce complete applications that demonstrate several related ideas. We have found that the fine-grained approach has many benefits for reducing development costs, enhancing reuse, and supporting flexible teaching and learning strategies.⁹

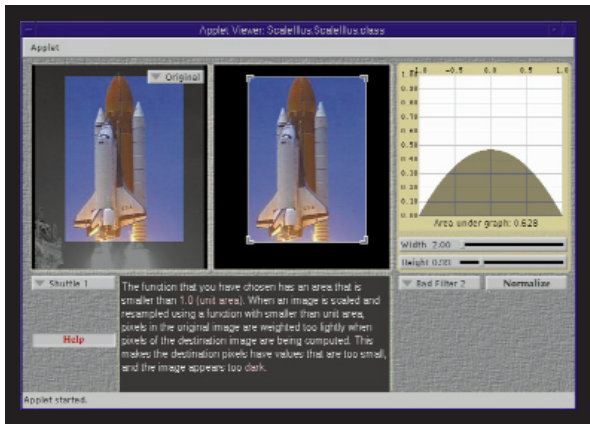
User Interface Widgets

As an example of the process of developing user interface widgets, we have documented the entire sequence of decisions, results, and analyses involved in deciding what kind of widgets to use to represent vectors for a set of basic linear algebra exploratories. This documentation will be exploited later in looking for patterns of effective UI design for different categories of exploratories, and is part of the very preliminary version of the Design Strategy Handbook.

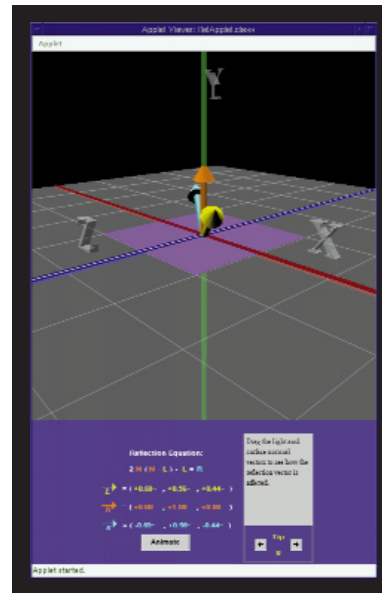
Future Work

The fluid, asynchronous, fragmented, and open nature of the Web as an environment forces us out of the fixed formats we have been used to in our classrooms and labs, and provides invaluable experience in rethinking educational goals, strategies, materials, place, time, and space. Our hope is that some of the patterns and strategies we develop in response to this environment may prove flexible enough to provide a foundation for working with the radically changed environments that students, of all ages, will be immersed in five, 10, or 20 years from now. To further this goal, we will: extend our testbed to other courses in computer science, the sciences, and the humanities; continue to investigate the work of others in this area;^{6,7,8} and explore adaptations of our current work to semi-immersive environments, such as the Responsive Workbench, and fully-immersive environments, such as the Cave.

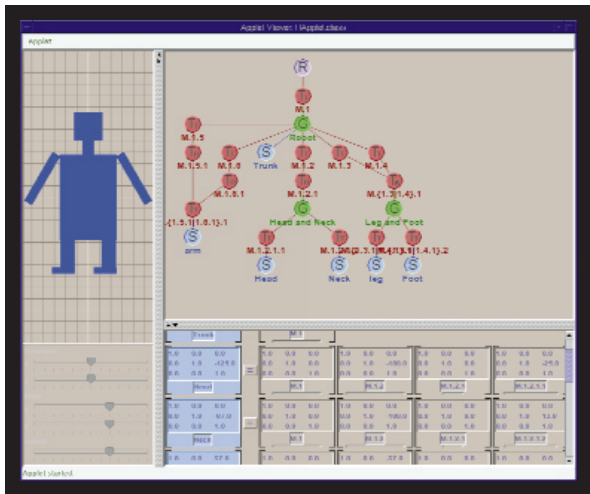




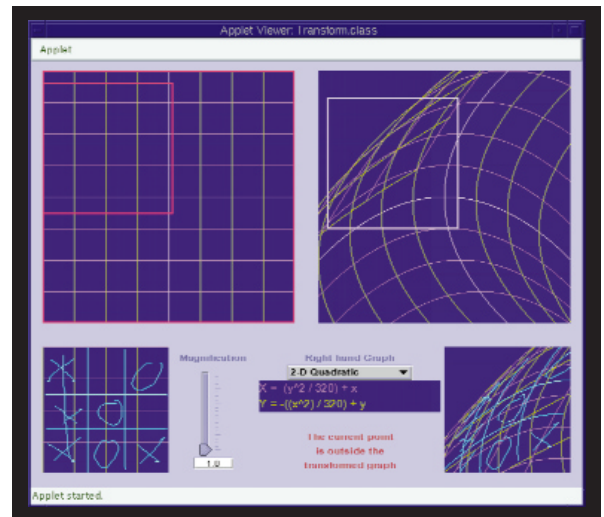
Scaling and Filter Shapes



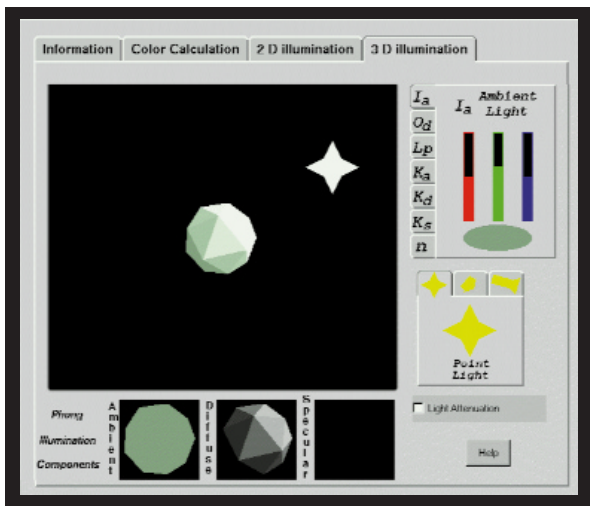
Reflection Equation



Scene Graphs



Coordinate Transformation, 2D Quadratic



Phong Illumination



Halftone

FELIX 3D Display

Implementation of a "holodeck-like" display device is a long-held dream of researchers in computer graphics and interactive techniques. The FELIX 3D Display represents another step towards this goal. Projected spatial images occupy a physical space of full height, width, and depth. The system can be viewed by any number of spectators from almost any angle in the room without cumbersome goggles. The images are created within the physical world of the observer, in contrast to virtual reality systems where the observer is placed within the virtual environment of the non-physical image space. Thus FELIX provides an important new option for team tasks and tasks requiring many simultaneous views of multi-dimensional data. Entertainment, air traffic control, molecular design, or computer-aided design are only a few examples of potential applications.

In the FELIX system, a two-dimensional pattern is generated on a passive projection screen. A rotating helix sweeps out a cylindrical envelope (frame rate is 20 Hz), providing a volumetric display medium through which scanned laser pulses are projected. The hitting laser beam is scattered from the rotating surface, which generates a visible light spot (voxel). The spatial position of the emanating voxel within the display is determined by the momentary location of the laser beam's intersection with the rotating helix.

Vector Graphics

The vector graphics are projected with galvanometric scanners. Colored images are realized by combining red, green, and blue lasers. Separate modulation of each component enables additional color mixing.

Random Access Graphics

In the random access approach, raster images are generated. But only the volume pixels (voxels) to be displayed are scanned, saving data resources. For this mode, a high-speed acousto-optic deflector is used. The advanced software concept is capable of displaying wire-frame graphics in standard data formats within the FELIX 3D Display. The resolution is about 10,000 voxels at 20 Hz.

The FELIX 3D Display project team has evolved from a scientific working group of students and teachers at the Vincent Lubeck High School in Stade, Germany. This group works together on projects that range from development of a cyberbike to multivision projections through development of the FELIX 3D Display. All these voluntary activities are aimed at combining aspects of science and art.

Despite minimal financial resources, the group has achieved considerable results. Since 1983, the FELIX 3D Display team has been working on various 3D display techniques (stereoscopy, holography, multiplanar displays, etc.) in collaboration with the Institute of Flight Guidance and Control at the Technical University of Braunschweig, which led to development of innovative volumetric 3D display concepts.

FELIX



Figures of Speech

Playground

Speech

What do you get when you put a group of 3D computer animation students together with an animation instructor and a great concept? Figures of Speech, an eight-minute animation that humorously depicts interesting colloquialisms of various origins. Imagine really seeing what happens when the "cat's got your tongue" or what it might be like when "it" really hits the fan.

Overview and Concept Development

The English language is full of creative associations from diverse origins. These catch phrases and colloquialisms are obvious because they translate poorly into other languages. The sheer volume of commonly used English figures of speech is quite remarkable. Choosing a commonly used figure of speech and successfully depicting it in a non-obvious manner is quite a task.

Figures of Speech follows the process of preparing a body of student animations in a form and quality suitable for screening at festivals and contests. Teaching an individual to reach an applied understanding of 3D computer animation is a delicate balance of several skill sets. Teaching an individual to animate, problem-solve, and specialize, as well as tell a story, is critical in an industry that expects high-caliber work. A mix of correct schooling in classical technique, practical use of newly introduced animation tools, story and concept development, and multiple critiquing sessions is necessary for a successful performance.

Real World Critiquing

Instructors will find it helpful to integrate the client/art director/producer relationship in the classroom. There's a feedback loop that's essential to producing good work. It is extremely useful to show students the ropes using a real-world approach. The limitation of this experience (as opposed to the real world) is that the student completes the work from start to finish.

In a small studio, the critiquing begins almost immediately. The same is true in the classroom. Ideas are deemed feasible or not, and the scope of the project is limited to an attainable amount of work. Multiple critique and feedback sessions are essential to push students beyond their current aesthetic and practical understandings of animation.

In the real world, you get feedback from the client, who expects a product to be developed on time and on budget. The producer's role is to see that the work is done according to plan and that resources are properly allocated. The art director oversees the aesthetic of the project. As an instructor, you get to play all three roles. And critiques are done with all three job descriptions in mind.

The feedback loop is started from day one of the project. A log of the feedback from each of the critiques helps students logically address suggested changes. During the second critique, a preliminary grade is awarded. Students can elect to make the additional changes before beginning their final render. A grade is given based upon completion of those changes. The final render delivery date is set, and students are penalized, or graded-down, if the deadline is not met.

This panel presents recommendations for the future of computer graphics education developed at the Computer Graphics and Visualization Education Workshop (GVE 99), co-sponsored by Eurographics and SIGGRAPH, 3-5 July in Coimbra, Portugal. The GVE 99 workshop brought together many of the leading computer graphics educators from around the world to consider the future of computer graphics education. Their recommendations will influence many ongoing efforts in developing curricula, including a developing ACM/IEEE computer science curriculum study and a new effort in creating curricula for computer graphics in engineering. More information on GVE99 can be found at: <http://www.ccg.uc.pt/GVE99/>

The focus for GVE 99 was two-fold:

1. Teaching computer graphics and visualization.
2. Using computer graphics and visualization for education.

Participants in GVE 99 were selected from position papers reviewed by a distinguished international program committee. While some participants were invited to expand their position papers into papers to be presented at the workshop, considerable time was spent in organized discussions of future directions in computer graphics and visualization education. The recommendations and predictions from the workshop were distributed at both SIGGRAPH 99 and Eurographics 99 as part of the joint celebration of SIGGRAPH's 30th year and Eurographics' 20th year.

Panel topics include:

- Hardware and software technology for educational workstations.
- Tools for educational applications.
- Cognitive aspects of visual teaching and learning.
- New teaching paradigms such as interactive classrooms, tele-immersion, and networked tele-collaboration.
- Developments in specific academic areas such as computer science, engineering, and art.

Panel Chair

Michael B. McGrath
Colorado School of Mines
Position Statement: Engineering Education in the Future

The practice of engineering has been radically changed by graphical tools that model the real world and allow interactive simulation. The challenge will be to develop teaching, instructional materials, and curriculum, that will take advantage of these tools. The emphasis in engineering will be on using graphics in education.

Panelists

Werner Hansmann
University of Hamburg
Position Statement: Computer Science Education in the Future

It can hardly be imagined that there could be any discipline that has no need whatsoever for graphics as a powerful means of communication. There is an increasing demand for graphical tools to become more and more sophisticated. Hence the challenge for computer graphics curricula (mostly as part of a general computer science education) is to provide appropriate foundations that will enable future developers of graphical tools to meet that demand.

Education

Dena E. Eber

Bowling Green State University

Position Statement: Educating Artists in a Digital Medium

Against all odds (and budgets), computers have recently emerged on the scene in university art departments across America. As part of this rapid change, unique digital art forms are taking shape. From imaging, digital painting, and 3D modeling to interactivity, video, and 3D animation, graphic technologies in the arts have the power to be an expressive medium, thus forging new educational and aesthetic concerns for computer art educators. Among many issues, three emerge as paramount. Art educators must:

1. Provide a balanced curriculum that includes traditional art, a computer art focus, and computer science.
2. Teach beyond what I call the "wow" factor.
3. Help students use the technology to express themselves through works of art.

Judith R. Brown

University of Iowa

Position Statement: Enabling Educational Collaboration


Tele-immersive collaboration can provide new learning experiences and insights by bringing together educators and students from remote locations in a shared virtual space. The visualization lab has long been the crossroads between research and education. As virtual reality technologies have recently enabled scientists to better understand their data through immersive exploration, so are these technologies enabling students to better understand new concepts. Inter-university student explorations are an exciting new way to learn. Enhanced computer graphics and tools for interaction, along with high-performance networks, enable a shared learning environment.

Jose Carlos Teixeira

University of Coimbra

The emergence of multimedia technologies and tools, the Internet, and the huge volume of educational information require new skills both for educators and students. Therefore, computer graphics is a crucial ingredient for the new generation of computer-supported learning material. In recent years, different approaches have been developed to answer technological and pedagogic challenges in education and training. It is important to review issues considered in some specific educational projects and apply them to different learning scenarios.

A similar panel is scheduled for Eurographics 99 in Milano, Italy, in September, for dissemination to and input from Eurographics educators.



What does it take to get a job at a visual effects, traditional animation or interactive company? This course presents the keys to opening the door to interviews, how to put your life on a one-page resume, and how to showcase your talent in a three-minute-or-less demo reel.

Resumes

If yours doesn't work, neither will you. As a recruiter and career coach I have seen literally thousands of resumes. On a slow day, I get three or four; so here is what bugs me the most:

Issue #1 - No phone number, wrong phone number, wrong area code, hard-to-find phone number, hard-to-read phone number.

Issue #2 - Name missing. Yes it has happened! But if I have a phone number, I can call it and leave a message for the person, so this gets the #2 position.

Issue #3 - Resumes with a type face that is impossible to read – too small or too ornate. Huge blocks of type that are a challenge to read.

Issue #4 - Hiding your skills. Don't make anyone read through a big paragraph to find your specialized skills such as C++ programming or knowledge of Alias. Highlight your skills under a separate heading.

Issue #5 - Resumes with multiple pages. If your resume is more than one page, put your name and phone number on each one.

Issue #6 - Resumes that fail to tell me who you are, what you know (skills), what you've done (accomplishments), and what you want to do (objective or goal). If you are changing careers, focus your resume on the job you want rather than the job you have.

Issue #7 - Paper that doesn't copy well. Test your resume. Copy it and make a copy of the copy. Surprised? Orange and dark blue paper turns black. Marbleized paper makes your resume look like someone poured coffee over it.

Issue #8 - Graphics or artwork on a gray scale behind the type. After doing the copy test, you'll find those beautiful graphics in the background are now some of the ugliest stuff you've seen on paper and what's more, you can no longer read your phone number or name, which looked so crisp in front of the graphic on the original. If you want someone to get a sample of your graphics include it on a separate page with your name and phone number.

Issue #9 - Typos and spelling mistakes. Proofread and ask a friend to proofread.

Here are several tips for a better resume:

1. List your skills and be specific.
2. Many companies scan resumes into computer databases. Select a font where the lower case l and number one are different enough that the computer won't confuse the characters.
3. If you have email, put your email address on your resume.
4. If your resume shows a variety of jobs, make sure you have an objective at the top that indicates what job you're seeking.
5. Review your resume every six months to update your skills and accomplishments.

Portfolios and Demo Reels

If you are an artist, it is essential that you have an outstanding portfolio and demo reel. The first step is to determine what your strengths and interests are. What kind of work are you suited for? There are many different jobs for artists, from animators to modelers to graphic designers to Web site developers to interface designers. You need to figure out what you like to do and what you are really good at. Assess your skills. Make sure your demo reel reflects the very best you can do and keep it short. Make them want to see more. Make sure the demo reel and portfolio are relevant to the job you want. If you want a job as a character animator, don't show only compositing work on your reel.

The purpose of the resume, portfolio, and demo reel is to get you an interview with someone who can hire you. Prepare your marketing materials with care. Have others take a look at them and give you feedback before you send them out.

For artists, a demo reel and portfolio are more important than a resume.

Your demo reel should:

- Be no longer than three minutes. It can be much shorter.
- Show variety
- Contain only your best work
- Be dynamic
- Be irresistible
- Be labeled with your name and phone number and email address if you have one. Include slates on your reel with this information in case the label falls off.
- Be a VHS cassette in NTSC format. This is the format almost all companies can deal with in the United States. If it's a PAL tape, be sure the company has a way to view it.
- Be representative of your recent work and show your skills and talent
- Be of high caliber and quality

Put the very best segment first. Include slates on the tape or a written outline that describes each scene and what you did for that segment. Remember, your audience sees lots of demo reels and portfolios. Keep it moving.

If you must have your work returned, include a self-addressed stamped container for return. Never send your only copy to anyone.

If you have worked on an interactive project and want to submit your portfolio in a digital medium such as CD-ROM, call the company before you send it to be sure they have the appropriate equipment to view it. Include a breakdown of how each piece was done and the constraints of production.

A portfolio of life drawing, illustration, photography (if you are interested in lighting), sculpture (if you are interested in modeling), character design, or color design is a big plus. Many

aspiring computer artists today have no foundation in fine art, and the lack of training in aesthetics limits their capabilities. It's easier to train someone to learn a software package than to learn to draw. If you have a fine art background, include some of the work with your reel. Portfolios should have no more than 25 pages of work, and remember to include only your best work.

Whether you submit a demo reel, CD-ROM, portfolio, or all three, remember to always include a resume with it.

Networking

I recently got a message in a fortune cookie: "A wise man knows everything. A shrewd one, everybody." That message is the essence of networking. As an independent recruiter, I have found that no matter what you do in the entertainment arena, networking is key.

Here are some networking tips to try out at the next function you attend:

- If you have trouble getting started, think of it as a game. Make a goal of meeting at least two people at the next party or meeting you attend. Years ago, I went to a party with three friends, one of whom issued us all a challenge: meet five people. Instead of hanging around together, we went off in all directions and reported back the results. We had all met five different people, so now our network had expanded by 20!
- If you are painfully shy, go to events with someone who is good at networking. He or she will take you around and introduce both of you to someone new.
- Listen and learn. Force yourself to eavesdrop on a group. Learn their names.
- If you forget someone's name, admit it and reintroduce yourself. If you dread doing that, if you have a friend with you, reintroduce yourself to the person and then introduce your friend. Then pause so the person can introduce himself.
- Be prepared to meet people, follow up, and keep in touch. Bring plenty of business cards and exchange them with everyone.
- In a group made up of strangers and acquaintances, talk to someone you don't know. Once you introduce yourself to a stranger, he or she is now an acquaintance and could be part of your network.
- You have something in common with everyone. Make it your goal to find out what it is. This is fairly easy to do especially at SIGGRAPH 99. Everyone here has a common interest.
- Never whine, gossip, or speak badly of a fellow artist or employer. Be nice to everyone. It's a small world, especially in the entertainment industry.
- Prepare for meetings by reading the trade journals or the program. Read the bios of the speakers who do the presentations before you attend the meeting. Find out about the people you are going to meet. Do your homework. It'll be easier to speak to people if you know something about them.

Get a Job!

- You don't have to wait for an event to try networking. Form a relationship with people in charge. Go to lunch with the boss. Network with people on other projects at your company. Network with people from other companies, too. Your next job may come from one of them.
- Be positive and flexible. Be a team player.
- There is no such thing as a small job. Do your best on every job you get and your circle of fans will grow.
- Everyone, yes, everyone is a potential job lead. Don't keep what you want a secret. Tell people what you are looking for. Ask them for help.
- The most important thing about networking is you must be prepared to give before you get. Find out what you can do for someone else. Perhaps someone is having back trouble (not uncommon in the animation industry!), and you know a good chiropractor or acupuncturist. Be ready to lend a hand, and hands will reach out to help you when you need it.

Whenever you attend an industry conference, trade show, or an association meeting or software user group, make a goal of meeting at least five people. Networking is one way to market yourself for jobs that may never be advertised. Build on these relationships.

Interviewing

Before you go for your interview:

- Research the company's products or services.
- Find out what their reputation is. Larger companies will have publicity materials. Study their press releases.
- Look for people in responsible positions that you respect.
- Determine the long-range prospects of the company by looking at their goals and target market.
- Look for a company that is growing.
- If it is a publicly traded company, check it out on the stock market report.
- Most importantly, look for opportunities to learn from exceptional people.

This research will pay off when you finally get an interview. You will impress them by your knowledge and enthusiasm for their company. Everyone wants to find someone who will fit in, and this research will help you convince the interviewer that you do.

After you get one foot in the door, how do you get both feet in and stay in? Interviewers want to answer three basic questions:

- Can you do the job?
- Will you do the job?
- Will you fit in with the other people at the company?

To get answers to these three questions, they may ask you many other questions. If you can convince the interviewer that you have the skills, (can you do it?); the willingness (will you do it?); and flexibility, and that you are a team player (fit), chances are you will get a job.

The last question (fit) is often the one that makes all the difference. If you have done your homework and research, you will likely be the one they believe will fit in. You know the other companies that do the kind of work they do. You know the company's work and something about the company's history and the key people. And you can speak with confidence.

But only one of the key factors in getting a job is the fit. The other is willingness. Attitude is of paramount importance, particularly in the computer graphics industry. There is little room for arrogance or a prima donna attitude. You are working closely with many personalities under often strict deadlines, and being a team player is essential to keeping a job and sustaining your career. So convince them that you can do the job, want the job and are willing to do it, and that you will fit in, and you will be getting an offer soon.

Once you get that job, remember to network, keep your resume, portfolio and demo reel up to date, and do the very best job you can every single day.

The Author

Pamela Kleibrink Thompson is an independent recruiter, career coach, and management consultant for clients such as Walt Disney Feature Animation, Digital Domain, Fox Feature Animation, Dream Quest Images, and interactive companies such as Lucas Learning, Activision, and Hollywood On Line. Thompson's adventures in animation include setting up a studio from scratch, hiring and managing a crew of 100 artists, creating and staffing a 3D art department, producing award-winning video games, commercials, creating and producing the 1998 Career Boot Camp, and co-producing the 1999 Career Boot Camp.

Thompson regularly speaks on career issues and consults with colleges and universities to design animation training programs. Her production credits include "The Simpsons," "Family Dog," and "Bebe's Kids." Thompson worked on the computer animated film currently running at Epcot Center's Universe of Energy Pavilion, which is possibly the longest continuously running computer animated film ever. She was recognized in February 1999 by Animation Magazine as one of the top recruiters in her field. Her articles on animation and business have appeared in over 40 periodicals. A member of the Academy of Television Arts and Sciences, Thompson is currently the subcommittee chair of LA SIGGRAPH and is active in Women in Animation and ASIFA.

Going Farther in Less Time: Responding to Change in Introductory Graphics Courses

Moderator

Rosalee Wolfe
DePaul University

Panelists

Steve Cunningham
California State University, Stanislaus

Scott Grissom

University of Illinois at Springfield

Lew Hitchner

California Polytechnic University

The field of computer graphics has matured greatly since the formal statement of the introductory undergraduate course was created for ACM/IEEE Curriculum 91, and courses must change accordingly. This panel will describe a philosophical basis for the changes and gives some examples of courses that are responding to that change.

Introduction

The panelists all teach computer graphics at medium-size institutions and teach courses whose details vary rather widely. However, all have addressed the changes in the field, and their discussions of their choices illuminate how the changes in the field are reflected in course design. The panel shows how several recurring themes appear in the courses.

One of the goals of the panel is to lay out an early form of a philosophy for the introductory graphics course. We hope that this philosophy will evolve into a basis upon which instructors can develop courses that fit their local needs while reflecting the changes in the field. Fundamentally, the philosophy is:

1. Computer graphics is inherently 3D and courses should be also.
2. The fundamental subject of a computer graphics course is geometry and how it is expressed in computational terms. Thus, geometry is a major part of the introductory course. Geometry is expressed in terms appropriate to the field, such as coordinate systems, transformations, and surface normals. The basic shape is the triangle. The mathematics of curved surfaces is typically treated in a more advanced course.
3. Computer graphics is intrinsically visual, and even the most technically oriented graphics practitioner must be aware of the visual effects of algorithms. Unlike other areas of computer science, algorithms must be considered not only for time and memory usage, but also for their visual effect.
4. Besides geometry, computer graphics is about light and surfaces, and about developing algorithms to simulate their interplay. Courses need to include material about light and surface properties, and about the distinction between the ways various algorithms present light and surfaces visually.
5. Computer graphics has matured to a state in which there is a small number of high-level API's that support all the fundamental concepts needed for early work. Courses should be built upon this high-level approach.
6. Computer graphics should be interactive. Courses should include interactive projects and cover event-driven programming. Because this explores new possibilities in curricula for computer graphics courses, this panel is designed to spark discussion and encourage involvement in this process.

Panel Members and Position Statements

Steve Cunningham

Steve Cunningham has taught at California State University, Stanislaus since 1982. He has worked in computer graphics since he developed a graphics-based general statistics laboratory in 1976-78. He has worked with SIGGRAPH since 1983 as chair of the Education Committee, chair of the SIGGRAPH 91 Educators Program, Director for Publications, and currently as Chair of the organization. He has presented two SIGGRAPH conference courses, co-authored books on user interfaces and electronic publishing, and co-edited two books on visualization and one on object-oriented computer graphics.

"In the late 1980s I contributed to Curriculum 91 by reviewing its recommendations in computer graphics and user interaction. As part of that review, I wrote the computer graphics course outline in that curriculum. However, the field has changed substantially and our graphics courses must keep up. I changed my introductory course to become fundamentally a course in computer graphics programming based on OpenGL. This

change showed me that students could succeed with less computer science preparation when they use a capable API. This makes the course accessible to more students, and I have begun to orient the course towards students in the sciences. I am now developing materials for such students. It is still important to offer a computer graphics fundamentals course, particularly for students who want to pursue a graphics career. Because the fundamentals course builds on students' OpenGL background when it is a second course, it can cover significantly more material than when it is the first graphics course."

Scott Grissom

Scott Grissom has been teaching at the University of Illinois at Springfield since 1993. He graduated from The Ohio State University with an emphasis on computer graphics and human-computer interaction. He is editor of the Visualization Resource Center, a collection of peer-reviewed teaching resources for computer science. He teaches Data Structures, Computer Science II, Object-Oriented Design, and Human-Computer Interaction in addition to the Computer Graphics course.

"We only offer one undergraduate course in computer graphics. So I try to expose students to a wide range of computer graphics. I have been using C++ and OpenGL for three years. I want students to create interesting and motivating images as early as possible. Using a high-level API allows them to do that. Towards the end of the semester, I briefly introduce concepts of ray tracing and have students use POV-Ray to render an image. POV-Ray requires an understanding of lighting models, view manipulation, and texture mapping, and is available on all platforms. The final project involves an interactive application on the Internet using JavaScript, CGI, VRML, or Java."

Lew Hitchner

Lew Hitchner has a varied background in applying computer graphics and virtual reality in academia, industry, and research. His PhD is from the University of Utah. He taught at the University of California, Santa Cruz for five years and currently teaches at California Polytechnic University. He has worked as a researcher for NASA, in industrial R&D, and as a private consultant in graphics and VR. His teaching experience includes introductory CS1 and CS2 and all levels of computer graphics courses (intro through graduate courses).

"The Cal Poly introductory computer graphics course is a practice-oriented curriculum that combines fundamentals with intensive laboratory exercises and programming assignments. Students learn to apply fundamentals through programs that use two high-level API's. Our one-quarter course covers 3D geometry and transformation basics, event-driven interaction, hierarchical modeling, camera and lighting equations, and rendering techniques (visible surface, texture mapping, etc.). Students use a high-level API in all assignments (Open GL and Open Inventor)."

Rosalee Wolfe

Rosalee Wolfe has taught graphics and human-computer interaction at DePaul University since 1987, after graduating from Indiana University. She served as the SIGGRAPH Educators Program Chair in 1996 and 1997. She has edited several SIGGRAPH Slide Sets and the Seminal Graphics book published in 1998. In addition, she is the education columnist for Computer Graphics and currently serves on the SIGGRAPH Education Committee.

"We have two introductory computer graphics courses. Computer Graphics Survey covers the entire discipline and uses high-level packages (Rhino, POV-Ray) to teach topics from the areas listed above, and to teach animation. This course is often referred to as the seduction course, because students taking this course often decide to embark on additional courses in graphics. The second course, Computer Graphics I, uses C++/OpenGL as a platform. Students are given a crippled wireframe browser, to which they add transformations, hidden surface removal, shading, texture mapping, and interactive elements. Although both are entry-level courses, many find that the survey course helps them when they take Graphics I."

Change

Hands-On Animation

Introduction to Animation using Maya 2.0

3D animation is an important subject in the film and video market. Everyone wants to animate or have some form of 3D animation in their productions. There are numerous tools available in the production environment for 3D animation, and they all work from similar concepts. In this session we will look at how one software system (Alias|Wavefront's Maya) handles the concepts of animation.

Maya Overview

Maya's Architecture

- A brief look at how objects are defined in Maya with nodes and attributes.
- Use of the dependency graph view to look at the relationship between nodes.

Animation in Maya

To animate something is to bring it to life. A simple cube can be considered a character and brought to life by a good animator. The animator is able to give it personality by carefully animating the cube's movements with subtle nuances. We no longer see the cube as an inanimate object but as a character with personality. In Maya, there are several different ways in which you can animate your scenes and the characters who inhabit them. In this session we concentrate on the basic tools used for animation.

Keyframes

To understand keyframes, you need to look at time. Within Maya, you are able to define what time is by using a frame rate. A certain number of frames will define one second of time (for example, 30 frames per second for video). Once you have defined your frame rate, you know how many frames make up a second (in this example 30). This is important because a keyframe is a key set on a specific frame. The key is specific information about the object that is held at that frame, such as position, scale, and rotation. At a different frame, the key can hold different information. If you wanted an object to move from position A to position B in one second, then you would put one keyframe on the object at position A, frame one and a second keyframe on the object at position B, frame 30. The system now has two keyframes that define the object's position at two different times. Since the positions are different from the first frame to the second, when the system moves to frame two, the object starts to move. As you progress through the frames, the object will continue to move until it reaches the position at frame 30 defined by the keyframe. When this animation is played back, the object moves from position A to position B in one second. The tools we will be looking at are:

- Set key
- Auto key
- Set breakdown

Editing Keyframes

When animating, you will generally take a first pass at an animation by blocking out positions. This is a rough first pass of an animation, setting initial keyframes to establish positions of your objects. Once you have established initial positions and keyframes, you will need to spend time editing them by moving keyframes and inserting new ones. There are a number of tools and windows that can help you when it comes to editing keyframes:

- Time slider
- Graph editor
- Dope sheet
- Curve tangents

Hierarchical Animation

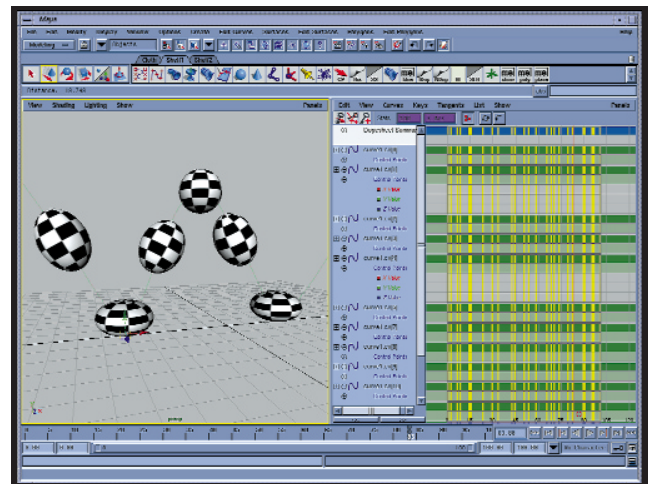
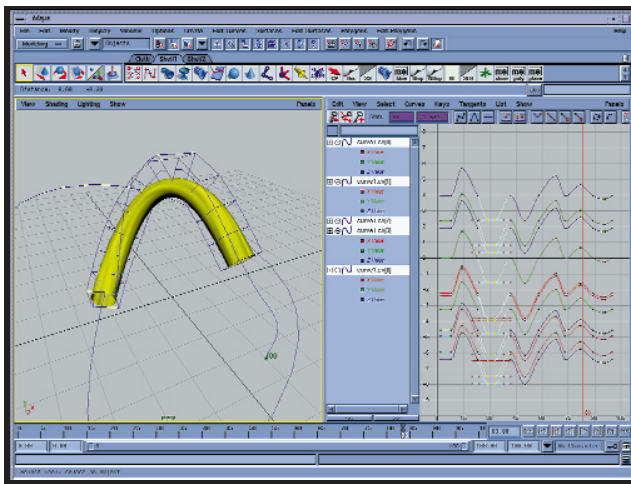
Hierarchical animation is used to add secondary animation to an object. For example, if you animated a bouncing ball, you might want to introduce some squash and stretch as well as some rotation. If you were animating a car, you would want the tires to rotate and the engine to run. To make these examples work effectively, you need to use a hierarchy.

Set Driven Key

This is a powerful tool that lets you connect the animation from one object to another object or objects. You can quickly set up an animation where you keyframe one object and use that keyframed motion to drive the movements of one or more objects.

Path Animation

Another important animation tool is the path tool, which allows you to define the motion of your object by using a path. First you need to draw a curve defining where your object will travel, and then you use the path tool to attach the object to that curve (for example, a car on a roller coaster). To further modify the animation of the object, you would then modify the curve and the keyframes generated along its path.



Presenters
John Refling
Carl Pennypacker
Lawrence Berkeley National Laboratory

Hands-On Universe: Teaching Astronomy with Java-Based Image Processing Tools

The goal of the Hands-On Universe Project is to teach observational astronomy and related science concepts within a collaborative and stimulating learning environment. This environment is enhanced by appropriate software tools and professionally produced curricula. The software tools consist of Web-based image processing, visualization, and computer graphics programs that provide various forms of image enhancement, qualitative and quantitative data extraction, and analysis functions.

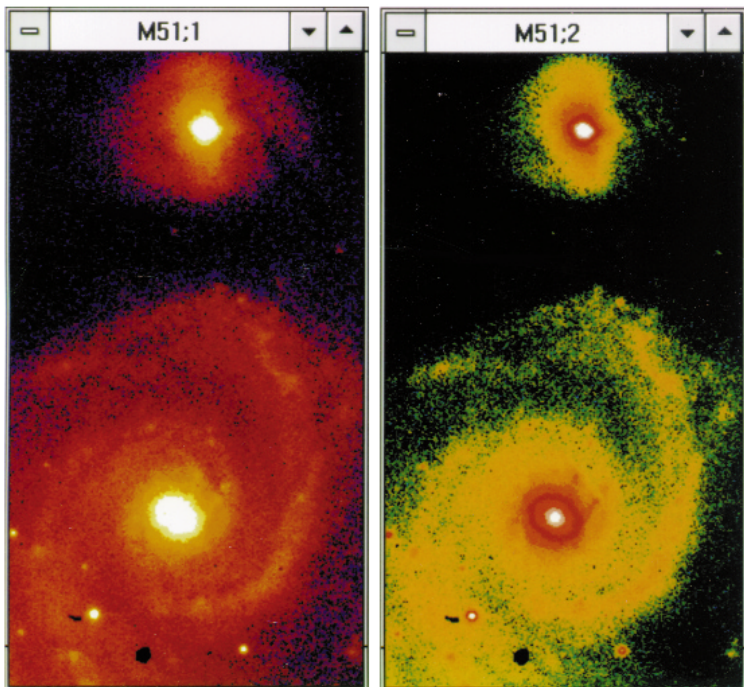
The project goes beyond the traditional static classroom approach to teaching by providing a networked environment in which students interact with other students around the world. They can have their questions answered by professional astronomers and scientists who are participating in the project. They can also participate in collaborative research with other institutions.

A key feature of this project is the use of networked, automated, research-quality telescopes. Via the Internet, participating students can make personal requests for a telescope to take an image of a particular feature in the sky. The students' request for images is the first important step in their subsequent application of the computer image processing tools and astronomy fundamentals. This step is crucial, as it provides a sense of ownership of the image, which in turn increases the students' interest and motivation in performing scientific investigations on the image. Since each image is unique, and since the universe is changing over time, there is a possibility that the image contains astronomical features that have never been viewed before, even by well-funded research institutions. This potential for discovery also enhances the students' interest in learning and participating in the project.

In addition to the networked research-grade telescopes, the underlying technology consists of remote and local file servers that provide archived data and number-crunching capability in response to interactive requests from the participant, and Java-based graphics terminals that provide an interactive visual interface between the user and the rest of the system via the World Wide Web.

The main focus of this Electronic Schoolhouse workshop is to examine the project's Web-based image processing tools, which visually manipulate and analyze astronomical images, and to explore the curricula that accompany the software. Specific subjects include: CCD camera fundamentals, image enhancement, visualization and extraction of measurements from the image (for example, measuring feature sizes on the Moon), exploring photometry, correlating orbits of objects with their visual images and the underlying orbital mechanics principles, and studying changes in objects in the sky.

In the most compelling demonstration of the project, the workshop demonstrates how students in rural America were able to use the software to detect a minute change in images of the sky over time, leading to their discovery of a supernova. They are currently collaborating in a search for comets.



This paper examines the issues involved in the use of high-end interactive media, computer graphics applications, and virtual reality technology in museums. As museums adapt advanced digital media for use in exhibitions and public programs, new relationships develop between the audience, the venue, the virtual representation, and the real object or fact. While use of state-of-the-art technology can effectively shape how museums deliver public education, issues of high cost and maintenance of such technology, larger and diverse audience throughput, and difficulty in content development present important drawbacks. Both the benefits as well as the problems caused by deployment of technology in the museum will be analyzed. Examples will be presented of special museums worldwide that use technology in innovative ways for educational and artistic purposes. Particular focus will be given to presentation of the projects created by the Foundation of the Hellenic World, a cultural heritage institution in Greece, which uses immersive virtual reality, VRML, and 3D graphics for reconstruction of archeological sites, historical interpretation, and education.

Technology in the Museum

Recent technological developments have provided designers with new tools for introducing media into the museum setting. Many museums are now beginning to utilize more recent information technology for internal and external organizational purposes, while more and more interactive exhibits are incorporated into galleries in order to enhance the visitor experience.

Museums realize that they are one among many components in a panoply of cultural amenities and that computer technology can help them quantitatively and qualitatively expand, deepen, and enhance the museum experience for their visitors. A growing number of museum educators regard new media as tools that can offer unparalleled opportunities for learning. Through educational Web sites and CD-ROMs, museums have enhanced their role as providers of informal education. Educators respond to such efforts favorably, as they provide alternatives to restricted curriculum material and allow for more exploration and ownership of the learning process. Finally, museum visitors, especially non-frequent and novice audiences, appreciate and benefit greatly from additional forms of information that make the museum a more accessible and attractive place for them to spend time in.

What are the most common ways that media technology is used nowadays in the museum? One can broadly define the following ways that museums incorporate media in their daily function:

- Audiovisual media used for passive presentation in an appealing way. This, for the most part, consists of video presentations on simple monitors or wall projections in special rooms intended to cover needs in public spaces where staffing is minimal.
- Guided presentation with the help of audio guides, video projections, and other means to accompany visitors throughout their tours, offered as alternatives to the popular tours usually given by museum docents.
- Interactive browsing stations with information on museum collections and educational programs (usually kiosks with "press-a-button," easy-to-learn interfaces).
- Environments that provide opportunities for direct creation or production, take-away experiences, interactive experiences, and innovation.

Current museum theory and practice suggests that technology, as incorporated in today's exhibitions, should generally evolve through two successful formats: inquiry-based guided tours and interactive hands-on exhibits. In this paper we focus on the latter types of technology use, particularly on interactive media that move beyond the point-and-click of common multimedia and information stations. These media may include interactive installations, simulation environments, interactive film, large format theaters or small-scale exhibits, and virtual reality. The explosive development of information technology and the increasing confidence on the part of the museums to



incorporate it in their setting has encouraged many institutions to adopt these sophisticated technological means, innovative environments, and equipment.

Of particular interest to museums is the use of virtual reality (VR) displays and computer-generated interactive experiences that aim at allowing visitors to travel through space and time without stepping out of the museum building. The potential to transcend the physical location of the built environment and the growing sense of the educative function of the museum juxtaposed with commercial pressure has led museums to consider virtual reality as a necessary component in the arsenal of tools to educate, entertain, and dazzle. Although virtual reality suffers immensely from media hyperbole and thus has not lived up to its promises, development of VR systems has matured enough to find its way out of the research realm and into public settings. Introduction of projection-based VR systems has shifted the format from the one-person VR experience with bulky headsets and equipment to the slimmed down, more comfortable visual displays for multiple simultaneous participants.

Multimedia and Internet Stations in a Museum Café

Issues and Challenges

Introduction of high-end or virtual technologies in museums runs up against a number of issues that must be considered. Specifically, high-end technology in the museum must take into account the physical context of the public space, support the conceptual and aesthetic standards of the exhibition or learning purpose, and be functional and accessible to its intended audience.

Whether it's the novelty of interactive technology and virtual reality exhibits or the compelling nature of the applications themselves, visitors flock to museums to see things that are new and cutting edge, even if the content remains relatively unchanged. This generates added worries for museum practitioners, who must design keeping their educational role in mind, yet provide the added novelty, accommodating an increasing range of educational experience, and enhancing their visitor demographic.

Attention can be focused on the following set of issues that previous experience has led us to recognize as the more critical ones in the process of introducing highly interactive technologies in public spaces:

1. Bring technology into context.

Technology in a museum exhibit can not stand alone. To serve museum needs, it must address the specifics of venue and audience, and be used in the context of the total museum experience. The physical context provided by the museum as a public space, the exhibit in its entirety, and the interaction with the exhibits, the other visitors or the museum staff are all important parts of the museum whole. The best media are integrated architecturally and conceptually into an overall exhibition narrative and created directly for the spatial and thematic context of

the museum itself. Effective deployment of technology relies upon the degree to which it can be thoroughly embedded in its context of use, including the ordinary, practical competencies of those who are to work with it.

Museum visitors vary in age, level of education, interests, and learning styles. To address these differences, good exhibit design recognizes the social activity and nature of museums by providing space for more than one person to get involved and by developing exhibits and media experiences that often serve as prompts for social interaction. Media productions in the museum should not come into conflict with the natural, often non-linear, ways a visitor interacts with the museum space. The multi-modal nature of interactive technologies provides ideal opportunities for development of flexible tools for both group and personal experiences.

2. Technology must be seamless.

A well-designed environment is one that is consistent, predictable, and transparent. New media tools are just elements to enhance the educational, informational, or aesthetic goals of any given exhibit and should be seamlessly incorporated. We do not want the media experience to seem like an add-on, an extra to the exhibition, but an integral part of it. Hence, the technology should be designed to "disappear" during the experience. It should be non-obtrusive. Projection systems and computer equipment, for instance, can be hidden from view and acoustically isolated. Visual cues associated with the projection surface or other equipment can be minimized.

3. Provide immediate feedback yet prolonged engagement.

Donald Norman identifies four principles of good design that can be applied to an information environment to make it self-explanatory and not frustrating: visibility, a good conceptual model, good "mappings," and feedback. The format of a typical museum experience involving technology is inevitably controlled, structured, and brief. Hence, responsiveness and feedback are particularly important in museum media settings where we encounter only brief interactions, even in structured events. Users of multimedia installations expect a response immediately; otherwise they walk away. This may be more difficult with highly interactive environments where interactivity is not easily controlled. Experiential-based installations and participatory environments must demonstrate different strategies to include and involve the observer in the events on a long-term basis. Advanced forms of interactive technology, such as virtual reality, allow us to tie together the time and space-breaking nature of interactive media with the time and space-bound nature of the site-specific museum environment. The key is to provide immediate responsiveness and then prompt for deeper involvement on the part of the user.

Overall, experience using virtual reality with large groups of children has shown that the effect can be much more powerful if it possesses two important attributes: an interacting mentor and long-term engagement with the learner.

Interactive

4. Design with content in mind and involve the experts.

In many cases, high-end technological innovation often creates the suspicion of high cost and low content. Current applications seem unable to step out of paradigms created by other media. "Old media" do not always translate gracefully into new media environments, producing, in many cases, outcomes that seem fragmented. It is also often the case in advanced media environments that novelty overshadows content. Technological developments may often be associated with disappointing gains for users whereas compelling content will most likely engage the visitor no matter what the form of presentation may be.

As a result, most educated audiences are skeptical regarding the added value and appropriateness of advanced technology applications in the public domain. Based on the above, it is important to use the involvement of content experts when designing interactive and virtual environments. As difficult as it may be for educators, artists, designers, archeologists, historians, architects, doctors, scientists, and technicians to speak the same language, it is nevertheless essential for creation of sound and complete environments. Collaboration must definitely take place amongst those who are concerned with how things work and those who are not concerned with the technical details but with what is delivered to the public.

The designers of new media must understand the medium to achieve the perfect blend of form and function. They must attract visitors, and they must meet expectations both in terms of their innovation and their delivered content. Questions should be asked regarding the underlying principles that are guiding the development of content, and conventions should be set for structuring and delivering this content. The museum must provide the best combination of technological innovation and educational content and create a shared critical context within which to understand the work aside from understanding it technically.

5. Be concerned about physical and accessibility issues.

Clearly, an important point of particular relevance to high-end technology is usability. Public viewing must be considered in the context of hundreds of people who will visit a site each day, more so if the site is set up for visitor interaction. Practical issues and problems are especially apparent when the apparatus is not designed with novice or special users in mind, as is the case with most experimental high-end computer technology. In the case of virtual reality, for example, it is common for most systems to cause motion sickness; active stereoglasses are too large for small heads, too fragile, and too expensive for use by excited visitors, let alone a child. Children must use both hands to operate hard-to-use interaction devices, hold the stereoglasses with special ties, or even deploy support systems to stand up higher in order to achieve the correct viewing angle.

Issues of high cost and maintenance of advanced technology, larger and diverse audience throughput, and difficulty in content development present important drawbacks for public venues. Prohibitive costs of such technologies and concomitant

inaccessibility, staff development, operation, and maintenance can find no place in dwindling museum budgets overwhelmingly dominated by human resource costs.

Good sight lines, ample seating where applicable, comfortable viewing for extended periods, good field of view, and ergonomics are some of the issues that must be addressed when designing a unique, high-end environment. The interactive experience must also have an easy-to-learn and simple-to-use interface that is accessible to a wide range of skill levels and requires virtually no visitor training.

In brief, accessible new media must be characterized by:

- Attractive designs
- Rugged engineering
- Accessibility
- Practical maintenance

Despite the above issues, there are ambitious and significant efforts taking place in museums worldwide that have used high-end technology to complement their exhibitions. Some examples of innovative media use in museums are mentioned below.

Innovative Examples of Media Museums and Museums Using Media

Science museums, hands-on children's museums, large "edutainment" venues, and recreation parks have traditionally embraced new media first. Museums such as the Exploratorium, the Tech Museum of Innovation, or the Computer Museum in the United States employ fascinating and sophisticated interactive installations and have been presenting up-to-date results on the creative use of technology.

Aside from existing museums that have employed higher-end technologies, a new form of museum has emerged, as exemplified by a growing number of media museums worldwide such as the Ars Electronica Center in Austria, the ICC center in Japan, and the Center for Art and Media (ZKM) in Germany. These are new types of museums devoted to depiction and presentation of new media, or "productive media art museums," a term coined in ZKM's mission statement. They are true pioneers in the use of high-end interactive technologies open to the public. The Ars Electronica Center, for example, was the first museum worldwide to install and open to the public a CAVE, an immersive multi-person virtual reality room, where high-resolution, 3D video and audio are projected on three walls and the floor.

In the more traditional fields of fine arts, archeology, and history, there are hardly any museums that have moved beyond multimedia presentations (which are usually placed outside their main exhibition galleries). Unique is the case of the Foundation of the Hellenic World (FHW), a non-profit cultural institution in Athens, Greece. The foundation's mission is to preserve and disseminate Hellenic culture, historical memory,

Acknowledgements

The Foundation of the Hellenic World is a privately funded, non-profit cultural institution based in Athens, Greece. The authors wish to thank the staff of the museum, who have worked hard to achieve its innovative interactive programming and vision.

Notes/References

1. Ars Electronica Center, Linz, Austria, www.aec.at/
2. Brooks, F. "What's Real About Virtual Reality?" Keynote speech at IEEE VR '99, Houston, Texas, March 13-17, 1999.
3. Dietz, S. *Curating (on) the Web*, Proceedings of Museums and the Web, Toronto, April, 1998.
4. Foundation of the Hellenic World, www.fhw.gr/
5. ICC Intercommunication Center, *ICC Concept Book*, NTT Publishing Co., Ltd., Tokyo, 1997, www.ntticc.or.jp/
6. Norman, D. A. *The Design of Everyday Things*. Doubleday, New York, 1988.
7. Thomas, S. and Mintz, A. *The Virtual and the Real: Media in the Museum*, American Association of Museums, 1998.
8. ZKM - Zentrum für Kunst und Medientechnologie, Karlsruhe, www.zkm.de/

and tradition, as well as project the contribution of the Hellenic spirit in the development of civilization. Its main plan for operation is the use of advanced technological methods to promote this understanding of the past, as a point of reference for shaping the present and the future. The foundation is housed in a large innovative museum/cultural center on an old 18-acre industrial plot in Athens, where historical exhibits on all periods of Hellenic culture are presented using the latest audiovisual technology.

Consistent with the foundation's decision and commitment to use state-of-the-art technology, normally found only in research labs, is the plan to develop a type of "VR center" focused on the cultural domain. Not only is this blend of the traditional with the most advanced unique as far as museums are concerned, it is also a one-of-a-kind endeavor for Greece.

Specifically, the Foundation of the Hellenic World's high-end technology endeavors involve immersive virtual reality. Generally, VR technology in FHW functions in two basic ways: as an educational/entertainment tool and as an instrument of historic research, simulation, and reconstruction. Examples of applications include an educational environment related to traditional Greek costume throughout time and reconstruction of and a journey through the ancient city of Miletus by the coast of Asia Minor, as it was in antiquity. Guided by a virtual "time-machine," our travelers are able to explore the city as it unfolds through time, and experience the life of its architectural glory, its people, and their customs, habits, and way of life. To present such projects, the Foundation has purchased a roving immersive projection table, the ImmersaDesk. This virtual reality environment "on wheels" has the ability to move about the museum campus. It thus enables creation of a truly flexible virtual reality setting for the museum and allows for customization of the virtual reality exhibit to any current exhibition or program.

While this type of setup can afford unique educational and recreational opportunities, equipment of this kind is inevitably in an experimental stage and suffers from a number of drawbacks in terms of usability – essentially, many of the practical problems mentioned above. The roving nature of this particular virtual reality system is a relative advantage as the equipment is fragile and requires special handling; once installed, the virtual reality equipment requires a specially designed place where metal structures are not present. The display system must be designed to withstand breakage, short attention spans, greasy fingers, and large numbers of visitors. The special glasses are expensive and can break easily as their sizes are not made to fit everyone's head. Equipment must be placed out of reach yet remain accessible to the technical staff. Finally, experienced guides who have both technical skills and museum education background are always required.

The above efforts set standards in the direction of technology use in informal public settings, yet the contours of the new media landscape are just now becoming visible. No one yet knows what will be successful or how visitors will ultimately use and interact with these emerging technologies. We must, therefore, continue to see that the insights gained through experienced use are adequately translated into the public space and are both inquisitively and critically examined.



This Playground exhibit demonstrates how a unique, cross-age, multicultural, and, most importantly, cooperative learning environment allows for producing well-prepared animation artists as well as production of animation itself that is tremendously engaging (and fun!).

For the computer graphics student, having a strong drawing and traditional animation background provides the strongest possible foundation for developing a career in computer animation. Conversely, computer animation experience is critical to the learning needs of traditional animation students. Therefore, the Abram Friedman Occupational Center gears its program toward the holistic integration of both disciplines to provide learning for students who are concentrating solely on a career in one or the other.

This exhibit shows how traditional animation and computer animation classes collaborated to produce an animated short. The exhibit emphasizes how computer animation techniques for producing such things as lighting, environments, and particle systems were married with traditional 2D techniques such as hand-drawn character animation to produce a seamless animation that highlights the best qualities of both.

All steps of the collaboration are demonstrated, from developing the story in discussion to how it was storyboarded on paper, from building and animating characters using computer animation to more traditional techniques that utilize paper, cels, disks, pencil test machines, and animation stands. The use of pencil test machines, scanners, and computers to create animation tests for characters, backgrounds, and environments will also be demonstrated. How the students rendered scenes, edited, and assembled footage to develop a final product will also be shown. Team building will be demonstrated and discussed, as representatives from each class "walk and talk" observers through the exhibit. The exhibit also includes a documentary showing how the animation was made.

This exhibit is an opportunity for students to talk about both the learning experience and how it prepares them for work in the industry, whether they are computer graphics artists or traditional animators.

Animated

Incorporating Principles and Examples from Art/Design and Film/Video into a CS Computer Graphics Course

Introduction

Computer graphics is one of the most technology-dependent areas in computer science. Creating and storing realistic 3D images and animations requires substantial processing speed, main memory, disk space, and graphics display capability. Fortunately, we are in a field where incredible advances in these technologies are an accepted fact. Thus, I have changed my course constantly over the past few years.

A non-technology change that I have made is to incorporate established principles from art/design and film/video and into my computer graphics courses. Artists have been producing images for purposes of visual communication for over 10,000 years. They have established a rich set of principles and guidelines for image composition and have worked out many of the problems that have been faced in computer graphics. Similarly, film has a history of over 100 years and has added extensions and modifications to the art principles. Our computer science students can benefit from exposure to these principles. In addition to teaching them algorithms, we should teach them some of the aspects of visual communication. I also think that they would benefit from exposure to examples from art that illustrate some of our core principles.

In this paper, I will briefly describe the particular environment in which my course is taught and the changes within my course over the past several years. I will discuss how these principles and examples from art/design and film/video can be incorporated into a computer graphics course. Most of the examples given below are in HyperGraph, which is available in the Curriculum and Materials portion of the ACM SIGGRAPH Education Committee Web site (www.education.siggraph.org).*

Background

Georgia State University is primarily a commuter university with approximately 24,000 students, many of whom live 30 to 45 minutes away from campus. The undergraduates have a sound mathematical background and are familiar with vectors and matrices. The prerequisites to the first graphics course include the first three computer science courses, covering programming principles and data structures in Pascal and Ada.

Many of our students work part time or full time. Almost all of the students have their own machines; almost universally these are IBM PC compatible machines. This leads to several constraints. First, I cannot force the students to use Java, C, or C++. Second, it would not be fair to the students to require them to spend a large amount of time on campus. They should be able to do most of the work at home.

Although the course material and teaching methods change every time I teach the computer graphics courses, my objectives have not significantly changed. These objectives include both computer graphics and non-computer graphics issues. The graphics objectives are that the students should learn about graphics systems, algorithms, and the process of creating graphical images. I have emphasized 3D graphics and the 3D graphics pipeline for several years, although I still cover 2D graphics algorithms. Although I cover different modeling methods, an emphasis has been on different image synthesis techniques.¹ Perhaps this is because of my background in physical science, whereas someone who has a background in mathematics might tend to emphasize modeling techniques.

I have changed the projects for the students over a period of years, always adhering to the constraints mentioned above. Initially, I had students develop a scanline graphics system that culminated in display and manipulation of a simple 3D faceted shaded object such as a parallelepiped. Then GSU obtained a license for the RenderMan development system, and I began using this in my courses, in addition to the programming projects.²

As computers became faster (with the introduction of the Intel 386 chip) it became feasible to teach ray tracing. Students developed a ray tracing program over the course of the quarter. I still used RenderMan, since I wanted the students to be able to compare scanline graphics images with ray traced images.

With the advent of VRML, I switched from RenderMan to VRML.³ Unfortunately, I have found that I need to periodically change my assignments as the students tend to copy the previous year's assignments. Over the past year, I have used Larry Gritz's Blue Moon Rendering Tools (BMRT) program. This freeware package adheres to the RenderMan interface specification and creates both ray-traced and/or radiosity-based images. So, instead of writing a ray tracer from scratch, my students create Renderman RIB files and then use BMRT to render the images. This allows them to compare a system designed for high quality image output using global illumination methods (BMRT) with a scanline system designed for interactive speed (for example, VRML).

GSU has recently revised its computer science curriculum so that the first programming courses are in Java and then C++. This means that in the near future I will be able to explore the use of other tools for programming assignments, such as OpenGL and Java3D.

My courses have also been changed in that I have incorporated principles from other, related fields. This is the primary topic of this paper and in the next few sections I will give examples of how I have done this.

Image Composition and Graphics Design

Our students are going to be creating graphical user interfaces (GUIs), Web pages, and computer images. Therefore, it is important that they understand some of the basic principles of design and image composition. A human computer interaction (HCI) course might discuss these topics, but HCI courses usually emphasize cognitive aspects rather than the visual display aspects of interfaces.

Some of the image composition topics that I cover include the following:

- General concepts such as: line and contour, value, shape, texture, space (both positive and negative), color
- Unity and harmony
- Ways to achieve emphasis: contrast, tangents, isolation, shape, motion
- Ways to achieve balance including the concept of visual weight: value, color, shape, texture, position, eye direction
- Scale and proportion
- Repetition and rhythm

These can be somewhat specific to Web design⁴ or they can be more general.^{2,5,6,7} Scene composition is implemented in computer graphics by the choice of a virtual camera placement, orientation, and field of view. However, there are other aspects of camera usage that we can learn from film by looking at the purpose of different types of camera shots. For example, what is the purpose of a long shot versus a close-up? What are we trying to accomplish with these different types of camera shots?

P. Turner presented some similar ideas in the context of an art computer animation course at SIGGRAPH 98.⁸ For example, depth of field can be used to shift the emphasis from one character or area to another. In this first scene, the creature and Debbie are having an innocent conversation, with the center of focus and attention on them. Next, we switch to focusing on the evil alien as he covertly observes their conversation. (Both images are on the Conference Abstracts & Applications CD-ROM)

Lighting Considerations

In discussing the lighting of computer images, I use the example of lights from film and video,⁹ and how these principles can be applied to computer images.^{5,10} The different types of lights used in film (key light, fill light, etc.) are discussed. Then the kinds of lights available in CG systems, (point lights, distant lights, spotlights, etc.) are defined, and we discuss how we can use the CG lights to emulate the real lights used in film. The importance of lighting for scene composition and emphasis, is covered. We also discuss how different types of lighting can create different types of moods or emotional impact.

Using Examples from Art to Illustrate Computer Graphics Concepts

There are several places in my computer graphics courses where I incorporate examples from art. I do this to illustrate to the students that these are not really new concepts, but are problems that were long ago solved by artists. Here are some examples of this.

Perspective

The architect Filippo Brunelleschi discovered linear or scientific perspective during the Italian Renaissance.^{2,11} This allowed artists to use geometric methods to project a 3D space into a plane. Masaccio was the first painter to use this technique, and I use his painting, along with other examples from art, as shown in illustration 1.

Two other examples I use (with the images on the CAA CD-ROM) are as follows: a painting (The Piazza of St. Mark, Venice) done by Canaletto in 1735-45 in one-point perspective; a painting (Sunlight in a Cafeteria) in two-point perspective by Edward Hopper (1958).

Another example illustrates how well-known classic CG images were inspired by art. The environment depicted on illustration 2 was inspired by the painting *Lady and Gentleman at the Virginals*, by the 17th century Dutch painter, Jan Vermeer. Illustration 3 is the original, painted in 1662-65. Vermeer is well known for the depiction of light in his paintings. A modified radiosity solution was ray traced to produce the specular highlights on the floor. The image is from the 1987 paper "A Two Pass Solution to the Rendering Equation: a Synthesis of Ray Tracing and Radiosity Methods" by John R. Wallace, Michael F. Cohen, and Donald P. Greenberg. The image appeared on the cover of 'Computer Graphics: Principles and Practice' by Foley, van Dam, Feiner, and Hughes.

Contrapposto

Another example is in the modeling and posing of human figures. Here we can use examples from sculpture. Early sculptures of human figures, while anatomically correct, appeared stiff and unnatural. The classical Greeks progressed to where they were able to model the human form in a nonsymmetrical, relaxed stance that appeared much more realistic. This is described by the Italian word "contrapposto" (counterpoise). This technique was lost during the Dark Ages and rediscovered by Donatello during the Italian Renaissance. All images can be viewed on the CAA CD-ROM; they are from Mark Harden's art archive.¹²

There are three examples that I use (with the images on the CAA CD-ROM): here is an example of Egyptian sculpture from about 1920-1880 B.C.¹² Notice the unnatural stiffness of the figure. Here is an example of Greek sculpture from about 440 B.C. Notice the nonsymmetrical, relaxed stance, which appears much more natural. Here is Donatello's *David* from 1444-46. As with the Greek statue above, it is relaxed, nonsymmetrical, and realistic.



Illustration 1
Here is the first ever painting (*Trinity with the Virgin, St. John and Donors*) done in perspective by Masaccio, in 1427.

Conclusion

Going beyond a computer graphics course that is purely algorithms by teaching computer science students aspects of visual communication, imported from other but related fields, is something we should seriously consider. Our students benefit from an exposure to elements of art history because they see that artists, painters, sculptors, and filmmakers have encountered and solved many of the same problems that we find in computer graphics.



Illustration 2



Illustration 3

References

1. Owen, G.S. Teaching Image Synthesis as a Physical Science, *Computers & Graphics*, Vol. 18, No. 3, pp. 305-308, May/June, 1994.
2. Ocvirk, O.G., Stinson, R.E., Wigg, P.R., Bone, R.O., and Cayton, D.L. *Art Fundamentals: Theory and Practice*, Brown and Benchmark, 7th edition, 1994.
3. Owen, G.S. Using VRML in an Introductory Computer Graphics Course, *IEEE Computer Graphics and Applications*, (First WWW-based issue announced in Vol. 16:3, pp. 16-17, May 1996 and activated in September, 1996, URL: <http://computer.org/pubs/cg&a/cged/>).
4. Mitchell, B. and Weinman, L., *Creative Design For The World Wide Web*, SIGGRAPH 97 Course 2.
5. Kahrs, J., Calahan, S., Carson, D., and Poster, S. *Pixel Cinematography: A Lighting Approach for Computer Graphics*, SIGGRAPH 96 Course 30.
6. Glassner, A., Callender, J., Gleason, M., Kerwin, B., Mahoney, J. *Art for Computer Graphicists*, SIGGRAPH 98 Course 30.
7. Lauer, D.A. *Design Basics*, Holt, Rinehart and Winston, Inc., 3rd edition, 1990.
8. Turner, P. *The Language of Cinema and Traditional Animation in the 3D Computer Animation Classroom*. pp. 80-83, *Conference Abstracts and Applications*, SIGGRAPH 98 Educators Program.
9. Gross, L.S. and Ward, L.W. *Electronic Moviemaking*, Wadsworth, 2nd edition, 1994.
10. de Leeuw, B. *Digital Cinematography*, AP Professional, 1997, *Inside 3D Studio Max Volume III: Animation*, Maestri, et al, New Riders, 1997.
11. Janson, H. W. and Janson, A. F. *History of Art*, Prentice Hall, 5th Edition, 1995.
12. Mark Harden, www.artchive.com

*A more extensive version of this paper, with more images, is contained on the Conference Abstracts & Applications CD-ROM.

Classroom

Panelists

Taylor Gutermute
California Department of Education

Lynn Hickey
Los Angeles Unified School District

John Hughes
Rhythm and Hues Studios

Alan Warhaftig
Fairfax Magnet High School for the Arts

This panel discusses issues related to how and why arts instruction in California public schools can be supported with contemporary technologies to enhance students, learning and achievement, and their preparation for life after graduation. The discussion includes the unique issues and concerns of a classroom teacher and the Visual Arts Specialist in the Los Angeles Unified School district, the Arts Education Consultant for the California State Department of Education, and the President and CEO of Rhythm & Hues, Inc. They review the findings of the State Superintendent's Task Force on the Visual and Performing Arts published in "Artswork: A Call for Arts Education for all California Students." The Task Force stated the following overarching goal for arts education:

All students in California public schools will have high-quality arts education programs from pre-kindergarten through grade 12. All students will:

- Develop and demonstrate literacy in and through dance, music, theatre, and the visual arts.
- Participate in arts-related school-to-career experiences.
- Have access to the arts through a variety of educational experiences and technologies both in and out of school.

Opportunities and constraints experienced by teachers in the Los Angeles Unified School District provide case studies for the panel discussion.

Lynn Hickey

Visual Arts Specialist, Los Angeles Unified School District

With graduate degrees in the visual arts, education, and school administration, Lynn Hickey has been involved as a visual arts teacher, school administrator, provider of professional development programs, program developer, and advocate for the use of educational technologies in the Los Angeles schools, the California Department of Education, and several campuses of the California State University.

While developing aesthetic sensibilities, an arts education in the new millennium must further students' abilities to participate in an information age and within a global economy. Through a sequential visual and performing arts curriculum supported by contemporary technologies, students must learn to:

- Communicate effectively using a variety of symbol systems.
- Understand people and traditions in a diverse and complicated world.

Taylor Gutermute

Consultant, Visual and Performing Arts, California Department of Education

Taylor Gutermute does developmental work in the visual and performing arts, most recently coordinating the development of the Visual and Performing Arts Challenge Standards. She serves as a liaison to the arts education field both statewide and nationally.

Arts education in California public schools has traditionally focused on the disciplines of dance, music, theater, and the visual arts. For each of these disciplines, arts educators are asked to include four components in their curriculum: artistic perception, creative expression, historical and cultural context, and aesthetic judgment. At this moment in time, with the onset of new technologies, arts educators are asking and are asked questions such as: Who teaches computer graphics? Is video production visual arts or theatre? How do teachers who have taught traditional visual arts techniques get training in new technologies? Is the funding available to purchase necessary equipment for art teachers or those in traditional career education programs?

While these questions appear mundane, finding answers will help to get effective programs implemented. The recent California Arts Task Force report, *Artswork*, emphatically requests that there be an "updating of the arts" to embrace the field of computer graphics and interactive techniques.

John Hughes

President and CEO, Rhythm & Hues, Inc.

Rhythm & Hues produces animated visual effects for feature films, theme park attractions, music videos, commercials, and interactive video games. In 1995, Rhythm & Hues won the Academy Award for Best Visual Effects for "Babe."

The K-12 educational environment is challenged to prepare a literate populace, one that can appreciate the visual and performing arts as well as participate in a variety of businesses and industries dependent upon communicating with digital media. It has been demonstrated that students in schools with a strong emphasis on the arts experience greater meaning, excitement, and depth in what they learn. They are more motivated, engaged, and eager to learn. Even if they do not envision themselves in an arts-related career in the future, all students can experience the joy and inspiration of the arts, understand the connection of the arts to their lives, and appreciate excellence in the arts.

Most specifically, we have very particular needs in my industry, for example, for people who are knowledgeable about history and culture and have the ability to draw, design, and communicate ideas using digital media. Industry leaders have both ideas and the responsibility to share them with the education community. Together, for the benefit of all, we must re-conceptualize the components of a quality arts education and ensure its delivery to all students.

Alan Warhaftig

Fairfax Magnet High School for the Arts

Alan Warhaftig teaches American literature and first-year algebra at the Fairfax Magnet Center for Visual Arts. A graduate of Stanford University, where he majored in Social Thought and Institutions and specialized in Caribbean Studies, Mr. Warhaftig has been deeply involved in technology, arts, curriculum, and professional development issues in the Los Angeles Unified School District. As co-chair of the Technology Focus Group, he wrote two widely-disseminated Discussion Papers reflecting his concern that the value of computers and the Internet for K-12 education has been wildly overstated and expressing his doubt about the institutional ability of schools to cope with technology's costs and other complications.

K-12 visual arts education in the new millennium will still require foundation training along the lines of that offered at the Fairfax Magnet Center for Visual Arts, where ninth and 10th graders take introductory classes in drawing, sculpture, photography, and computer graphics. In 11th and 12th grade, they take more advanced courses in these areas. Computers are part of our training, but it's inconceivable that they could substitute for training in the other areas without damaging the development of our student artists.

As a teacher of American literature, my job is to teach students to read critically. We closely study challenging works. Writing follows reading, and in my classroom writing is about clear expression of ideas. Technical issues are important but secondary. My most precious resource is instructional time, and I would regard the introduction of computers in my classroom as a distraction.

There are certainly valuable uses for computers and the Internet in K-12 education, and a sensible approach would be to identify these excellent uses, on a grade-by-grade, course-by-course basis, and teach teachers how to use technology to improve curriculum delivery. This should happen before schools spend huge sums on technology that they are not sure how they'll use.

The Integration of Graphics, Video, Science, and Communication Technologies

Panelists

Thomas D. Cauffield
Glenn Dame
Kevin J. Meehan
Robert Wickman
Forest Hill Community High School

At Forest Hill Community High School in West Palm Beach, Florida, a unique program that integrates technology, science, video, and the graphic arts has been developed to provide students in grades nine through 12 with employment skills for the information economy. This modern multimedia school-to-career academy blends the skills of students and teachers with diverse multimedia interests to provide real-world experiences that support local government and international environmental issues.

As part of a Sister City project between the City of West Palm Beach and the Tzahar Region, Israel, the students at Forest Hill Community High School are developing an international Internet site that will allow science students, teachers, government officials, university personnel, and environmental scientists and engineers to share data, graphics, video, and artwork about the South Florida Everglades and the Hula Valley in Israel. This international project uses student and teacher talent in the following engineering and graphic application skill areas: Microsoft NT, 3D Studio Max, HTML, XHTML, Painter, PhotoShop, PageMaker, Frontpage, Macromedia Authorware, Dreamweaver, linear and nonlinear video editing, and environmental science.




FHCHS graphics and engineering students work together with 3D Studio Max.



West Palm Beach community members study graphics and video editing in an evening program at Forest Hill.



Students prepare video title pages and graphics for their Web site (www.fhchs.com).



Introduction

Storytelling is an effective educational medium. Imagine the possibility of having a conversation with someone who has been dead for almost 100 years. Imagine the tales and everyday life stories this person could tell. Imagine seeing that person's world and being free to explore it. Now, give this experience to a classroom of students and watch them become active learners without them even knowing it. They will be enthralled with the possibilities set before them. The Interactive Learning Environment (ILE) is designed to do just this.

For our first application, ILE focuses on the fourth grade-classroom. Students study California through a local landmark, the Bidwell Mansion in Chico, California. ILE allows the learners, both teacher and students, to explore the mansion through computer-generated, 3D environments. Integrated lesson plans support the curriculum required of a fourth grade classroom.

Presenters
Stephen Detwiler
Elizabeth Padilla
California State University, Chico

Jonathan Hendryx
Arizona State University

The Theory Behind the Educational Method

ILE merges group work and traditional individual learning methods into a process called communal work. In this approach, group development is based on individual knowledge and participation. Communal work can be surprisingly helpful, enlightening, and beneficial to all students. The key to success is finding a way that not only works in theory but in the classroom as well. A good group allows for all levels of learning and a shared goal, yet the learning and the objectives should also be individual. The work of a group should be stronger, and the goal should be harder to reach than one person alone could achieve. It becomes pointless to have a set of minds working for a goal that one alone could do. It is essential all the members work toward the goal and share the information they have accumulated. The connection of group members must ensure that competition for the best answer and superiority over other members is eliminated.

With ILE, learners reach out to others and interact, not only with their groups, but with computer-generated people as well. The people the learner comes into contact with allow information to pass from one person to another, making the learning experience a two-way exchange. The learner can then become not only a receiver but also a transmitter.

ILE is not a replacement for the classroom. Children will not be placed before a screen and asked questions and lectured to by an electronic teacher. ILE is a tool for implementing all the functions for learning. Students are able to walk away from the machine and apply the knowledge that they have gained to the other activities that correspond to a themed learning unit, whether it is history or chemistry.

The classroom environment is not regimented by ILE. It should still reflect the teacher's own style. ILE is just another component, an extension of learning through hands-on experiences such as field trips and museums. These are ways to expose children to art, history, and other elements of cultures. ILE immerses learners into the environment, giving them a deeper understanding and comprehension.

The Technical Design

From the beginning, ILE has been designed to work on various types of computers. Recognizing that not every school can have the latest and greatest computers, ILE automatically adapts itself to the machine that is currently running it. The program was built using industry-standard technology components. The 3D content is streamed from an ILE geometry server to a Java applet on the client and then rendered to the screen via Java3D. By tracking each client computer's performance, the server only sends as much information as the client can process. The geometry server streams the geometry in levels of detail, and the client only processes as many polygons as it can handle at a given time. By this method, the client workstation determines how detailed the environment will be.

ILE is a multi-user system. A primary goal of ILE is to ensure that students are not isolated from each other. Allowing multiple students to explore the shared environment has several advantages. Students can share their experiences in real-time. Also, students can move in groups and interact with the environment together. Students can talk with each other through a text-based chat system, and the teacher can talk with them to give advice or ask questions. When students enter an environment, they can be assigned to a group by the teacher and given a task to complete. Their progress is tracked by the server and can be monitored by the teacher at anytime.

Each student's progress is stored on the server in a database. By accessing each student's performance, the system can also recommend different areas that they should explore and study. Messages from the server are delivered in a variety of modes. Some messages come in textual form, but the goal is to have all messages come from computer-controlled characters that approach students to ask questions or give advice.

The Windows NT-based server is an open, object-oriented system. Topic modules can be loaded onto the server, allowing any number of subjects to be presented. New technologies can be added in the future without having to modify the older code. The multimedia content is obtained through Triadigm Technology's FutureArts server. The text, images, audio, and video are stored in a database and extracted as needed. This allows ILE to present information to students in various forms. By tracking student performance based on presentation method, the server tries to give information to students in the way they learn best. For example, a student who seems to learn better by reading will be given textual descriptions first. Audio narration is provided for students who appear to learn best by listening. The ILE server uses TCP/IP network protocols, so it can reside locally or be accessed remotely via the Internet. By providing an open and versatile system, ILE can be adapted to many different learning environments.

Conclusion

Integration of computers into classroom can increase the learning potential of a student. Computers can provide immediate feedback on student progress and allow a teacher to better monitor progress. The Interactive Learning Environment is a step in the right direction, toward more than Internet workstations. Computers become tools for cognitive development and support participants in a rich learning environment.

The Process

Though it may appear to be filled with complexity, there are no great secrets in the fascinating world of 3D animation. To put it into familiar terms, to produce successful 3D animation, you must visualize the computer and your 3D software as an empty soundstage and complete seven basic steps:

Script/Storyboards

The more time and energy you invest in planning and preproduction, the more successful the final product will be. A good script and finely detailed storyboards can work wonders in communicating the director's vision to an animation team.

Modeling

Modeling is the process of "building" the geometry in a 3D package. Models are defined in terms that allow the software to simulate the sets and characters that populate the scene.

Materials/Textures/Lights

Completed models and characters appear to be constructed of gray (or blue, etc.) cardboard. Applying materials and textures to models is the process of defining the surface characteristics of those models. Are they red? Are they shiny? Are they made of plastic or metal? Are they smooth or rough? Do the walls have wallpaper? Materials and textures are vital in making the 3D environment "real." So is lighting. Just as in theater or live-action filmmaking, lighting can make or break a mood. Properly defining lights, placing them, and creating realistic shadows is just as important in the 3D world.

Animation

By definition, animation is the process of presenting a series of still images in succession so that the viewer perceives continuous motion. In 3D, animation is the process of adding motion or change to the scene and saving it. With current software, most elements in a scene can be animated: objects, characters, cameras, lights, materials, and effects.

Rendering

Because animation is the sequential display of still images (30 images per second for video, 24 per second for film), the 3D environment built on our virtual sound stage must be converted to 2D images. This process is called rendering and can be extremely costly in terms of time and computer resources. The more complex a scene, the longer it takes to render a frame. Because of this, all previous steps in the 3D process are executed with rendering in mind.

Compositing

3D work is frequently created for use in a 2D scene. So compositing, or combining two or more 2D images to create a new image, is a part of completing a 3D sequence. For instance, 3D F-14 fighters might be composited over live-action footage of the California desert.

Output

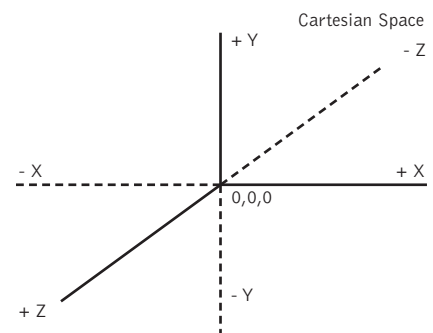
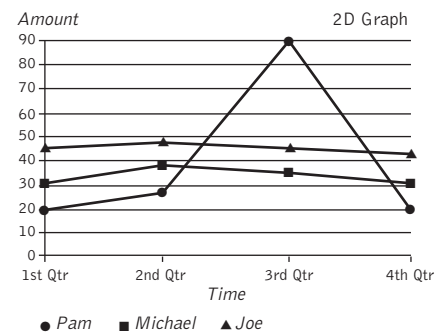
After it is rendered, a sequence must be output to a medium for distribution such as film, video, CD.

3D Space

Global Space

In a 3D program, you are working in 3D space based on a system called the Cartesian Grid. Cartesian space is defined by a series of axes and planes.

When you are working in 2D space (for example, on a 2D image or a graph), the



space is limited to, and defined by, two directions; left to right, and up and down. These directions, or axes, are usually defined as X and Y. In the example on the previous page, the X axis (left to right) displays the time and the Y axis (up and down) displays the amount. Each of the axes are divided into even segments, making it possible to determine an exact amount at a specific time.

This system is the basis of Cartesian Space, with the addition of a third axis to define depth. Although axis definitions may vary from one program to the next, most computer animation systems define left/right as X, up/down as Y and forward/backward as Z.

All three axes continue past their intersection, defining a 3D space that is infinite. The intersection of the three axes is referred to as the global or universal origin or center, and the numeric designations for each axis begin at this point. The center is at zero for each axis: X0, Y0, and Z0 or 0, 0, 0.

With this system it is possible to define exactly where in 3D space a specific point lies (for example, X3, Y4, Z5), making it possible for computers to simulate 3 dimensions.

Local Space

Cartesian Space is the entire space in which the computer creates your simulated environment. The origin, or center, of that space is the global center. As the global space has a center, so does each object in your scene. This is the local center. The local space is the area that is inhabited by the object.

The Building Blocks

3D software packages define and manipulate simulations of shapes using the mathematics of shapes, or geometry.

Polygons

When you can define specific points in space, you can build models. Defining two points in space gives the coordinates to draw a line. If you continue to define points (control points or vertices) and connect them with lines, you can, with a minimum of three points, build a polygon. A polygon with a surface, or face, is the primary building block of polygonal modeling programs. Together, multiple polygons build a 3D model. To help organize the large numbers of polygons in complex models, programs commonly define the model as a mesh, or grid, of polygons.

Splines and Patches

Because polygonal modeling cannot recreate curved surfaces, researchers developed the curved line, or spline. Splines are curved lines with control points, or vertices, that determine the shape of the curve. How the vertices affect the curve is determined by the type of curve you have drawn. With splines, curved surfaces or patches can be generated. The shape of the patch can be changed by editing the control points.

Primitives

Specific 2D and 3D shapes that are so commonly used in 3D can be created automatically. Referred to as primitives, these shapes include such objects as cubes, spheres, cylinders, squares, circles, toruses, cones, and grids. Since so many objects are comprised of these basic geometric shapes, primitives are very useful starting points in modeling.

Transformations

When working with 3D geometry, you frequently edit an object to get the model you want, or you change it to create motion. There are three basic changes, or transformations.

Scale

Scaling changes the size of an object. It is expressed in terms of a ratio of the new size in relation to the original size of the object (for example: the cube is now 2.75 the size of the original cube). As with everything else in 3D animation, the changes are made in relation to the X, Y, and Z axes.

Rotation

Rotation changes the orientation of the object around the X, Y, and/or Z axes. Rotation is expressed in degrees and, like scaling, the changes are made in relation to the axes.

Translation

Translation changes the location of, or moves, an object. Translation is expressed in units (inches, miles, kilometers, or the default unit of the software with which you are working) and, as always, is in relation to the starting location of the object along the axes.

Although there are more sophisticated processes for building models or making motion for animation, all changes to the geometry of an object are based on scale, rotation, and translation. You may be rotating only one surface of an object, translating the points that control the shape, or scaling just part of your model, but all changes are based on the major transitions.

Summary

So now you know the basics. Armed with this foundation, you can begin to understand and work in the 3D computer animation world.

Have fun!

CGI (Computer Generated Images) – 3D computer animation

CGI creates the entire image in the computer. A true 3D model of an object is built in the computer, the object is animated, then the computer “renders” the object (computes the finished image). This technology permits creation of virtually any imaginable image without the practical constraints of the real world – a huge advantage in the world of special effects.

The types of objects that can be created fall into several categories. Character animation is where a character of some kind is created and animated, such as the T. Rex in *Jurassic Park*. Effects animation is used to create computer versions of snow, rain, tornadoes, tidal waves, and other effects. Props such as spaceships and fighter planes can also be created and animated. Even entire virtual sets can be created and live action actors placed within them.

The first step is the modeling process, where the 3D shape of the object is created. Points in 3D space are defined, then they are connected together to describe a surface. When finished, the “wireframe” view of the object shows the object’s shape and can be viewed from any angle.

The next step is the surface attributes: color and texture. Texture maps can be applied to the surfaces that give the look of a painted surface. Bump maps give the surface texture. Environment maps create reflections of the environment on the surface of the object. There are other types of maps that can be applied in combination until a very realistic surface is built up.

Lighting is achieved in ways similar to the real world. Several computer “lights” are placed within the scene to illuminate the object. Of course, the rules for the computer lights are much more generous than the rules for real lights. Computer lights have no real world limits and can do remarkable things like shine through walls or maintain full brightness without dimming over distance if the artist wishes them to. While one artist is working on the surface attributes and lighting, a 3D animator can be working on the animation of the object.

There are basically three types of animation methods. Keyframe animation is where the object is “posed” at a few key frames in a scene, then the computer interpolates all of the in-between positions for each frame of the scene. Motion capture uses live actors wired up with sensors so that a computer can track the position of multiple body parts simultaneously. This captured motion data is then fed into the computer character, which duplicates the moves. Procedural animation (sometimes called “dynamics”) uses rules and math formulas to describe some action, such as falling snowflakes. Each flake is given a falling rate of speed, a wind factor, and a turbulence factor, then the computer figures out the motion of each flake individually.

The final step is rendering, where the finished image is computed. The computer combines all the information from the shape of the model, the various surface attribute maps, the lights in the scene, and the object’s animation, locates the camera’s position, then renders the scene as seen by that camera. This is the most compute-intensive part of the job. Using very powerful computers, render times of several hours per frame are not unusual. Typically, each object in the scene is rendered separately, then combined afterwards in a compositing operation.

Effects

Compositing

Compositing is the process of combining pre-existing images, or elements, into a single picture. The background might be a live-action scene, the character might be a live-action actor shot on a greenscreen stage, and a prop such as a car might be created with CGI. Regardless of how the elements are created, it is the compositing operation that puts them all together and color corrects them to look right.

When compositing one element over another, the foreground (top) element will always require a matte layer that tells the computer which parts of the foreground image are kept and which parts are thrown away. These mattes can be created several ways. With a greenscreen (or bluescreen) shot, the computer creates a matte from the green part of the screen. CGI images, however, are rendered with a matte automatically so that the compositor does not have to create one. When all else fails, a matte can be painted by hand (a "roto" matte), but this is slow, expensive, and difficult to do well.

In addition to combining and color correcting the various elements into a finished picture, compositing can add some special effects tricks of its own. Two-dimensional animation can be used to rotate, scale, or move elements around the frame. Elements can be "cloned" and multiple copies placed all around. Images can even be warped, and warped images can be morphed.

Paint Systems

Computer paint systems have a variety of important applications in digital effects. Matte paintings, which are entire background scenes that are typically painted on a large piece of masonite and then filmed, are more and more often created on computer paint systems such as Adobe Photoshop. Frame fixes can also be done, where small defects in individual frames of the film are touched up by hand. Done on a large scale, these frame fixes become entire film restorations. Wire removal, where the paint system is used to paint out supporting wires, rods, or robotic armatures, is another very important application.

Digital Ink and Paint

All feature film and most broadcast animation is painted, composited, and filmed digitally. The actual animation is still drawn on paper, but the paper drawing is then scanned into the computer, and from that point forward the entire process is digital. In addition to the character drawing, there is a second drawing called a "tone layer" that outlines the highlight and shadow areas of the character. This tone layer is painted as a mask, which is then used by the computer to highlight and shade the painted character at compositing time. Backgrounds and overlays are usually painted conventionally, then scanned into the computer for camera moves and compositing with the animation layers.

Input

Digitizing film into the computer is done with a machine called a film scanner. These are very large and expensive machines (\$100,000 and up) that typically digitize the film at 2048 pixels across and 1556 scan lines down. Scan times are in the range of two to four frames per minute, and the file size for one frame of film will be about 10 megabytes. Large digital effects studios usually own their own film scanners, while smaller companies take their film to be scanned at service bureaus and get the digitized film back on data tapes.

Digitizing video is considerably simpler, especially since most video formats today (D1, D2, DigiBeta) are digital formats to begin with. The videotape is loaded into the appropriate videotape player, which sends the digital video data to a DDR (Digital Disk Recorder) in realtime. A DDR is actually just a high-speed, large-capacity disk drive with some video circuits. From the DDR, the video data is transferred to the workstation over Ethernet. The video frames are 720 pixels by 486 scan lines, and each frame of video is about 0.7 megabytes in size. Most studios have their own DDRs and tape decks, but any video post-production house will do the transfer for a fee.

Output

Outputting the digital film data to 35mm film requires another large and expensive machine called a film recorder. There are two basic types of film recorders, laser and CRT. The laser film recorders read the digital data and "burn" the image directly onto the film with three colored laser beams: red, green, and blue. The laser film recorders are fast, very high quality, and expensive. The CRT film recorders place one color "channel" at a time, for example the red channel, on a black-and-white CRT with a red gel in front of the lens. The red layer of the film is exposed, the film is NOT advanced, then the green and blue channels are exposed. When all three passes are complete, the film is advanced to the next frame. The CRT film recorders are slower, lower quality, but cheaper.

Outputting the digital video back onto video tape is done with the same equipment used to digitize the video in the first place. From the workstation, each video frame is transferred to the DDR. When the entire shot is on the DDR, the videotape deck is rolled and the DDR dubs the shot onto the video tape in real-time.

Math

Introduction:

Behind all great (and not so great) computer graphics images and/or animations stands a great lady. Her name? Mathematics. Whether it is a spline or parametric equations that generate the curve, the various geometries that provide the illusion of 3D, or the vector theory used in reflections, rotations, and shading, just about all aspects of a computer-generated image rely on mathematics. This paper discusses some of the different mathematics used in generating graphics images, from basic Euclidean geometry to the basics of spline interpolation. We begin with an overview of the simpler concepts in computer graphics: line drawing and transformations. In section 2, we examine how the illusion of 3D is created. In section 3, we examine the mathematics involved in lighting our scene. And finally, in section 4, we examine creation of objects that do not have a simple geometric representation.

Section 1: The Basics

Computer graphics techniques have always relied heavily on mathematics for their implementation. Simple line and conic drawing algorithms^{1,2,3,4} show the importance of a thorough understanding of basic algebra, geometry, and multivariate calculus. Calculus? Yes, gradients are used to determine vectors perpendicular to the tangents to conics other than circles. This is used when determining which pixel to "turn on" to create the most realistic-looking curve. Even pattern filling and line clipping rely on these basic concepts. Then, of course, clipping requires knowledge of the windowing system and coordinate systems in general. So we're back to geometry again. When creating an object to be represented, a mathematical model representing this object must be created. Whether it is a simple 2D or 3D geometric figure, or a more complicated curve or surface, each object must have a mathematical representation. This representation is given in "real-world" coordinates: real values for which the mathematical model is valid. To create the graphics image to represent our object, we must change the real-world coordinates (window or user coordinates) of our object to the normalized device coordinates and/or the viewport coordinates. This change requires what is commonly referred to as a window-to-viewport transformation. This transformation of user coordinates (x_{user}, y_{user}) to normalized device coordinates (x_{norm}, y_{norm}) is a simple linear function that is a combination of scaling and translation:

$$x_{norm} = \frac{x_{user} - w_{x_min}}{w_{x_max} - w_{x_min}} (v_{x_max} - v_{x_min}) + v_{x_min} \quad y_{norm} = \frac{y_{user} - w_{y_min}}{w_{y_max} - w_{y_min}} (v_{y_max} - v_{y_min}) + v_{y_min}.$$

Here, w_{x_min} , w_{x_max} , w_{y_min} , and w_{y_max} , are the minimum and maximum window (user) coordinates in the x and y directions, and v_{x_min} , v_{x_max} , v_{y_min} , and v_{y_max} are the minimum and maximum viewport coordinates in the x and y directions.

Once we have the ability to represent our object, we may wish to change its position or size or inclination. Each of these transformations of an object, whether it is a translation, rotation, or scaling is accomplished with the assistance of vectors and matrices (linear algebra). These 2D transformations are given on the Conference Abstracts and Applications CD-ROM as Exhibit 1.

To rotate about a point other than the origin, say point P, we first translate P to the origin, rotate the object, and then translate back.

In order to treat all three transformations equally and to be able to easily compose them, we extend our 2D object to a 3D representation, by creating what are known as

Math

homogeneous coordinates: $(x, y)^T \rightarrow (x, y, 1)^T$. This is done so that the different functions that are to be applied to our object (or primitive) can be composed into one matrix transformation, and this one transformation is applied, rather than incurring the cost and time of applying the separate transformations to the object. The individual transformation matrices are given on the Conference Abstracts and Applications CD-ROM as Exhibit 2.

When working with 3D objects, we simply extend our 2D ideas to 3D. These transformations are applied to similar homogeneous representations: $(x, y, z)^T \rightarrow (x, y, z, 1)^T$. These forms are given on the Conference Abstracts and Applications CD-ROM as Exhibit 3.

Finally, to handle the generic case of rotating by an angle θ about an arbitrary direction given by the line segment beginning at $P_1 = (x_1, y_1, z_1)^T$ and ending at $P_2 = (x_2, y_2, z_2)^T$ we apply a composition of transformations. Set $(a, b, c)^T = P_2 - P_1$ and $L = \sqrt{a^2+b^2+c^2}$ and $p = \sqrt{b^2+c^2}$. Step 1: Translate the line and the point to be rotated to align with the z-axis. Step 2: Rotate the line around the x-axis until it is in the xz plane. Note that right-triangle trigonometry is used to determine the values for cosine (c/p) and sine (b/p). Step 3: Rotate once more to put the line on the z-axis. Step 4: Rotate through the desired angle θ . Step 5: Reverse steps 3, 2, and 1 to place the point back to its relative position.

Section 2: It's All an Illusion


The idea of trying to present the illusion of a 3D picture on a 2D surface had its beginnings in the Renaissance. The painters of the period utilized the idea of a "point at infinity" to represent the place where parallel lines seem to intersect. Mathematicians, in their attempts, to prove Euclid's parallel postulate developed non-Euclidean geometries^{5,6} – geometries that do not assume one or more of Euclidean Geometry's postulates – provided firm mathematical foundations for this idea of a "point at infinity." When we attempt to represent a 3D object on a graphics device, we are utilizing these concepts. There are two types of projections that are commonly used to create these representations: perspective and parallel. If the distance from the center of the projection to the projection (or view) plane is finite, then the projection is perspective; if the distance is infinite, the projection is parallel. Perspective projections do not preserve angles and distances, while parallel projections do. Figure 1⁷ provides an illustration of the two basic projection types. Since it is the perspective projection that creates the more "realistic" representation, the illusion of 3D, it is this projection method that we will discuss here.

We start by assuming that the projection plane is perpendicular to the z axis and is located at $z = d$. The center of projection is at the origin. To determine the projection $P_p = (x_p, y_p, z_p)^T$ of a point $P = (x, y, z)^T$ we use similar triangles to get $x_p = x/z/d$, $y_p = y/z/d$, and $z_p = d$. The division by z causes the projection of more distant objects to be smaller than that of closer objects.

A more general formulation was developed by N. Weingarten.⁷ We summarize this formulation here. We still assume that the projection plane is perpendicular to the z axis and is located at $z = z_p$, but now the center of projection is a distance Q from the point $(0,0,z_p)^T$, the intersection of the projection plane with the z axis. If the normalized direction vector from $(0,0,z_p)^T$ to the center of projection is $(d_x, d_y, d_z)^T$ then it can be shown that the projection P_p can be computed by a simple matrix product. This formulation is given in Exhibit 5 on the Conference Abstracts and Applications CD-ROM. This matrix transformation provides a one-point perspective projection. The vanishing point (or the point at infinity) is given by $(Qd_x, Qd_y, z_p)^T$.

There are many variations of this type of projection; for example, those that allow for multiple vanishing points. A summary of the mathematical effort required to transform a 3D world coordinate into a 2D device coordinate is given in Figure 2⁷.

Figure 1
Perspective Projection



Parallel Projection

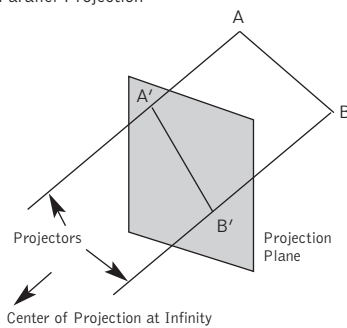



Figure 2



Section 3: A Little Light on the Subject

When considering the lighting of a scene, one must consider the interaction of light with the surface of an object in the scene and model this interaction. One must consider the light emitting sources (a point source or distributed source) as well as the light reflecting sources (ambient light or background light). There are also two types of reflection to be considered: diffuse and specular. In diffuse reflections, incoming light that is not absorbed is reflected off the surface in random directions, while in specular reflections the light reflects in a nearly fixed direction without any absorption.

In order to compute the overall illumination of a given point $P = (x, y, z)^T$, we need the following information:

1. The direction of the light source L (there could be more than one)
2. The normal vector N
3. The reflection vector R (there could be more than one)
4. The viewing vector V

If $S = (S_x, S_y, S_z)^T$ is the position of the light source and $E = (E_x, E_y, E_z)^T$ is the position of the eye, then $L = S - P$ and $V = E - P$. Each of L and V must be normalized, i.e., made to have length one. The computation of N and R can be complicated depending on the type of surface. N and R must also be normalized. Once L, V, N, and R have been determined then $\cos\phi = L \cdot N$ and $\cos\theta = V \cdot R$ and the Phong⁸ model of illumination can be computed by⁹:


$$I_* = I_{a*}k_{d*} + I_{p*}(k_{d*}\cos\phi/d + k_s\cos^n\theta)$$

where

- * = r, g, or b (for red, green, or blue)
- I_{a*} = intensity of ambient light for *
- k_{d*} = diffuse reflectivity of *
- k_s = a constant that estimates the specular reflection coefficient ($0 \leq k_s \leq 1$)
- I_{p*} = intensity of point source of light for *
- n = measure of "shininess" of the surface, very shiny = large n (> 150)
- d = distance to the point source or distributed source of light

There are many other models that could be considered, for example, GOUR71¹⁰, WARN83¹¹, VERB84¹², and NISH85¹³.

Figure 3



Section 4: The Spline's the Thing

When trying to represent an object mathematically, quite often it is impossible to construct an appropriate model using only simple geometric shapes. If this is the case, then arbitrary curves need to be created to model the object. The curves of choice are splines and other parametric curves. These curves are used in animation to generate in-between frames, or to create curves when only a few datapoints describing the curve are known. Yet another application of these curves is the generation of arbitrary surfaces. Each of these applications has essentially the same premise: from a few data points, generate a smooth curve based on this data. This curve should be simple to compute and easy to modify if a change in data is necessary.

Parametric equations are used to represent curves that are not easily represented or are impossible to represent by a function. Examples of this would be curves that have loops or places of infinite slopes. In this case we use a parameter, say t, and describe each variable x and y as functions of t, $x = x(t)$ and $y = y(t)$, over some range of t. If we are working in three dimensions then each of x, y, and z would be parameterized. The slopes of tangential lines for parametric curves are simple to compute. For example, $dy/dx = (dy/dt)/(dx/dt)$, i.e. the chain rule of differentiation becomes simply a quotient.

Splines are piece-wise defined polynomial functions of, usually, low degree. If you have cubic points $(x_0, y_0), \dots, (x_n, y_n)$, x_i s distinct, then the Natural Cubic Spline, $S(x)$, will interpolate this data (i.e., $S(x_i) = y_i$) and is given by cubic polynomials defined on each of the intervals $[x_i, x_{i+1}]$, $i = 0, \dots, n-1$. To determine the coefficients of this spline, one must solve an $(n-1) \times (n-1)$ tridiagonal linear system of equations. This spline, though historically significant, is not very useful in computer graphics. A unique spline does not exist if there is a repeated x value and if a data point needs to be changed then the entire spline will need to be recomputed.

To overcome these problems, splines that do not necessarily interpolate the given data are used. The most often used splines are the Bézier^{14,15} and uniform B-splines.¹⁶ Splines can be defined for any degree, though the cubic splines are usually sufficiently robust to handle the smoothness requirements in graphics.

The uniform cubic B-spline, b, stretches over five data points, called knots. The distances between the knots are all taken as 1 and the point 0 is put in the middle of these five knots to achieve symmetry. Formally this is given by:

$$\begin{aligned} b(u) &= 0 & \text{if } u \leq -2 \\ b(u) &= (2 + u)^3 / 6 & \text{if } -2 \leq u \leq -1 \\ b(u) &= (2 + u)^3 / 6 - (2 + u)^3 / 3 & \text{if } -1 \leq u \leq 0 \\ b(u) &= 2(1 - u)^3 / 3 - (2 - u)^3 / 6 & \text{if } 0 \leq u \leq 1 \\ b(u) &= (2 - u)^3 / 6 & \text{if } 1 \leq u \leq 2 \\ b(u) &= 0 & \text{if } 2 \leq u \end{aligned}$$

References

1. Bresenham, J. E. *Algorithm for Computer Control of a Digital Plotter*, IBM Systems Journal, 4(1), 1965, 25-30.
2. Bresenham, J. E. *A Linear Algorithm for Incremental Digital Display of Circular Arcs*. Communications of the ACM, 20(2), February 1977, 100-106.
3. Pitteway, M. L. V. *Algorithm for Drawing Ellipses or Hyperbolae with a Digital Plotter*, Computer J., 10(3), November 1967, 282-289.
4. Van Aken, J. R. *An Efficient Ellipse-Drawing Algorithm*. CG&A, 4(9), September 1984, 24-35.
5. Hartshorne, R. *Foundations of Projective Geometry*, Benjamin/Cummings Publishing Company, 1967.
6. Wylie, C. R. Jr. *Foundations of Geometry*. McGraw-Hill Book Company, 1964.
7. Foley, J. D., A. van Dam, S. K. Feiner, J. F. Hughes. *Computer Graphics: Principles and Practice*, 2nd Edition, Addison-Wesley Publishing Company, Inc., 1990.
8. Phong, B.-T. *Illumination for Computer Generated Images*. Communications of the ACM, 18(6), June 1975, 311-317.
9. Pokorny, C., Gerald, C. F. *Computer Graphics: The Principles Behind the Art and Science*, Franklin, Beedle & Associates, 1989.
10. Gouraud, H. *Continuous Shading of Curved Surfaces*. IEEE Trans. on Computers, C-20(6), June 1971, 623-629.
11. Warn, D. R. *Lighting Controls for Synthetic Images*. SIGGRAPH 83, 13-21.
12. Verbeck, C. P., D. P. Greenberg. *A Comprehensive Light-Source Description for Computer Graphics*. CG & A, 4(7), July 1984, 66-75.
13. Nishita, T., I. Okamura, & E. Nakamae, *Shading Models for Point and Linear Sources*. ACM TOG, 4(2), April 1985, 124-146.
14. Bézier, P. *Emploi des Machines à Commande Numérique*, Masson et Cie, Paris, 1970. Translated by Forrest, A. R., and A. F. Pankhurst as Bézier, P., *Numerical Control -- Mathematics and Applications*, Wiley, London, 1972.
15. Bézier, P. *Mathematical and Practical Possibilities of UNISURF*, in Barnhill, R. E., and R. F. Riesenfeld, eds., *Computer Aided Geometric Design*, Academic Press, New York, 1974.
16. Curry, H. B., & I. J., Schoenberg. *On Spline Distributions and Their Limits: the Polya Distribution Functions*. Bull. American Mathematical Society, 53, Abstract 380t (1947), 109.
17. Bartels, R. H., J. C. Beatty, B. A. Barsky. *An Introduction to Splines for Use in Computer Graphics and Geometric Modeling*. Morgan Kaufmann Publishers, Inc. 1987.

If we need to draw a curve based on knots $P_0 = (x_0, y_0, z_0), \dots, P_n = (x_n, y_n, z_n)$, then one way to construct the approximating curve is to form the linear combination:

$$C(u) = \sum_{i=0}^n p_i b_i(u)$$

where p_i stands for x_i, y_i or z_i . The most convenient way to draw the B-spline curve is to use a matrix product. This product expresses the value of C in any subinterval $[u_i, u_{i+1}]$ with i not 0 nor $n-1$ and is given as Exhibit 6 on the Conference Abstracts and Applications CD-ROM.

As the parameter u changes from 0 to 1 the formula generates the curve from P_i to P_{i+1} .

The other most commonly used cubic spline is the Bézier spline. The main difference between the two is that the Bézier curve is globally, not locally, controlled by the points P_i . Changing one of the control points affects the entire Bézier curve, but its strongest affect is in the neighborhood of the point. The matrix formulation of the Bézier curve is given on the Conference Abstracts and Applications CD-ROM as Exhibit 7.

To generate spline surfaces, one simply takes the Cartesian product of two splines. For example, we define a B-spline surface $b(u,v) = b(u)b(v)$. The matrix formulation of this surface is given by $S(u,v) = 1/36 UPM^T V^T$ where $U = (u^3, u^2, u, 1)$, $V = (v^3, v^2, v, 1)$ and

$$M = \begin{bmatrix} -1 & 3 & -3 & 1 \\ 3 & -6 & 3 & 0 \\ -3 & 0 & 3 & 0 \\ 1 & 4 & 1 & 0 \end{bmatrix} \text{ and } P = \begin{bmatrix} p_{i-1, j-1} & p_{i-1, j} & p_{i-1, j+1} & p_{i-1, j+2} \\ p_{i, j-1} & p_{i, j} & p_{i, j+1} & p_{i, j+2} \\ p_{i+1, j-1} & p_{i+1, j} & p_{i+1, j+1} & p_{i+1, j+2} \\ p_{i+2, j-1} & p_{i+2, j} & p_{i+2, j+1} & p_{i+2, j+2} \end{bmatrix}$$

A similar formulation for Bézier surfaces exists.


There are, of course, other splines and generalizations that can be used to model curves or surfaces, e.g., Hermite, non-uniform B-splines, and Beta splines.¹⁷

Section 5: Other Neat Stuff

We have only begun to scratch the surface of the different types of mathematics used in computer graphics, and we have omitted many of the details of the concepts presented in this paper. For example, most of the calculus concepts involved in the illumination section were omitted. However, it should be fairly obvious that geometry, linear algebra, and calculus play fundamental roles in the representation and manipulation of images. Other aspects of computer graphics often utilize more complicated mathematics. For example, antialiasing uses digital signal processing and hence Fourier transforms, and physically based modeling requires the solution of partial differential equations. To truly understand fractals, complex arithmetic and a little analysis will go a long way. In order to fully appreciate and understand the finer aspects of computer graphics one must at least appreciate, if not understand, the finer aspects of mathematics, as you would not have one without the other.

Many types of math (geometry, algebra, and calculus, among others) are crucial “tools of the trade” for people who create computer-generated effects. For example, animating people and creatures calls on methods from geometry and trigonometry. Rendering actual images requires algebra and arithmetic, among other things. And simulating physical phenomena like cloth and water uses knowledge from calculus. This presentation shows some examples from film and commercial effects that illustrate the use of these mathematical tools.

Tools



Minotaur: A Tactile Archaeology Game for Kids

Playground

This tactile educational game on Greek mythology for children aged 6-8 is designed for presentation in a museum setting. The intent of the project is to combine the tactile nature of museum interaction design with the computer's ability to create a stimulating audio-visual environment.

Unlike a traditional classroom setting where the course of learning is pre-structured and linear, an interactive environment allows students to choose their own paths, creating unique experiences. In essence, they become active participants in the learning process, not unlike archaeologists who unravel ancient mysteries with their brushes and spades. In the same spirit of exploration, the users of this game learn about ancient Greece and its myths through sight, sound, and touch.


The story of Theseus and the Minotaur was chosen for the game because it has all the ingredients of an adventure: a monster, a hero, romance, conflict, and the triumph of good over evil. This myth is unique in that it introduces users to Theseus, the greatest hero of Athens; Minos, the greatest ruler of Crete; and their respective cultures.

In *Minotaur*, users participate in an archaeological dig for treasures of ancient Greece. Their task is to find the six missing pieces of an amphora at the site and re-construct it. Each piece of the fragment represents a character of the myth depicted on the amphora. When a fragment is found and put in the right place, a mythical character, such as Ariadne, appears in the video projection, unveiling history, art, or myth-related information about the ancient world. When all six pieces are found, a short 2D animation of the myth is unfolded to the user as a reward.

Tactile interaction is made possible by embedding magnets in the fragments and magnetic switches in the amphora. An ADBIO box then relays the information to the computer, directing it to project short animated sequences onto the screen.

Reconstruction of the amphora mimics an archaeological dig; the users learn through physical interaction with the object, and segments of history are unveiled. By direct manipulation of the game (amphora fragments) rather than the mechanics (for example, the computer mouse), cognitive distance between intent and execution is minimized. Combining the tactile nature of museum design with the computer, *Minotaur* blurs the line between the digital and physical world, creating a more layered and fuller user experience.

This game is a graduate thesis project produced in partial fulfillment of the requirements for a M.F.A. degree from Pratt Institute.



Carla Roth
Think Jacobson & Roth
croth@primenet.com

*Classroom
Playground*

Panelists
Liz Caffry
Elisabeth Cameron
Los Angeles County Museum of Art

Carla Roth
Think Jacobson & Roth


Bill Zullo
Z-Digital

Museums are discovering that well-designed computer games that require active decision-making capture and keep children's attention. This panel looks at how curators, writers, interactive designers, producers, and educators, each with different agendas, work together to create an entertaining learning environment for young museum visitors.

The panel focuses on multimedia productions created specifically for a museum environment using game design methods often reserved for entertainment titles. By combining the entertainment value of new media with specific educational goals, museum multimedia games are powerful learning tools in an exhibition environment.

The production "Secrets of the Ancestors" is discussed by the curator, educator, interactive designer, and producer. This engaging computer game helps children decipher meaning in objects from around the world. The production, part of the Los Angeles County Museum of Art exhibition "Ancestors: Art and the Afterlife," is played at two kiosks and in a new interactive theater for school groups in the museum's experimental education gallery.

From multiple and sometimes conflicting points of view, the panel reviews the challenges and pitfalls of developing interactive game-based media for a museum environment.



Organizing Summer Computer Graphics Camps

Classroom

This presentation explains how the annual Computer Graphics Summer Camp at Purdue University is organized and operated. The discussion includes how students are identified for inclusion in the program, funding issues, computer project sessions, student participation, software and hardware uses, tours, and entertainment. While this camp was previously operated as a regional resource, it is now open to students from all U.S. states and territories. The camp's success has led the sponsoring department to extend the program to include professional educators in computer graphics, math, and science in their own camp to increase their technological literacy.

Camps

People in the Past: The Ancient Puebloan Farmers of Southwest Colorado

Theresa Breznau
Living Earth Studios, Inc.
livearth@lasal.net

Playground


People in the Past: The Ancient Puebloan Farmers of Southwest Colorado is currently installed as an exhibit at the Anasazi Heritage Center in Dolores, Colorado, and was exhibited at the Houston Museum of Natural Science during the summer of 1998. It also won both first and second place in the 1998 National Association of Interpretation Awards (first place for the CD-ROM and second place for its companion teachers guide, "Classroom Activities"). This program is available as a CD-ROM for both Windows and Mac. The "Classroom Activities" guide for grades 4-12 is a 160-page book with lesson plans related to the program contents for language arts, math, social studies, science, and art. The CD-ROM is included with the book.

The program explores Lowry Pueblo, a prehistoric Anasazi village in southwest Colorado, during the last days of an archaeological excavation. The Pueblo and the entire site have been recreated in 3D. As assistants to the lead archaeologist, users explore the area, the Pueblo, and the archaeologist's tent, helping out with such tasks as piecing ceramics together, taking photographs, and excavating artifacts. In QuickTime videos, several archaeologists speak to the assistants as they explore. Hopi and Tewa people, descendants of the Anasazi who are the Puebloan people of today, speak throughout the program about their ancestors who lived in the village. They provide insights into their daily lives and culture.

In many areas of the program, the assistants are able to go back in time, to AD 1125, and see life as it might have been lived through narrated 3D re-enactments and animations. Special features of the program include an animated corn grinding scene, recreations of rooms in the Pueblo, recreations of both the plaza and the fields throughout a full year's seasons, and a nature walk through the canyons where participants can find the ancient Puebloan's water source, a rock art panel, and native plants. The archaeologist's tent is filled with books, maps, a microscope, a computer, a CD player, and a VCR, all of which the assistant can explore to uncover the methodology of contemporary archaeologists.

"People in the Past" was created for the Anasazi Heritage Center with grants from the Colorado Historical Society. It was produced by Paradox Productions and Living Earth Studios, Inc.

Ancient



The Area Map of the Entire Site



Entering the Archeologist's Tent

Proposal Writing 101: Ensuring Your Submission is Understood

Classroom

Few things are as frustrating as submitting a proposal to a conference and being turned down because of the ways in which your ideas were presented. Often times, the ideas or projects themselves aren't the problem. It's the way the proposals are written that frequently disqualifies a submission. I know the feeling. When I've had projects rejected, I wonder, "Why didn't the reviewers realize how perfectly my work fit with their conference?" or "Gosh, was my idea that bad?"

After serving on numerous selection committees for major conferences, and playing a role in the SIGGRAPH conference selection process for the past three years, I have some suggestions that will help submitters increase the chance their proposals will be accepted to a conference. Drawing upon many of the concepts I teach in rhetoric courses, this essay describes some practical ways to ensure that your proposal will be considered on the merit of its content for inclusion in a conference such as SIGGRAPH.

There are two major sections to this paper. First, I'll describe how you can use the Call for Participation to tailor a proposal to be a "perfect fit" for both the conference and program. Secondly, I'll highlight some basic writing techniques that will help reviewers quickly grasp your main ideas.

Using the Call for Participation

This may sound silly to some, but a key way to make sure your proposal will "fit" within a conference program is to create a checklist of what needs to be included. At SIGGRAPH, the program chairs already create that checklist for you: the Call for Participation (the Call).

The Call is created to invite proposals and to describe the content the program chair is seeking for the conference. The Call is designed to help streamline the process of reviewing submissions. Contrary to being an arbitrary set of rules that you can ignore, the Call is important because it asks for exactly the information the program chair needs in order to make a decision. Contributors to a conference such as SIGGRAPH are asked to adhere to a set of specific guidelines; the Call is the first of these requirements. Your proposal is the first indication of how well you will pay attention to details later in preparation for the conference. Submitting a proposal that is difficult to understand, poorly written, or that fails to include all requested information decreases your chances of being accepted regardless of how good your idea is.

When writing a conference proposal, perhaps the most important consideration is making sure you complete the requirements put forth in the Call. It's surprising how many rejected proposals are simply due to a failure to pay attention to the Call. Don't fight or ignore the Call. Instead, use it to your advantage!

Perhaps the best way to use the Call to your advantage is to make sure you fulfill all of its requirements. I suggest treating the Call as a strict set to guidelines. Always keep it in front of your eyes, and even commit parts of it to memory.

Step One: Obtain the Call for Participation

Make a copy of the Call. Use it as a draft to help you prepare your proposal. Different conferences may ask you to submit your proposal online or to mail a paper copy.

How can you use the Call to help you prepare? It's simple: write all over it! Grab a highlighter before you ever start writing, and highlight all of the requirements put forth in the Call. If you'd like, you can even create a checklist and cross each item off the list once you've completed it. When I'm writing, I tend to highlight the requirements on a copy of the Call, and make notes to myself about what I think that means.

Step Two: Identify the Requirements

Distinguish between the types of requirements. Generally, the requirements for the SIGGRAPH conference fall into two categories: technical requirements and content requirements.

Meeting Technical Requirements

The technical requirements have little to do with your ideas, and much more to do with the format in which they are presented. To determine what the technical requirements are for the program to which you are submitting, start by looking at the actual checklist within the Call. These tell you what is required in order for you to have a complete submission. Without a complete submission, it's highly unlikely that your proposal will be accepted.

First, and perhaps a bit too obvious, make sure that you provide the complete contact information for yourself and everyone who will participate in your presentation. You'd be surprised how many submissions include incorrect or incomplete information. The more ways you provide for the program chair to contact you, the more interested you seem in having your submission included in the program.

Second, be sure to include a signed copy of the Submission and Authorization form. Also known as the permission-to-use form (PTU), it gives SIGGRAPH permission to print the information in your proposal in the conference publications. In order for a submission to be included in the program, it must have a signed PTU on file.

Next, most programs require at least two abstracts, a long and a short version. Be sure that your abstract provides a brief yet complete summary of what you'd like to cover in the presentation. Also (very important), make sure that your abstracts fit within the required word (or page) limits. It's very difficult for reviewers to sort through abstracts that are too sketchy (they don't provide enough information to tell what will be talked about) or too detailed (remember that reviewers may be evaluating many proposals so even the long abstract should be clear and concise). The assigned limits should be adhered to pretty strictly. That way you are consistent with what is requested in the Call.

Fourth, if the program to which you are applying requests supplemental materials, be sure to follow the guidelines within the Call for those materials. If you deviate from the acceptable formats, or don't provide enough copies, it is likely that your wonderful supplements won't even reach the reviewers. Also, since the timeline for reviewing is often very tight, if you would like to have your "extras" included in the review process, you should make sure that they arrive at the address in the Call on or before the due date. Otherwise, once again, it is likely that the reviewers won't see those materials.

Meeting Content Requirements

Meeting the technical requirements is only one part of using the Call to your advantage. Once you are sure that you've made a checklist for what is required, now it's important to make a list of how things are required: the content requirements.

What I mean by content requirements is this: program chairs have certain things they are looking for in a proposal. Generally, this information can be found within the part of the Call that describes the program itself. Often, the types of content are directly stated when the program is described. For example, the 1999 Electronic Schoolhouse Call asks potential contributors to "share how you teach computer graphics and/or use computer graphics to teach at all levels and across all disciplines." This tells us directly that in this program, proposals must at least deal with computer graphics and their relation to education. Make a note to include a discussion of that relationship. While it sounds obvious, a number of submissions to the Electronic Schoolhouse did not include a reference to one (computer graphics) or the other (education).

Proposal

Beyond making note of the obvious clues within the Call, you can also look for particular words within the description of the program that seem to stick out when you read them. For example, you might notice that the authors use words like “innovative” or “cutting-edge.” These words give you clues as to what types of submissions will be chosen. Add these words to your list of things to include in your proposal. We’ll come back to them later when we begin the actual writing of the proposal.

Step Three: Post your Lists

Before you begin writing, I suggest completing one more task. Take the lists you’ve made of technical and content requirements and put them somewhere so you can see them while you’re writing. I actually put the Call (or list) on the wall over my desk so that every time I look up from my computer, I see the requirements. That serves as an effective reminder that I have to work within the guidelines that are constantly in front of my face.

Step Four: Write the Proposal

At this point, all the necessary leg-work to write is complete. We’ve completed lists for both technical and content requirements for the program, so we know what types of information we should include, some key words or phrases to use, and how much space (or how many words) we have to work with while writing the proposal. Now let’s work on style.

Develop a Clear Thesis

One of the most common complaints I’ve heard from reviewers over the past several years is: “From this proposal, I can’t tell what this person wants to do.” In spite of telling submitters what they wanted to hear in the Call, reviewers and program chairs still don’t receive the types of proposals they need in order to make decisions. I believe this complaint is due to the writing style of the submitters, perhaps even more than the specific content. This final section suggests some stylistic approaches to proposal writing that will immediately answer the “what does this person want to do” question.

The most important part of your entire proposal is a short statement (the thesis statement) that summarizes what you would like to present. The shorter and clearer you can make this statement, the more effective its impact will be. Often times, the “I can’t tell what this person is saying” proposals either lack a thesis statement, or state it in a very complex or convoluted manner. Developing one short, clear statement that tells the reader what to expect avoids any confusion. Generally, I will spend up to an hour working on a thesis statement that will work for the given essay. Time spent developing a clear thesis will always pay off.

Write Deductively

Related to developing a clear and concise thesis, the order in which you present your ideas is also very important. Since proposals are usually read very quickly, it's really important that they be written deductively, and not inductively. This means that you place the main point of each paragraph in the first sentence. Writing deductively allows the reader to quickly glance at your paper and still "get the gist" of what you're trying to say even if she or he does not have much time to devote to your proposal. After the main point of the paragraph there should be an explanation of why that point is important: expand the statement you made in the first sentence. Next, provide support or examples that "prove" to the reader why what you've said is true. Finally, after the examples, re-state the main point and link it to the next point you'd like to make. As a general rule, here's the formula for deductive writing: Make a statement. Explain what it means. Support it with examples. Summarize and move to the next statement.

Use the Appropriate Vocabulary


Another frequent problem in proposal writing is that an author tries to write using a particular type of jargon. Of course there is nothing wrong with using technical words in a proposal, and a conference like SIGGRAPH needs to include certain jargon. But it is important to realize that the more technical your word choices, the more likely a reviewer may not understand your topic. So if you decide that you need to use technical words, be sure to explain what a given word means immediately following the first time you use it. Only then can you be sure that your audience will understand your words.

The most important thing to focus on when making specific word choices is clarity. Can someone understand what you "want to do?" If the words you choose to use don't help make your proposal clearer, then you're using the incorrect vocabulary.

Step Five: Using Your Checklist

After you've completed your writing, making sure that you have a clear thesis, you've written deductively, and you've chosen an appropriate vocabulary, it's time to revisit the checklist you created when you first read the Call for Participation. Double check to make sure that you've included each point on the list. If you haven't, you need to re-work your proposal so that the missing item is included. Once an item has been accounted for, cross it off your list. You know that the proposal is completed when you have no more un-crossed items. The advantage of this system is that you can constantly see where you are in the process of developing your project.

Once you've completed these steps, the proposal will be complete, and hopefully easily read. At the very least, by following these guidelines you can be confident that your submission contains all the necessary content and technical elements, and that the reviewers will at least take your proposal seriously enough to read it. Then you know that you have given the proposal your best effort. By following a systematic approach to proposal writing, the chances of reviewers being able to quickly and easily understand your proposal dramatically increase. Now your proposal can be measured on the strength of your ideas. Hopefully that means I'll see you at SIGGRAPH 2000!



The Round Earth Project: Collaborative VR for Elementary School Kids

This paper discusses deployment of a collaborative VR environment in an elementary school to help teach children that the Earth is spherical.

The concept of a spherical Earth and the implications of that fact are well-represented in the AAAS Project 2061 Science for all Americans report, and the difficulty of teaching it has been documented by Nussbaum, Vosniadou and Brewer^{3,4}. It is not a simple concept for children to acquire. Their everyday experience reinforces their deeply held notion that the Earth is flat. More precisely, their mental model of the world separates "sky" and "earth" into two parallel layers, one "above" the other; the two directions "up" and "down" are absolute. Telling young children that the Earth is round does not cause their intuitive model to be replaced by a spherical conception of the Earth. Instead, children assimilate the new information into their prior knowledge and often conclude that the earth is flat and circular.

The spherical Earth is a simple example of a deep idea, a fundamental concept that lies underneath our extensive system of domain knowledge and influences how experience and discourse are conceptualized. Revising these deep ideas, these core concepts, is difficult because new knowledge is assimilated in terms of existing knowledge. When the new knowledge is both different from and more fundamental than the existing knowledge, the typical outcome is distortion.

The Round Earth Project is a collaboration among researchers in computer science, education, and psychology investigating two alternative pedagogical strategies for teaching children that the Earth is spherical, and the implications of that fact. The transformationalist strategy attempts to effect conceptual change by evidencing a breakdown in the children's prior models. The alternative displacement strategy attempts to effect learning in an alternative setting free of pre-existing biases, and to relate that learning back to the target domain: the Earth. In the transformationalist approach, VR simulates the launching of a spacecraft from the Earth's (apparently flat) surface and subsequent exploration within a fixed-height orbit. In the displacement approach, VR simulates a small-diameter asteroid where the learner walks on a curved horizon, sees objects appear from below the horizon, takes a long walk around the globe and comes back to the departure point.


The virtual worlds are collaborative. One child experiences the surface of the world as an astronaut, and the other, acting as mission control, sees the avatar of the first child on the spherical world. The kids are given a task that requires the astronaut to move around the spherical body. This way, the astronaut is often "upside down" on the sphere but "rightside up" on the surface. The task fosters positive interdependence because neither child can perform the task alone. They need to cooperate and communicate, and through this communication the children must reconcile their different views.

The children are first given a five-minute orientation to each view, which shows them how to use the controls and points out features of the environment that are important to the learning goal (the ability to circumnavigate the globe, seeing the tops of objects over the horizon first, not feeling "upside down" when you are at the South Pole). Each of the children then experiences each role for 10 minutes.


The displacement strategy requires a second step. This new knowledge, established at the alternative cognitive starting point, must be brought into contact with the child's prior knowledge. The point is not just to know what it would be like to walk on a spherical planetary body, but to understand that the Earth is such a body. We call this second step bridging activities. It involves talking to the children about their VR experiences for 10 minutes using a physical model of the Earth and the asteroid.

We have previously described our pilot studies, which were conducted by bringing children to VR equipment in the laboratory ^{1,2}. For these actual studies, we worked in close cooperation with the teachers and administration of a local elementary school. An ImmersaDesk driven by a Silicon Graphics desk-side Onyx and a stereo-capable monitor driven by a Silicon Graphics Octane were brought into a classroom in the school for two weeks, and studies on the displacement strategy were conducted onsite. The ImmersaDesk was used for the astronaut view, giving the user a wide field of view on the surface, while the stereo monitor was used for mission control.


A guide introduces a student to his role as the astronaut at the ImmersaDesk.



A guide introduces a student to her role as Mission Control at the stereo monitor.



The astronaut approaches one of the crystal forests on the asteroid.



The happy astronaut.

The Round Earth Project: Collaborative VR for Elementary School Kids

There are 84 second graders in four classrooms at this school. Seventy-six permission slips were returned (a pretty high return ratio), and all of these children took the 20-minute pre-test. The pre-test consisted of 18 questions spread over five topic areas: the sphericity and support of the Earth, the relativity of up, circumnavigation, occlusion, and egocentric vs. exocentric perspectives. These questions were asked three ways: verbally, with 2D paper drawings, and using 3D PlayDoh models. This was done to minimize representational bias. We developed a simple scoring system and divided the children into three groups: the high group answered 14 or more correctly, the intermediate group answered 11-13 correctly, and the low group answered 10 or fewer correctly.

The 29 children in the low group were chosen as the treatment group for the VR experience. From our previous experiments with third-grade children at another school, we expected to have a larger subject population. Because we only had 14 pairs of children, we had them all experience the displacement-based asteroid world. One week later, randomly chosen pairs of these children went through the 30-minute VR experience and the 10-minute bridging activities. They were given the post-test on the next day.

During development, we were concerned about whether we would be able to create a compelling alternate reality. This did not seem to be a problem, as several of the children reported being scared when they first stepped onto the asteroid in front of the ImmersaDesk, thinking that they would fall off the nearby horizon. One of the children was unable to continue as the astronaut. Another child reported being dizzy at mission control, but wanted to continue. Overall the children had little difficulty using the apparatus, but the adult guides who handled the orientation session stayed with the children during the VR experience to provide guidance.

As with our previous studies, the children became very focused on completing their task before their time ran out. Because of this, we waited until after the orientation phase to give them their mission. In the case of the asteroid, the two children are told that their spaceship has crashed on this asteroid, and they need to recover fuel cells scattered about the surface to allow their spaceship to leave. Once they had their mission, the children's dialogue tended to focus on getting to the next fuel cell. Mission Control would frequently tell the astronaut to "go up" or "go down" with reference to the sphere, which then initiated a conversation as the astronaut tried to map this direction into moving forward, turning left, or turning right on the surface.

The 22 children in the intermediate group became the quasi-control group. These children were given the post-test without having the intervening VR experience. To be fair, once these post-tests were performed, the children in the intermediate and high groups were given a chance to experience the VR worlds.

VR

The score for the treatment group increased from a mean of 7.3 correct out of 18 to 12.9 out of 18, and this change was statistically significant. The questions on the post-test were identical to those on the pre-test, so there is the possibility that the test itself contributed to learning. The quasi-control group, who didn't have the VR experience, saw their scores increase from a mean of 12.2 on the pre-test to 14.0 on the post-test, and this change was also statistically significant, but the magnitude was considerably smaller. Additionally, the difference between the treatment group and the quasi-control group was statistically significant on the pre-test, but there was no significant difference between the two groups on the post-test. This suggests that the combination of the VR experience and bridging activities brought the treatment group up to the level of their classmates in the quasi-control group.

Compared to our previous studies in the laboratory, this study in the classroom went much faster and required fewer personnel. Our previous experience with taking ImmersaDesks to conferences made the deployment to the school quite straightforward, and the 15-year relationship one of the investigators had with this school provided an environment of trust that was invaluable.

This study also showed that the children learned more than in our pilot studies, which we believe is a direct result of the changes made during those pilot studies. For this study, we were able to use the adult guides more effectively at the beginning of the experience to point out features of the environment that were important to the learning goal. We believe that this orientation helped the students bridge the gap between the two representations of the astronaut on the asteroid, making it easier for them to relate the two heterogeneous perspectives. We also modified our tests to ask the same question three different ways to avoid the biases that the media introduced into the questions and get a better idea of the child's model.

The children seemed very excited by the experience, and as word spread through the school many children and teachers from other grades came by to see what we was going on. We were concerned that the children might be jaded by their familiarity with video games, but several children favorably compared our setup to a Sony Playstation. There was great interest among the children, teachers, and staff in having us return to the school for future work, and we plan to return in the late spring to continue this investigation. This will include giving a delayed post-test for the children who participated in this experiment, comparing the transformationalist to the displacement approach used here, and investigating the relative influence of the VR experience and the bridging activities.


Acknowledgements

This research was supported by funding from the National Science Foundation (award EIA 9720351 - Deep Learning and Visualization Technologies).

The ImmersaDesk is a trademark of the Board of Trustees of the University of Illinois.

References

1. Johnson, A., Moher, T., Ohlsson, S., Gillingham, M. *The Round Earth Project: Deep Learning in a Collaborative Virtual World*. Proceedings of IEEE VR99, Houston, March 13-17, 1999.
2. Moher, T., Johnson, A., Ohlsson, S., Gillingham, M. *Bridging Strategies for VR-Based Learning*. Proceedings of CHI 99, Pittsburgh, May 15-20, 1999.
3. Nussbaum, J. *The Earth as a Cosmic Body*, pp. 170-192. *Children's Ideas in Science*. Milton Keynes, UK: Open University Press, 1985.
4. Vosniadou, S. and Brewer, W. *Mental Models of the Day/Night Cycle*. *Cognitive Science*, 18: 123-183, 1994.




The SIGGRAFFITI Wall: Multi-Input Painting

The SIGGRAFFITI Wall is a 3D virtual canvas that allows simultaneous input by multiple artists through various input devices. Some input devices are 2D, while others are 3D. The 3D input devices are used to create a virtual sculpture, while the 2D input devices are used to create texture maps on transparent polygonal virtual walls surrounding the sculpture.

One input device is a large wall space allowing artists of all ages to make their “marks” by throwing beanbags at the wall. As each beanbag hits the wall, a paint-splat primitive with time-varying characteristics appears on the transparent, 2D texture-map. Meanwhile, an artist using a glove input device is creating the 3D sculpture in the center of the virtual creation space, which onlookers view through the paint-splat wall. Other users may be finger-painting on the 2D texture maps from touch-screens, or drawing with lasers. The works of art can be saved and printed out by the hour or the day.

This is an ongoing creative process, created live at SIGGRAPH 99, just for the fun of it!

Painting



Introduction

In most classrooms, subjects such as Shakespeare and philosophy are associated with the words "boring," "haughty," and even "incomprehensible." Using the Internet and 3D graphic technology, students can now be transported to a 3D interactive world where these subjects become "real."

Shakey's Place 3D (library.advanced.org/10502/index.htm) and The Lighthouse (library.advanced.org/18775/index.htm) are Web sites designed to breathe life into Shakespeare, philosophy, and the students who study them. The sites' use of 3D graphics creates striking environments meant to peak interest and increase the educational value of the interactive devices they contain.

Background

In 1996, three students decided to attempt to free students from "painful" lectures and readings about Shakespeare by creating a 3D world that any student could visit and, more importantly, learn from. As part of the ThinkQuest Project, William Shakespeare's Globe Theatre was rebuilt in 3D and placed online. Students and teachers alike could visit and understand that Shakespeare, although a stereotypically boring subject, could be interesting. The result: Shakey's Place 3D.

SP3D was a success, not only in the ThinkQuest Project, but also in its intent. It received second place in the ThinkQuest competition, earning \$48,000 dollars in scholarships and grants. More than three million visits have been logged since its release in August 1997. Teachers, students, academics, and enthusiasts have walked through the theater, and its influence is still growing.

With the success of SP3D, the idea was expanded to philosophy, as a new entry in the ThinkQuest contest: The Philosophers' Lighthouse. Here, the graphics were improved, and, once again, a topic was revitalized with 3D technology and Web design. The Lighthouse has prompted a great amount of discussion and response since its release in August 1998 and has been selected as an Honorable Mention in the 1998 ThinkQuest contest.

In-Depth


The success of these Web sites has prompted a new method of teaching students, where field trips can be an inexpensive treat, where new ideas are bolstered and encouraged, even where collaborative study moves beyond the boundary of the classroom and on to a global realm. As Shakey's Place decrees in its introduction: "It's not what you need to learn... it's what you are inspired to learn." The worlds created in these sites draw the user in beyond the boundary of stereotypical "forced" education. Users are enticed further into the site, to explore. They can then use the tools without pressure or boundaries, thereby strengthening the educational value of the interactive devices. A few of these devices include:

- Madlibs (SP3D), where users are prompted to type random words, then the Web site replaces keywords in, for instance, a soliloquy by Hamlet with the random words provided. This shows the importance of iambic pentameter, poetic meter, and especially how words "paint" emotions. It is a perfect example of "effortless retention." Students are much more likely to remember if they have a "fun" example to remember, rather than just their notes.



SP3D and the Lighthouse: Explorations in 3D Internet Learning

- Online Auditions (SP3D), where users choose a "face" from a gallery of graphics, then decide which part they would like in which play. This creates not only a one-on-one relationship with the student, but also provides teachers and researchers statistics on which characters certain age groups associate with. Thus, not only does the Web site create a personal link with visitors, but it also gives teachers that opportunity as well.
- Voting Booths (Lighthouse), where visitors can vote on which philosophy or philosopher makes the most sense to them. This gives value to visitors' opinions and causes them to think, compare, and contrast, then vote.
- Bulletin Board Systems (both Web sites), where visitors are encouraged to collaborate on topics and formulate opinions. These bulletin boards broaden the horizons of the sites, perhaps even bringing to light topics that were not approached on the sites before.




These tools would be useless, however, without a strong hook to interest the user. The 3D environments draw the user in by making navigation easy, minimizing unnecessary reading, and simulating a field trip. Teachers can lead their students into a computer lab or classroom and walk them through an exciting world, teaching them and entertaining them at the same time. Enthusiasts of all ages can research or discuss their topics without being slowed by continual searching or reading. SP3D and The Lighthouse also use technology to encourage users to respond in a way that other sites and media cannot offer. These sites do not tell users what to think. SP3D and The Lighthouse bring a whole new meaning to the word "interactive."

Plans

Using the same concept, these principles of design and education are being applied to another site. In this trio of sites, each site will be unique in its subject matter, but they will share the same ideal: using 3D graphics as an interface to bring difficult or unpopular subjects to light. The new site is displayed alongside SP3D and The Lighthouse as a case study, to show how these sites come into existence.

The software, hardware, and World Wide Web issues involved are explored in an effort to encourage similar efforts and use of the Web in education. Perhaps someday trips to ancient Rome, colonial America, or even mystical Atlantis will be possible from a classroom, and learning will be inspired, not required.



This workshop demonstrates the Northwestern University Collaboratory Project's MediaSpace and other collaborative environments, and provides participants with the chance to develop a multimedia "hub" in MediaSpace.

The Collaboratory Project (collaboratory.nunet.net/) is a Northwestern University initiative funded by a grant from the Ameritech Foundation. Its staff provides consulting, training, technical support, and information services to education, cultural, and nonprofit organizations interested in using network technologies to advance education. The Project's goal is to establish an easy-to-use, network-based collaborative environment that enables organizations in the greater Chicago region to work together to share information, resources, and expertise.

As part of the University's Information Technology organization, the Collaboratory Project draws on the experience and technical expertise of the entire organization to support its mission. It works with individual educators, school project teams, and multi-school collaborations in the Chicago Public Schools system and surrounding school districts. In addition, the Collaboratory Project works with museums, libraries, and cultural institutions to develop innovative Web-based educational resources.

Within the project's collaborative environment, various collaborative communities enable people and organizations with common interests to share projects, ideas, information, and resources. The communities are Web-based activities that participants can join as part of established curricular activities. They enable educators and students to share projects through dynamically created Web pages and discuss their ideas using Web-based discussion and chat resources. They also provide gateways to Internet resources and facilitate access to experts in the region.

An educator can join a collaborative community from the Collaboratory Project Web site and use it to meet specific curricular needs. Students can participate with anything from a single computer and a modem to a fully networked computer lab or school. Through a network connection, they submit text, graphics, sound, and video files from a standard Web browser to a database on a Collaboratory Project server. Educators can manage the content their students submit through other Web pages that provide password-protected access to the database. Web pages are created dynamically from the database.

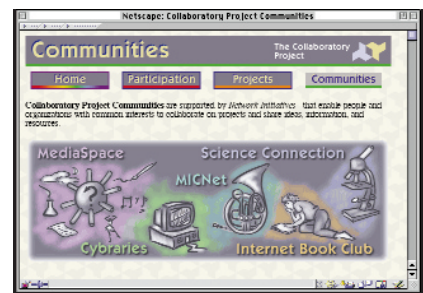
Collaborative communities available on the Collaboratory Project Web site include:

Cybraries

A Cybrary (Cyber Library) enables educators and librarians to identify, organize, and share Internet resources that support curricular activities. Teachers, librarians, and students can review and contribute useful URLs to a shared virtual library or create a new one to meet their own needs. The Collaboratory Project provides training on how to locate Internet resources, create online references, and prepare materials for network access and distribution as a Cybrary.

Internet Book Club

Designed for the K-12 curriculum, the Internet Book Club has four project areas where students and teachers can share language arts activities. Students can post book reviews, stories, essays, and original compositions and poems. Teachers can share project ideas and curricular materials in the Teachers Lounge. In particular, they can search participant databases to find other teachers who are using the same books and plan threaded discussions and online chat sessions for their students.



Online

The Science Connection

The Science Connection brings together science activities and resources. Northwestern University scientists and engineers, who answer questions from teachers developing classroom science activities, support "Ask a Scientist." Observation, data collection, and reporting projects that support scientific inquiry are being developed with Chicago museums. Teacher resources for science fair projects, reference materials, and more, are also available.

Music Internet Connections (MICNet)

MICNet is an Internet music collaboration project that supports school music teachers interested in exploring music composition activities with their students. Students can exchange music compositions in MIDI files and play and discuss those compositions from networked computers. Professional composers provide feedback to student composers.

MediaSpace

Media Space is a collaborative environment designed to encourage sharing of multimedia information, projects, and activities. Participants create "electronic multimedia postcards" using text, graphics, sound, and/or video files they have created. They can contribute a project to an existing theme or start a new theme and invite others to participate. MediaSpace supports a number of projects, for example:

- Community-based research projects for schools along Interstate 57 in Illinois.
- Fairy tales written and illustrated by students around the world.
- Student explorations of their musical heritage.
- Reports on places students have visited.
- Documentation and discussion of a teacher's visit to Swaziland, via a live link to an Internet café in Swaziland.
- Creating and "documenting" imaginary cultures.

Educators can join an existing activity or start new MediaSpace projects from the Collaboratory Project Web site and use it to meet their particular needs, such as science reports, multimedia essays, history reports, personal journals, etc. MediaSpace emphasizes the use of images and sound in preference to text.

Within the Collaboratory Project's shared environments, responsibility for content is distributed to educators and other interested parties while the server manages page layout, graphic design, and navigation using templates created by project staff. Each of the environments provides an easily extended and replicated model for networked collaboration. Because collaborative communities are participant-driven and used in projects that are closely tied to the curriculum, they provide fertile environments for developing innovative, collaborative projects and activities.



In his article "As We May Think," Vannevar Bush proposed a device called the MEMEX as a solution to the problem of information overload. Today, those teaching computer graphic design, animation, non-linear video production, multimedia and Web page design, sound editing and composition, or interactive art face a related problem of the information age: ongoing obsolescence of knowledge, skills, expertise, even aesthetics, and the imperative to continually upgrade and acquire new skills and knowledge to keep pace with the march of technological innovation and cultural fashion. In institutions of higher learning, Moore's Stairmaster favors the young, the fit, and the tireless. The World Wide Web may be a partial realization of Bush's vision for managing the explosive growth of information. For the educator, there appears to be no relief from the perpetual scramble to match teaching and research to the demands of a client-centered model of higher education.

Many credit Bush with the invention of hypermedia technology, which later inspired Tim Berners-Lee's proposal for the World Wide Web. Motivated by concern for the loss of scientific information, Berners-Lee explains, "...the dream behind the Web is of a common information space in which we communicate by sharing information." Yet this vision of the Web as a solution to the problem of information overload, management, and retrieval offers little relief for the educator suffering a mid-life crisis.

Moore's Law is unforgiving. As the power and complexity of silicon chips doubles every 18 months, each previous generation of computers and the software applications designed for them are rendered obsolete. The marketplace drives the proliferation of new features in upgrades and proclaims: upgrade or die! Software development resembles an autocatalytic chemical process, which increases in speed according to the volume of products it has created. With each hardware and software upgrade, there can be a concomitant obsolescence of skills. Educational institutions are hard-pressed to find the funds to cover costs of the next cycle of upgrades. Administrators are often reluctant to pay to retrain faculty and find it more expedient to hire recent graduates or expatriates from industry who are already expert users.

High-end software packages from Softimage, Discreet Logic, and Alias|Wavefront have steep learning curves that demand a tremendous time commitment. For a mid-career educator with full teaching, administrative, and personal responsibilities, this time is hard to find. Specialization is nearly unavoidable. If one chooses to maintain skills in Web site development using HTML, JavaScript, Lingo, Shockwave, Flash, VRML, CGI, and Java, there may not be time to tackle Softimage. A few years ago, it may still have been possible to maintain up-to-date working knowledge of a number of commonly used packages. Today there are Adobe Photoshop experts who may have never heard of VRML. Like increasing the resistance on the Stairmaster, the need to simultaneously maintain and reinforce existing skills and acquire new ones is an uphill struggle.

The client-centered model of education tends toward training on demand. Especially with complex software, the desire to learn how to do something often upstages questions about what to do and why. From the viewpoint of protocol analysis, much end-user software expertise is essentially scripted or procedural knowledge. Almost all tutorials follow this step-by-step cookbook recipe approach. This "training" approach lends itself to distance-learning strategies, which may have the secondary effect of reducing a teacher to the role of a technical troubleshooter. For some administrators, this diminished role is a cost-cutting opportunity.

Stairmaster

The Teacher's Mid-Life Crisis: Moore's Stairmaster of the Fittest

Theory is compressed into practice. Critical thinking is often shunted aside because students know what they like. The cultural paradigms that drive young imaginations are increasingly drawn from popular culture. There is a clear feedback loop from industry to the schools. The window on popular culture is a shifting frame of reference that increases the gap between the educator and the student regarding what is "kewl" to do. Like Dawkins' selfish-gene, new aesthetic memes are rapidly eclipsing the cultural icons of earlier generations. Influences can be traced from games and films like *Mortal Kombat*, Japanese anime culture, or the continued predominance of DJ club culture. Interest in these stylistic approaches shows no sign of abating and, in fact, continues to spread among the youth cultures of the world. Rather than hoping to be a Rembrandt or Beethoven, today's students are more likely to emulate a graffiti artist or a DJ. Increasingly, educators must acknowledge that this represents a legitimate cultural shift that offers opportunities for critical consideration and a rich resource for artistic creation.

I, among others, have argued elsewhere that learning by doing is by far one of the most effective methods for teaching the use of computer technology in fine art and design. The way to learn to draw with a piece of charcoal is to do it. Likewise, you learn how to create and edit NURBS through the experience of doing it. There is a tendency to confuse art made with computer with a kind of objectivity that belongs to science and engineering. In order to make art with a computer, you must conform to the dictates of the interface, which requires a deliberate and planned step-by-step approach. This is a rational process that yields to scripting and recipes. Yet this objective knowledge is no substitute for the subjective know-how and understanding that only comes from repeatedly using the menu-driven tools to achieve different aesthetic ends. In order to get over the hurdle of learning the interface, there must be the desire and infatuation with both the technology and its aesthetic outcomes. The "kewl" factor as a motivator is not to be underestimated.

New programs, curriculum development, and the needed implementations brought about by the demands of clients, both students and industry, affect hiring practices and conditions of employment for mid-career and entry-level educators. Could a revolving door between institutions of higher education and industry benefit the educational process and enrich the industry as well?

This hands-on workshop guides educators and students through the process of setting up an animatic and teaches them how to create a short story using traditional and computer-aided methods. The presentation includes:

- Story – how to structure around a given theme or music.
- Storyboarding – written as well as drawn visualization.
- Key framing – posing out the characters and scenes.
- Timing – guessing at what it should be (theory vs actual).
- Premeire – using the software to marry the visuals with sound; putting “theory vs actual” to work.
- Scheduling – how do you estimate the time it takes to do the work?
- Budgeting – if this were a real job what would it cost?
- The business – how students should think about marketing themselves.

The workshop also teaches the teachers what they need to set up a similar situation for their own teaching purposes. It includes samples of completed works and works in progress, to illustrate the process more completely.

Animatics

ThinkQuest: Students and Teachers Exploring a Global Web-Based Education Project

ThinkQuest is an educational initiative committed to advancing learning through the use of computer and networking technology. ThinkQuest challenges teachers and students of all ages to use the Internet in innovative and exciting ways as a collaborative, interactive teaching and learning tool.

Because of the Internet, groups of people from diverse locations, backgrounds, experiences, and nationalities are able to work and communicate in meaningful ways. ThinkQuest capitalizes on this new form of communication by inviting learners to explore topics they choose, collaborate to achieve a goal, and add meaningful content to the Internet.

An exemplary model of student-directed, project-based learning, ThinkQuest promotes an "Internet style of learning" – an interactive, participatory method that encourages students to learn by doing and take advantage of the Internet as a constantly growing source of information.

ThinkQuest teams have created a broad and valuable collection of Web-based educational resources for use by others around the world. In this way, ThinkQuest allows students to move from passive information consumers to active knowledge producers whose work is used and valued by millions. The initiative is comprised of several programs.

ThinkQuest Internet Challenge

The ThinkQuest Internet Challenge is an international program for students age 12-19 that encourages them to use the Internet to create Web-based educational tools and materials. Students form teams with colleagues from around the block or around the world and are mentored by teachers or other adult coaches. Competing for scholarships and awards totaling more than \$1 million, student participants learn collaboration, project management, leadership, and critical thinking skills that help raise their level of academic and technological prowess. Teams submit their entries in one of five categories: Arts & Literature, Science & Mathematics, Social Sciences, Sports & Health and Interdisciplinary. Finalist team members and coaches from each category are invited to the ThinkQuest Awards Weekend. Participation has grown at more than 40 percent per year, with over 10,000 students and coaches from 64 countries participating in 1999.

ThinkQuest Junior

ThinkQuest Junior is a classroom-based competition that encourages girls and boys in grades 4-6 to take a meaningful interest in computers and technology. ThinkQuest Junior teams create educational Web sites on a variety of subjects that make learning fun and contagious for other students of the same age. More than \$250,000 in cash, computers, and networking resources are awarded to winning students, teachers, and their schools. Participation in ThinkQuest Junior has more than doubled in its second year, with more than 1,000 teams participating from across the United States.

ThinkQuest Educational Technology Conference

This three-day conference for teachers and educational leaders is being held in Los Angeles, 10-13 November 1999, in conjunction with the ThinkQuest Awards Weekend. At the conference, noted technology and educational leaders share their visions and knowledge of technology and its implications for learning in the next century. Through presentations, panels, and small group discussions with conference attendees, issues critical to teachers and educational leaders, such as the future of the Internet, integrating technology into the classroom, and emerging Internet applications for education are explored in depth. For more information, check the ThinkQuest Web site: www.thinkquest.org

The University of Missouri-Columbia's Advanced Technology Center presents Virtual Harlem, an innovative use of virtual reality in an American literature setting. This virtual world allows students to visualize the setting and context of fictional texts in a computer-generated environment. Students are able to navigate streets, interact with historical characters, and participate in the virtual world's design. Approximately 10 square blocks of 1920s Harlem (the Harlem Renaissance) have been reconstructed to give students an unprecedented view of the cultural wealth and history of this area. It took over six months to gather material, maps, photos, and films to begin "building" Virtual Harlem. The environment was constructed to specifications detailing the exact lengths and widths of streets and placement of buildings. Currently, students are expanding this base and adding extra dimensions to the project.

According to Thomas Nagel (1974), research in cognitive science has demonstrated that experiential elements offered by various environments enhance acquisition of knowledge. Expanding environmental effects on learning, William Winn (1993) proposed that a different type of learning occurs in a virtual environment, called "constructivism," (which was actually an expansion of earlier constructivist theory proposed in the early 1980s by R.M. Gagne et. al.), in which students actively engage in the learning process. Rather than passively receiving knowledge, students must navigate and make decisions based upon various options presented to them. This active decision-making gives students the feeling of not only participating in a real-world environment, but also turns learning into exploration. Because they feel a sense of agency, students are more likely to engage in learning. The cognitive research of Bartlett, Neisser, and others has shown that, because of the experiential effect of VR and the interaction that users have with other users and objects within the environment, meaning within a virtual environment is very closely tied to that gained from the real world.


When students read novels, they are experiencing, second-hand, the experiences of the author. Although this second-hand account of an author's knowledge of a place or event may lead to a good understanding of what the author is trying to express, VR affords the student an opportunity to experience the same environment first-hand. Students may discover entirely new or additional meanings after "seeing" the scene themselves. In this way, VR becomes an intertextual engagement, allowing students to engage in a simulated environment, filled with music, photos, and dialogue. This intertextuality allows students to have a sense of engaging with the characters or author of a text rather than simply receiving information second hand. Students can then compare their experiences to those of the author, shaping their own reading of the environment or text, as well as better understanding the basis for the author's interpretations. VR promises not only to open the text to new possibilities, but also to revolutionize the notion of what it means to "read" or experience a text. Because of recent currents in literary criticism focusing on intertextuality, pastiche, and reader response, this aspect of VR should be vital to the future of literary criticism.

One of the primary reasons for using virtual reality in this course is to engage students on a technological and critical level. By entering a realistic rendition of the environment that inspired a story, a work of fiction can be better understood and more critically evaluated. Opening a novel spatially in this fashion helps students renegotiate their understanding of how fiction should be read. In Virtual Harlem's virtual world of sights and sounds, period photos were used for the reconstruction, as well as narrative and music from the 1920s. Participating in the creation of this literary environment encourages students to become more engaged in the class as a whole. Quantitative and qualitative assessments support this claim.

We have developed an advisory committee comprised of scholars from various disciplines, such as history, sociology, art, music, and psychology, all of whom see a relationship between Virtual Harlem and units they teach within their areas of specialization. The cross-disciplinary incorporation of VR technology in the humanities encourages students to explore relationships among art, history, economics, music, literature, and other fields that are usually taught separately.



Playground



A research sample is just forwarded into our Virtual Laboratory from the transport platform.



An analysis procedure: titration with a burette.

Virtual Science Laboratory is an interactive kiosk program with 2D character animation, 3D computer graphics, and actual tools and samples, imaginatively combined for the purpose of letting ordinary people, especially junior high and high school pupils, experience how scientists analyze environmental pollution, the hygienic condition of various food products, and other matters related to public health. Users are presented with a variety of actual samples or their mock-ups (for example, a piece of food, river water in a flask, etc.) and asked to analyze one of them in a scientifically appropriate manner. When they pick up the sample or mock-up and place it in front of our computer set-up, they are taken into the interior of the virtual (3D CGI) laboratory and greeted by a delightful pair of 2D animated characters: a seasoned scientist named Dr. Hygiene and a little girl, Ms. Ecolo. The participants' samples are represented in 3D CGI, and Dr. Hygiene shows them how to work on their samples in an interactive fashion. Dr. Hygiene is a bit absent-minded, and his instructions might sound a little confusing now and then, but when that happens Ms. Ecolo comes in to help the visitors.

In order to analyze their samples, visitors must interact with a variety of experimental tools in multimedia representations. They can see live-action video footage of their operation in a first-person point of view if they handle the equipment correctly. The result of the analysis differs depending on how visitors conduct it, so they feel as if they are real research scientists.

For difficult technical terms and measures, visitors refer to the Terminology Dictionary that is linked in the computer network and further linked to a variety of Internet sites related to the laboratory's subject-matter, so that visitors have access to more detailed information. Sometimes, a quiz is presented to visitors so they can confirm the results.

This program was originally developed for the Gunma Prefectural Institute of Public Health and Environmental Science in Japan, where scientists are actually involved in analyzing air, water, food products, and pathogenic microbes. In an exhibit hall at the Institute, this kiosk informs the general public about the institute's activities and purposes. The program is intended to make those visitors, especially junior high and older pupils, more conscious of the environment and hygienic conditions in their communities and help them to lead a healthier life.



Dr. Hygiene and Ms. Ecolo in the Laboratory; the animated characters who guide our visitors.

Classroom of the Future

This Workshop covers computer animation for the artist with disabilities who has a desire to enter the computer animation effects field. As new technologies emerge for the disabled computer user, the animation software being utilized by some of the major animation companies becomes more accessible for wheelchair-bound animators. The workshop stresses customization of the computer interface to the user, techniques for creating visual effects for various media, and computer animation as a viable career opportunity for the differently abled computer user.

Topics include:

- Unique computer interfaces, such as voice command, and alternative input devices.
- Pen-based interfaces that allow for more intuitive programs.
- Animation basics within popular animation systems and interfaces appropriate to creating animation within a production environment.

With the advent of Web-based informational content, many of the techniques and instructional resources for visual effects can be accessed through Web interaction and research. The workshop also summarizes these online resources, system integration, and research techniques.

Visual

Walking the Tightrope: Balancing Digital and Traditional Skills in Undergraduate Education

Panelists
Jeremy Butler

Kate Francek
Kathy Griswold
Jeffrey Lerer
John McIntosh
School of Visual Arts

Joel Sevilla

Ground-zero: there is a collision between the visual arts and emerging digital technologies:

It's not "On The Road." It's "In The Studio" where streams of consciousness are running wild. Small animation teams hidden away in sprawling lofts. Clandestine jungles of endless cables, technology, and meteoric talent. Here, coffee is considered a vitamin supplement.

OR

Independent developers, outside providers of content sitting in nearly hermetic isolation. Quietly, they shuttle between home and studio. Much of their content transmitted via modem, they may never meet their clients face to face. Defying the 24-hour day, they have usually surpassed the 40-hour week by Wednesday morning.

OR

Game developers? We won't even talk about THEM. Forget about it! Lunatics!

OR

Feature film production work. Four-year target strategies, two-year projects. Hundreds of artists pushing the stone and today only a few full-time contracts in the whole house.

OR

Independence! Experimental artists pursuing the festival circuit and grants. Perpetually plagued with no money, driven by an internal vision, sneaking around the periphery of a multi-billion dollar industry.

These are a few examples of how digital technologies are redefining art and the role of the artist. A hyper-evolution of tools and concepts is underway, and academia must evolve to meet the challenge and adjust to the needs of emerging digital artists.

The concepts of art training have been evolving for millennia. There is a continuity. Traditions handed down from one generation are broken by the next generation, and then juggled and revived by subsequent generations. Styles change, materials change, but the academic foundations that are most helpful to artists, in their formative years, remain unscathed. Identifying the critical foundations and insuring their survival in the digital environment are vital and essential.

With the advent of digital technology, most visual artists found the digital medium too expensive, too hard to maintain, too complex, too slow, too visually limiting. Computers were anything but inspiring to most visual artists. Complicating the matter was the limited amount of training that was available, and most of it was geared toward basic computer competency and navigating software interfaces.

As the computer has evolved into an affordable, user-friendly, outrageously powerful artist's tool, the focus on software operation in the early years of computer education has endured today, at the cost of our students' potential and the medium as a whole.

Traditional

First, Art; Second, Computer Art; Third, Industry Needs

In the ever-increasing volume of academic programs based in digital media, a closer connection with the tradition of art school training needs to be rediscovered and re-established. It is not sufficient for a school to provide a software training site or even a level of proficiency for students to enter any of the commercial industries mentioned earlier. Our mission is to expose students to the values and concepts of art in addition to the mechanical mastery of computers and software.

Have educators lost sight of their mission by focusing on vocational training in a response to the false challenge of an entertainment industry (promising to hire thousands of computer animators) to train artists as fodder for their now often empty production cubicles?

New curricula in the digital arts will only be vital if all the major concentrations of a thorough art education are included and stressed. Certainly, a higher degree of critical rigor must be imposed on digital work. In recent years, the function of effective art criticism has been scaled down to make allowances for the fact that it was digitally generated, the aesthetic standard fuzzy, and lowered for the mostly non-artists driving the medium.

Traditionally, one of the most valuable aspects of art school is the exposure to working artists. Students begin to understand first hand what a career in art means and what an artist's identity implies by their contact with the faculty. Today, too many instructors are in fact actually not practicing artists, and they do not have the expertise necessary to be training in an academically rigorous environment. Talented, experienced digital artists are in great demand. Most schools do not have the ability, flexible schedules, or funds to attract the most experienced artists, even as adjunct faculty.

These, and more, are the challenges facing technology-driven programs in undergraduate education. But these are only the philosophical or administrative concerns that will haunt the educational process for years. In the meantime, we still are trying to provide the best education possible to our students. We point, where we can, to the skills that will serve our students best and over the longest time.

One truly visual skill that artists share is their ability to translate or interpret what they observe into a work of art. In animation, observation is paramount. The first phase of the translation from the world to the computer most often lies in a drawing or a sketch. Capturing texture, tone, gesture, expression on paper to be used at first to remember or explore, then as a character sketch to reference while building a 3D model and animating it on a computer.

The creative process of moving from observation to drawing to computer animation can take as many unique turns as there are artists who attempt it. Depending on the skills of the artist, any single element of this process can be critical. Imagination vs. observation, quick sketch vs. detailed drawing, the decisions and the style of a mature artist are the important discoveries for a talented student. It is the student's perspective we wish to explore in this panel. It is their process of discovery and knowledge that is our concern and our reward. Here are some comments from two graduating seniors from the School of the Visual Arts.



Walking the Tightrope: Balancing Digital and Traditional Skills in Undergraduate Education

Observation in Computer Animation

Jeremy Butler

"The importance of observing a motion, or performance, for animation purposes is essential for capturing true life in motion. Almost all animation, whether cartoon or realistic, relates to real-world physics and the laws of motion. The process of computer animation is not as simple as making a 3D character move. The true art and goal of animation is to bring the character to life. The animator wants the viewer to know, without question, that the character on the screen is moving on its own with its own motivation and purpose. Part of that believability comes from observation.

"Observing from life gives clues to the subject's manner, it's habits, and most importantly the distinctive motion that defines it. These clues are vital to the believability of the character no matter how stylized the final output is. Observing from life is ideal, since it's easier to see how the form relates to itself in three dimensions. Video reference is valuable because the actions can be repeated infinitely for longer studies.

"The notes from my studies are generally a mix of words and sketches, with the notes being similar to captions for the drawings. Besides the notes and sketches, the process I go through is aided by acting the motions out myself in a mirror. Acting out the motion is probably the best way to 'feel' it. This brings me to a better understanding of the motion, and from there I form the basis of my animation."

Animation through Observation

Joel Sevilla

"Because we all move, the ability to animate should be relatively easy. If you can get yourself into a position or pose, then putting your character into the same position is a piece of cake, right? Wrong! It is the understanding of how you got into that creation position that makes animation fluid and believable. Weight shifting, hierarchy of movement, and bone structure are important elements of animation.

"The best way to implement these elements and to understand them starts with drawing from life. Being able to put your ideas for an animation down on paper before ever touching the computer is infinitely important. It not only helps in understanding what poses and actions your character must perform, but it also allows for the building of personality for the character. Would the character be able to do backflips, or is the character someone who likes to spend hours in front of the tube with a bag of chips? These are questions that are answered in the development stage through writing and drawing.

"Observing emotions and feelings is also very important in the animation process. Showing a character go through a cycle of emotion will make that character believable. What does it look like when someone is worried, happy, or mad? These feelings are not only conveyed in one's face but in their body language as well. This coincides with the performance that the character must give. The animator becomes an actor who is giving a performance through the character. Like an actor, the animator takes his or her observations, feelings, and experiences in life and uses them in the performance."

The Panel

This panel features aspiring animators Joel Sevilla and Jeremy Butler, independent animator and faculty member Jeffrey Lerer, and the chair of the Computer Art Department of the School of the Visual Arts, John McIntosh. The panel presents a no-holds-barred conversation and debate on the experience and challenges of balancing digital techniques in a fine art curricula. Illustrations from the faculty and students demonstrate the creative advantages of applying traditional skills in developing the most elaborate computer animations.



Web Pages, Interactive Interfaces, and Worm Holes: The Next Generation of User Interface Designers

MaryEllen Coleman
IBM Corporation
mea@us.ibm.com

Classroom

IBM wanted to work with elementary school students to see how the next generation of computer users viewed state-of-the-art technology, Internet communications, and software user interface design. We hoped that this insight would help us improve our products for these future customers. As information developers, we are using the Web more frequently to deliver technical information to our customers. We also work on teams in order to complete all our projects. We decided we could learn how children view the Web by teaching them how to put together a Web site of their own.

Presenters
Carol Bahruth
MaryEllen Coleman
IBM Corporation

Through previous contacts, the IBM team leaders approached the principals at Regina Coeli and North Park Elementary Schools and asked them if they would be interested in working on this collaborative project. They accepted the offer and selected students and faculty advisors to join the project. Each school team consisted of eight accelerated students, boys and girls in the fifth and sixth grades, school librarians, and technology teachers.

IBM kicked off the project by meeting with all 16 students and advisors at Regina Coeli. We talked about the project – what we expected to create and what we hoped to learn along the way. After letting the students gather ideas, members of the IBM team met with them and their advisors at their schools. The students called out their ideas while the IBMers wrote them down on the board in a tree-diagram. Through this lesson, the students learned how to organize their ideas into a hierarchy of topics that can be navigated on the Web. They also learned how to critically evaluate an idea without shooting down the person who suggested it and how to build on others' suggestions. After the meeting, the students met on their own (with their advisors) to gather pictures and text to put on their Web sites. They involved their schoolmates who were not assigned to the project by asking them to write essays about their school, their classes, their projects, and their special days.

When the students had their content ready to assemble into a Web site, they came to IBM for two all-day work sessions. At IBM, they got their hands on the same equipment that the human-factors team uses to evaluate interface usability; that the graphic-design team uses to draw, scan, and enhance images; and that the multimedia team uses to record and manipulate sound, create animation, and code HTML. The IBM team helped guide the students in all aspects of the project, including group dynamics, Web page design, and the technology and tools used for creating Web pages. Most of all, we stressed teamwork. The students had 12 hours of lab time at IBM, but they also spent many lunch hours and free periods at their schools working on their Web sites.

The students also learned leadership and presentation skills when they selected two of their teammates to present their Web sites at IBM's "Interact'98 Ease of Use" conference in Yorktown Heights, New York, an annual gathering of usability ease-of-use professionals from IBM sites worldwide. Attendees were fascinated with the students' accomplishments and their enthusiasm for working with IBM on user interface design issues. Here, the students admitted some of the challenges they faced while working in a team and described how easy it was to assemble a Web site once they organized their content and learned how to use the tools.

Throughout the entire project, we videotaped our meetings and lab sessions so that we would have a record of our thoughts and experiences along the way. In the tape, one can watch the students move from being shy individuals to valuable contributors to the team effort. At IBM, projects are rarely completed by individuals. We were happy to give students the opportunity, time, and space to work on a truly group project, while also giving them a glimpse into what human-factors specialists, graphic designers, and multimedia programmers do.

IBM continues to work with these students and is looking for additional opportunities to work with other schools and communities to develop collaborative projects like this one. We hope that our experience of a successful collaborative effort between industry and academia will inspire educators and industry to work together to learn from each other.

Designers

Classroom
Playground
Workshop

In this student animation project, children tell stories, draw, and work with modeling compound to create characters and scenery, then produce their stories on the computer.

Children



Why is the Mona Lisa Smiling? (library.advanced.org/13681/data/davin2.shtml)

Steve Feld
John F. Kennedy High School
sjfeld@erols.com

Classroom
Playground
Workshop

My experience as a ThinkQuest coach has been extraordinary. I am a teacher at John F. Kennedy High School in Bronx, New York. In December 1996, I was asked by Gino Silvestri, the principal of the school, to attend a classroom session by Dr. Sheila Gersh at City College to learn about the ThinkQuest Project. ThinkQuest is an annual international contest that brings together distant schools to partner on Internet projects. At the time, I had little Internet experience, but a great desire to get my students involved in the international collaborative project. Because our computer classroom had rudimentary equipment, I knew we would face compelling challenges to succeed in this process.

Seven years earlier, my students competed in the Learning Technologies Fair in Albany, New York. They won the competition with *Da Vinci's Visions*, a student-created computer research project, which included a video component and a 24-screen computer game. At the fair, we met Dr. Lillian Schwartz, who was in the audience. She introduced herself and explained that she had researched Leonardo da Vinci and had made some remarkable discoveries. It is her theory that da Vinci painted himself as Mona Lisa. Although Dr. Schwartz's book, *The Computer Artist's Handbook*, was out of print, my students selected the theme *Why Is the Mona Lisa Smiling?* as their project for the ThinkQuest contest to revitalize her theory.



We needed to collaborate with another school, and first we used the ThinkQuest Team Finder (where participants can post a message on the ThinkQuest server) to seek partners for our project. Although we located a Canadian student to help us with team formation, she dropped out of the project just as the deadline approached. We ultimately found partners from Borlange, Sweden, using a Canadian listserv called INCLASS. We knew we had to work quickly because the deadline was only three months away.

The key component was to make the Web site accessible to all. Through their research, my students discovered that da Vinci wrote music. They found the music and digitized the score to make it Internet compatible. We also wanted to include an online interactive quiz, but we wanted it to be viewable with any browser. That meant that we needed to learn how to incorporate Java Applets and CGI scripts into our project. We secured permission to use copyrighted images in our project. My students coded the project using text processors, while our Swedish collaborators researched the Web to find the best da Vinci resources available. We communicated with each other through email.

Because we wanted our project to be accessible to schoolchildren, we included elements that would be fun to use. We added random research and digital postcard elements. We also included a guest book and site survey to encourage feedback from our visitors. We responded to those who provided contributions from the field.

Mona Lisa

Why is the Mona Lisa Smiling? (library.advanced.org/13681/data/davin2.shtml)

The response to our project has been phenomenal. On August 23, 1997, a visitor from Mongolia signed our guest book. Since then, our guest book has been signed each day for over eight months, a continuous chain of support from educators, art historians, and schoolchildren from more than 48 states and nearly 60 countries. On September 2, 1998, our site became a USA Today "hot site." On the same day, it was selected as a "cool site." As a result, our site received 2,000 visitors in one day! Our growth has been exponential. In October, Gloria Edwards used our site in a classroom integration session at Purdue University. In November, the newsletter Classroom Connect selected our project for inclusion in its Web guide resource.



In December, we presented our site at City College, and we were featured on the front page of two local newspapers. As a result, we met art historian and author Rina De Firenze who wrote *The Mystery of the Mona Lisa*. We collaborated with her on the subject of the painting and have added scientific inquiry to our project by comparing her theory with that of Dr. Schwartz.

In January 1999, our site was featured on Radio Net's "The Human Factor," a live broadcast that is also Web-based, and *Seeker Magazine* gave our site their "site of the month" distinction. The Getty Museum has recently placed our site in their new Digital Experience. We were also featured as a "hot site" for The School Page UK, and we formed a partnership with *MidLink Magazine*.

In March, we presented the project at the Make It Work Conference, sponsored by AT&T and the New York City Board of Education. We were also entered in Teachers' Choice at The Well Connected Educator, and the encyclopedia Microsoft Encarta included our lesson as part of their collection.

In April, we were invited to present our project at the School Tech Expo at the New York Hilton. As part of the Web publishing workshop by Caroline McCullen, editor of *MidLink Magazine*, we shared our project with school administrators and teachers. At the School Tech Expo workshop by Robert Sibley, the educational director of ThinkQuest, we expanded the interest in our project to potential ThinkQuest participants. My students also presented their project as part of the Virtual Classroom Experience at the expo.

As our project evolves, we continue to expand our partnerships and outreach. We were the *Computer Currents Magazine* featured link of the Week on May 11, 1998 and are now part of the permanent review section on their site. We were also a featured "site of the month" on Prodigy. We anticipate an article in the fall edition of *Technology and Learning*.



As the ThinkQuest head coach, it was my responsibility to keep the ambitious project on schedule and to assist with promotion of the site. The experience has brought exposure to our school, and we will be getting four new computer labs and a T1 line at John F. Kennedy in the fall.

The students who worked on the project experienced the power of the Internet for collecting links that are pertinent to a given problem construct. As our guest book grew, the students noted how the use of that Internet-supported vehicle could serve as a multigenerational, multisector, collaborative forum. They weighed conflicting evidence, perspectives, and issues evoked by the comments in the guest book and reflected on the differing theories showcased in the project.

The students also learned how a Web site not only can be a forum for problem construct inquiry and evaluation, but also can serve as a vehicle for project dissemination. The students' responsibility to code their site so that it would be accessible to all led to their familiarity with and fluency in HTML, Java Applets, and CGI scripting. As researchers, investigators, and Web builders, they also learned the importance of scheduling, staying on track, and daily monitoring of site functions. The students became sensitive and responsive to the evolution of the project in response to the feedback from our participants through our site survey.

Steve Feld

Steve Feld is a veteran fine arts instructor who has been infusing computer graphics into the curriculum for over 25 years. He has received numerous awards, including: an Impact II Developer Grant, a Computer Learning Foundation Grand Prize, and the New York State Learning Technologies Championship for his student-created project *Da Vinci's Visions*. He has helped the International ThinkQuest team create the award-winning Web site *Learning About Leonardo* (library.advanced.org/13681). His projects have also been funded by the William T. Grant Foundation and the Bronx Superintendentcy. He is the author of *Computers in the Art Classroom* published by the New York City Board of Education.




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

Panels



Contents

- 116 Introduction
- 117 Committee
- 118 Hot Topics in Graphics Hardware
- 119 How to Cheat and Get Away With It: What Computer Graphics Can Learn from Perceptual Psychology
- 122 CG Crowds: The Emergence of the Digital Extra
- 124 3D Tracking in FX Production: Blurring the Lines Between the Virtual and the Real
- 127 Natural and Invisible Human Interfaces
- 130 Research and Development for Film Production
- 133 Visualizing Large-Scale Datasets: Challenges and Opportunities
- 136 Scene Graph APIs: Wired or Tired?
- 139 Get Real! Global Illumination for Film, Broadcast, and Game Production
- 142 Experiential Computer Art
- 144 Visual Effects: Incredible Effects vs. Credible Science
- 147 How SIGGRAPH Research is Utilized in Games
- 149 Visual Storytelling
- 150 Function and Form of Visual Effects in Animated Films
- 153 Digital Watermarking: What Will it Do for Me? And What it Won't!
- 156 Mixed Reality: Where Real and Virtual Worlds Meet



Special Sessions

- 160 Star Wars Episode 1: The Phantom Menace
- 160 Web 3D RoundUP
- 161 A Visit With an Animation Legend
- 161 Fiction 2000: Technology, Tradition, and the Essence of Story



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


Panels

116

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.



Conference Abstracts and Applications



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



Hot Topics in Graphics 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.

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



The virtual reality lab at the University of Tübingen 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 Tübingen" 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

Psychology

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.

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

CG Crowds: The Emergence of the Digital Extra

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.

Digital

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 inter-mixed 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.

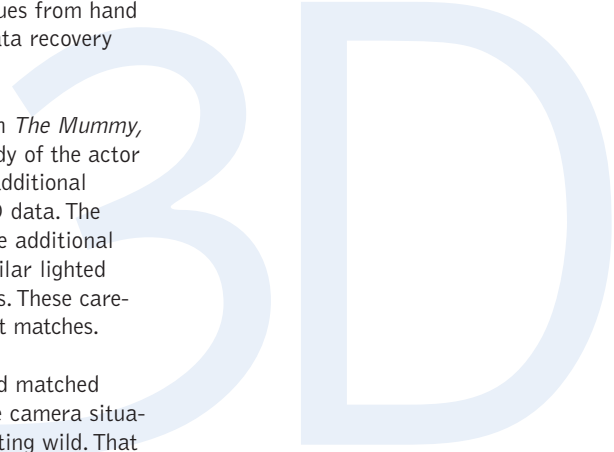
Recent projects at ILM have made use of all these techniques. On the film *The Mummy*, we provided CG prosthetics to replace portions of the face and upper body of the actor playing the title role. On set, he was fitted with LED-laden objects plus additional marks on his skin. In post production, these 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



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



Steps in Geometrix's fully-automated SoftScene™ process (top to bottom)

- 1) Automated feature extractions
- 2) Automated feature tracking
- 3) Camera path extraction
- 4) Scene geometry and texture extraction

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.

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

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, air-flow, and water movement. Here, we seek to communicate digitally mediated senses of activity and presence at the periphery of human awareness.

Panelists
Hiroshi Ishii
Massachusetts Institute of Technology

Caleb Chung
Giving Toys, Inc.

Clark Dodsworth
Digital Illusion/Osage Associates

Bill Buxton
Alias|Wavefront, Inc.

interfaces

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.

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 prod-



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

ucts 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.

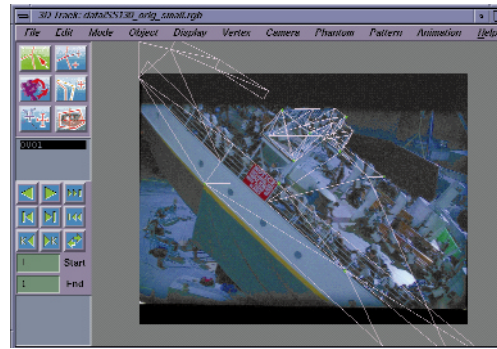
Ken Pearce

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.

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?

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.

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.



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.

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.

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.

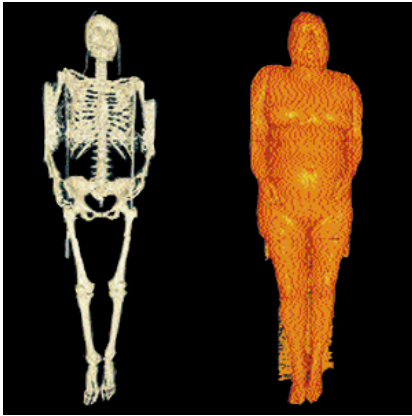
Using a full system design approach, visualization requirements are compared with technology trends to quantify the visualization system requirements. Researchers are exploring data reduction and selection, parallel data streaming, and run-time visualization techniques. The system performance goals require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


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.



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



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.

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.

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?



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 aneurysm on the right side of the image. Image provided by Christopher Johnson, University of Utah.

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 SGI's 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-

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

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.

Panelists
Craig Barron
Matte World Digital

Scott Lelieur
StudioDVP/Design Visualization Partners, Inc.

George Murphy
Industrial Light & Magic

Dave Walvoord
Blue Sky | VIFX



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.

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.

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.



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



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

Art

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.


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.



ClickScape 98, Thomas Lehner, Stadtwerkstatt



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.

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.

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.



George Suhayda



Rob Minkoff



Stuart Sumida



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

Moderator
Chris Hecker
definition six inc.

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.

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

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 indispensable 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.

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.



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

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.

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 unreality 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 photorealistic. 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 "unenanced" 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.

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.

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

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 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 (AR² 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

Organizers

Don Brutzman, Naval Postgraduate School
Timothy Childs, Oz

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.

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

Star Wars

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

Where digital and human
guests reconsider, rework, and
reinterpret everyday life in
alternative future formats.

Emerging Technologies: The Millennium Motel



Kathryn Saunders
Royal Ontario Museum




Contents

164 Introduction
165 Committee/Jury
166 Surface Drawing
167 The Luminous Room:
Some of It, Anyway
168 HandSCAPE
169 (void*): A Cast of Characters
170 Life Species
171 LEON
172 R O U T E 6 6
173 Curlybot
174 musicBottles
175 Ensphered Vision
176 Touchable 3D Display
177 Hologram/Head-Mounted Display
178 Shared Space; Collaborative
Augmented Reality
179 Head-Mounted Projector
180 Visual Conductor
181 Microworlds, Sirens, and Argonauts
182 Emergence
183 Water Display
184 VisiPhone
185 metaField Maze
186 HyperMask: Virtual Reactive Faces
for Storytelling
187 TV Guides
188 Digital Cloning System
189 City of News



Introduction

Kathryn Saunders
SIGGRAPH 99 Emerging Technologies Chair



Welcome to the Millennium Motel, where digital and human guests reconsider, rework, and reinterpret everyday life in alternative future formats.

In this parallel universe, only a quantum leap away from "old" routines, data spawn digital life forms that exchange genetic codes and reproduce in evolving mutations. Guests create 3D surfaces by moving their hands through space, interact with intelligent objects to assemble complex models, and conduct a virtual orchestra. They gather at The Pool to share stories with interactive digital images of their friends, enjoy a smart drink, eavesdrop on distant conversations, and observe Shooting Stars before they check out 21st-century NightLife.


The Pool

The occupants of the Millennium Motel include some of the world's most interesting researchers, whose dreams generate surprising departures in music, buildings, sculpture, and storytelling. Although their areas of interest range from ubiquitous computing to augmented reality, immersive displays, and intelligent characters, they all share a deep commitment to creating meaningful experiences and communicating their visions of future realities.

Visions

We all owe these explorers a debt of gratitude for their imagination, talent, dedication, and willingness to share their achievements in SIGGRAPH 99's Millennium Motel

Explorers



Emerging Technologies Committee

Committee

Chair

Kathryn Saunders
Royal Ontario Museum

Subcommittee

Mark Davies

Kirsten Douglas

Randy Dreager

Tara England

Drew Gauley

Rick Hopkins

CiCi Koenig

Richard May

J. Marshall Pittman

Preston Smith

Kevin Sugden

Special thanks to: Carolina Cruz-Neira

Jury Members

Bruce Blumberg
Massachusetts Institute of Technology


Jeff Close
ThinkOne, Inc.

Clark Dodsworth
Osage Associates

David Ebert
University of Maryland - Baltimore County

Linda Jacobson
Silicon Graphics

Ken Perlin
New York University




Collaborators
Cici Koenig
Peter Schroder


Drawing

Surface Drawing is a medium that enables creation of a wide variety of intricate, organic 3D shapes. Objects are created by moving the hand through space. The path of the hand forms surface pieces that seamlessly merge when they touch. Users can also erase, add details, and manipulate objects with a simple two-handed interface. Freely growing, joining, and erasing surface pieces enables rapid prototyping of freeform shapes.


This extension of traditional line drawing to 3D space allows creation of shapes without the creative constraints of a rigid mathematical structure, a large toolset, or a tedious construction process. The system is implemented with the semi-immersive environment of the Responsive Workbench. An 18-sensor CyberGlove measures hand configuration, which controls the shape of surface pieces.



A guitar is drawn in five steps by tracing its shape with the hand.



The Caltech Responsive Workbench.



A Surface Drawing in progress.



When the CRT breaks open and the pixels inside leak out to stain everything: one of the results can be a *Luminous Room*.

When graphical display is not only free to occur on any surface in the room but can also react to what's happening at those surfaces and within the space, certain kinds of usefulness may ensue.

When the behavior of these environmental pixels accretes especially around physical objects that act, to localize meaning and focus the expression of participants' intent, an interaction style called *luminous-tangible* prevails.

When one conclusion of all this is that a large class of spatially oriented design and experimentation activities can be served by luminous-tangible techniques, real architectural models begin to cast accurate shadows and hydrodynamically divert simulated airflow; cheap little models of lasers and mirrors and lenses begin to emit and bounce and spread visibly simulated beams of light.

Collaborators
Daniel Chak
Gustavo S. Santos
Jessica Laszlo
Hiroshi Ishii



Luminous


HandSCAPE

Collaborators

Jay Lee
Victor Su
Sandia Ren
Hiroshi Ishii
James Hsiao
Rujira Hongladaromp

HandSCAPE is an orientation-aware digital measuring tape. While a traditional measuring tape only measures linear distance, the addition of orientation sensors allows a vector measurement of both length and direction, and the tape can serve as an input device to computer drawing and modeling applications.

HandSCAPE provides a simple interface for generating digital models of physical objects. The interaction involves taking measurements of several physical objects and the distances between them. Once the model has been generated, the user can manipulate it in the digital domain. HandSCAPE preserves reliance on human senses and skills by referring to the familiar process of measuring objects and spaces.



(void*) is a novel gathering place that unites the physical and the digital, allowing people and a cast of directable and autonomous characters to interact with each other. The interactions focus on movement, groups, and body language. Building on Swamped! (SIGGRAPH 98), this installation continues the MIT Media Labs' work in intentional characters, sympathetic interfaces, and autonomous cinematography, and adds "dynamic music composition" as an integral part of the system. It also introduces a number of novel physical interfaces to allow simultaneous multiple-participant interaction.


Collaborators

Members of the Synthetic Characters Group
and the Responsive Environments Group,
The Media Lab, Massachusetts Institute
of Technology:

Bill Tomlinson
Michael Patrick Johnson
Song-Yee Yoon
Marc Downie
Ari Benbasat
Jed Wahl
Dan Stiehl
Delphine Nain
Joseph Paradiso



(void*)



Life Spacies

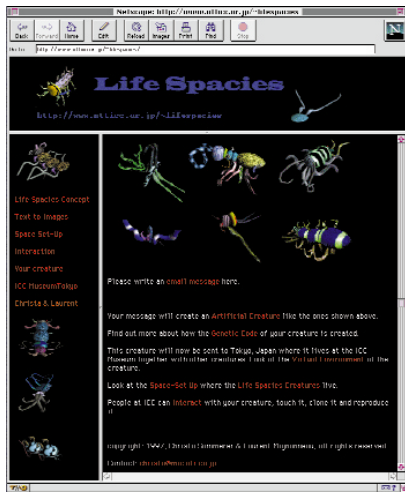
References

1. Life Spacies was produced for the ICC InterCommunication Museum, Tokyo, as part of its permanent collection.
2. C. Sommerer and L. Mignonneau, Life Spacies: a genetic text-to-form editor on the Internet. Proceedings AROB 4th'99, Beppu, Oita, pp. 73-77, 1999.

Life Spacies is an interaction and communication environment where remotely located visitors in a global environment (the Internet) and onsite visitors (in the local environment) interact with each other through artificial creatures.¹

By simply typing and sending an email message to the Life Spacies Web site (www.ntticc.or.jp/~lifespacies), visitors can create their own artificial creatures and receive their creatures' pictures and curriculum vitae. As soon as the site receives a creature-creation message, the new creature starts to live in the Life Spacies environment. Two independent interaction sites are linked together via a data line, allowing visitors at remote locations to be displayed in the same virtual 3D space.

Onsite visitors can directly interact with the creatures by touching and catching them. When a creature is caught by a visitor, it creates a perfect copy of itself. However, if two remotely located visitors are in the same virtual space, they can each catch a creature with their hands, which causes the creatures to mate and create an offspring by genetically exchanging their code. A special text-to-form editor enables translation of text into genetic code. The characters, syntax, and sequencing of the text are used to code specific parameters in the creature's design. Form, shape, color, texture, and the number of limbs are influenced by the text's parameters.²



LEON

Athomas Goldberg
IMPROV Technologies, Inc.
athomas@improv-tech.com

Lurking in the shadows of the Millennium Motel are some familiar faces (to anyone who attended the SIGGRAPH 98 Electronic Theater). Leon, last seen in Mitch Butler's "The Smell of Horror," and other mysterious characters have checked into the Motel, where they're checking out the guests who are checking out the fantastic array of emerging technologies.

"Leon in the Millennium Motel" is a joint project of IMPROV Technologies, Inc. and the Mitch Butler Company, and was created using the IMPROV Real-Time Animation System developed at New York University's Media Research Laboratory.

Executive Producers
Brian Blau and Athomas Goldberg

Artistic Director
Mitch Butler

Production support provided by students in the graduate animation program of New York University's Tisch School of the Arts.



Magritte's Cow is interested in the transformation of ideas. We look at the Web and see it as a place. We think about machines and envision them as bugs. We play with concepts and turn them into projects.

Magritte's Cow
Ronen Lasry
Daniel Szecket
Environments For People
info@magrittescow.com
www.magrittescow.com
+1.323.258.4371

When I was a little boy, my mother told me not to mix water with electricity. I have been doing it excitedly ever since.

I would like to thank Kathryn Saunders and Daniel Szecket, for giving me the opportunity to participate in SIGGRAPH 99, and my sister, Jennifer Holly, for perpetually working her mojo to inspire creativeness in our family. Also, thanks to Wendi-Mae Camara and Jerry Casilli for permission to use their Neon Fountain sculptures.

Neon Fountain, Liquid Light Sculpture Series, Cascading water, acrylic, liquid polymer, neon and electronics.

Portal to the Millennium Motel, Space Design, Spandex, lasers, intelligent light, high-voltage electronic devices

Dickinson Prentiss
Dreamlab
970 Lakewood Drive
Sunnyvale, California 94089 USA
+1.408.720.1450
Dreamlab@primenet.com

Jennifer Prentiss has been working in glass and light media for six years. She designs and constructs autonomous stained-glass panels, doors and entryways, lamps and lighting fixtures, mixed-media work, screens and room dividers, windows, wall pieces, 3D sculpture, and portraiture. She uses sandblasting, acid etching, mosaic, neon, copper foiling (Tiffany), and lead came techniques to create her art. Primarily, she is a commission artist specializing in architectural stained glass.

Her work at The Millennium Motel is stained-glass furniture. It was built with streaky and water glasses, using lead and zinc came.

Jennifer Prentiss
+1.323.258.4371
info@jenniferprentiss.com
www.jenniferprentiss.com

R O U T E 6 6


R O U T E 6 6 is a live, interactive, MIDI-driven, 3D world created for The Millennium Motel. It is the world outside the motel, an ever-changing environment that can be transformed by visitors using MIDI interfaces linked to SideEffects Houdini software and driven by Intergraph workstations.

The main display, located at the motel's entrance, is composed of five side-by-side projected screens that create a surreal panoramic desert landscape.

Some of the interactions include:


- A panoramic environment with a full 360 degrees of navigation
- Animated interactive elements such as UFOs, wildlife, etc.
- Changing landscapes
- Seasonal changes
- Eclipses
- Comets
- Weather changes (rain, snow, fog, lightning)
- Day/night
- Interactive trees (shedding and growing elements)
- Sound effects
- Overlay graphics

The R O U T E 6 6 landscape is interspersed with "windows" to video feeds. When visitors open the "windows," they reveal video loops, cameras pointing at other locations in the space (live video feeds of The Millennium Motel and other areas of SIGGRAPH 99), and typical motel TV. Fog machines create a 3D element in the foreground, allowing for 3D projected effects.



Recent trends of embedding digital technology in toys have led to greater possibilities for manipulation and interaction. Curlybot is a two-wheeled toy with embedded electronics that can record and play back motion. It remembers its change in position and replays its movements with all the intricacies of the original gesture. Every pause, and even the shaking in the user's hand, is recorded.

In this presentation, the interaction experience is augmented by projecting trails behind each toy to give participants a chance to compose a movement or a dance and create graphical patterns out of simple gestures.



Curlybot

Collaborators

Rich Fletcher
Jay Lee
Seungho Choo
Joanna Berzowska
Craig Wisneski
Charlie Cano
Andres Hernandez
Colin Bulthaup

MUSIC

musicBottles introduces a tangible interface that deploys bottles as containers and controls for digital information. The system consists of a specially designed table and three corked bottles that “contain” the sounds of the violin, the cello, and the piano in Édouard Lalo’s Piano Trio in C Minor, Op. 7. Custom-designed electromagnetic tags embedded in the bottles enable each one to be wirelessly identified. When a bottle is placed onto the stage area of the table and the cork is removed, the corresponding instrument becomes audible. A pattern of colored light is rear-projected onto the table’s translucent surface to reflect changes in pitch and volume. The interface allows users to structure the experience of the musical composition by physically manipulating the different sound tracks.




Opening a bottle to release the sound of the cello.(Photo by Joanna Berzowska)



The musicBottles table
(Photo by Seungho Choo)


Visual immersion plays an important role in virtual environments. Head-mounted displays (HMDs) provide a full solid-angle view of virtual spaces, but their optical systems limit their fields of view.

In this image display system, a large screen is used as an alternative to HMDs. The sphere is an ideal shape for a screen that encompasses the human visual field because it maintains a constant distance between the eyes and the screen as the viewer's head rotates. Ensphered Vision uses a single projector and a convex mirror to display seamless images. The optical system employs two mirrors: a plain mirror, which bends the light so that the viewer can see the image from the center of the spherical screen, and a spherical convex mirror, which diverges the light from the projector in the spherical screen. This optical configuration provides a seamless wide-angle image in a very limited space. The screen's field of view is 270 degrees horizontally and 100 degrees vertically. The image totally surrounds the viewer. The image viewing angle is much larger than a dome screen displaying a fish-eye-lens image.



Vision

Touchable 3D Display




Team
Koichi Oyama
Communications Research Laboratory and
Telecommunications Advancement Organization

Yoshiki Arakawa
Communications Research Laboratory

Makoto Sato
Tokyo Institute of Technology


In this reality-enhanced 3D display, Fresnel lenses are set so that parallax adds a sense of depth to a "real" image. Three-dimensional images appear very close to viewers (an effect that is difficult to achieve with conventional 3D displays), so viewers feel that their bodies are included in the 3D space. When a force-display SPIDAR is combined with the display system, users experience a reality-enhanced virtual environment with 3D images and force feedback. They interact with the images as if they were a part of the virtual environment.

Combination of Real Image and Parallax Presentation



Principle
A near object is presented with almost-normal parallax and focus, both of which are used to sense the depth of near objects. A far object is presented with proper parallax, which is mainly used to sense the depth of far objects.

Combination of 3D Display and SPIDAR




Appearance
Viewers can see, touch, and interact with 3D images in front of them.

This experimental see-through HMD (head-mounted display) uses holographic optical elements (HOE) instead of the half mirror that is usually used in conventional HMDs. Because it is grated, the system can produce images by diffraction, and it behaves like a heads-up display. It delivers the left and right images into both eyes, so users experience binocular parallax (stereoscopic) images.


Collaborators
Eiji Shimizu
Hideya Takahashi
Koji Yamasaki
Toshiaki Matsumoto

When the system is used in interactive augmented-reality environments, its laser-illuminated LCD display mixes virtual and real worlds. Users see clear laser-illuminated images floating within wide, bright real environments.



Prototype of HMD using holographic optical elements.

Hologram



View through a holographic optical element.

Collaborators
Hirokazu Kato
Richard May
Stefan Kraus

Shared Space merges real and virtual worlds in a way that can radically enhance face-to-face and remote collaboration. By wearing a lightweight, see-through head-mounted display, users see 3D virtual images overlaid on the real world and attached to real-world objects.

For face-to-face collaboration, this allows users to see each other at the same time as the virtual images between them, which supports natural communication between users and intuitive manipulation of virtual objects. For remote collaboration, Shared Space overlays life-sized live virtual video images of remote collaborators on the local real environment, supporting spatial cues and removing the need to be physically present at a desktop machine to conference. In both cases, computer vision techniques are used to precisely register virtual images with physical objects, extending the "tangible interface" metaphor.


At SIGGRAPH 99, Shared Space allows face-to-face and remote users to create interactive art together using virtual animated characters and props in a real tabletop environment.



The view through the head-mounted display in a face-to-face setting. Users can see their collaborators and virtual objects between them.



View from outside the interface. The virtual objects are only visible to those wearing the head-mounted displays.



A virtual video window of a remote collaborator.



A virtual video window of a remote collaborator and shared virtual image attached to a real-world object.


This project proposes a head-mounted projector (HMP) using the X'tal Vision (Crystal Vision) technology that was demonstrated in Enhanced Realities at SIGGRAPH 98. With X'tal Vision (a projector with a small iris and a retroreflective screen), users can observe stereoscopic images with an almost-correct occlusion relationship between the virtual and the real environment.

Collaborators
Naoki Kawakami
Dairoku Sekiguchi
Taro Maeda
Susumu Tachi

AT SIGGRAPH 99, the HMP is demonstrated in three applications.

1. Virtual images of a skeleton that make a patient's body appear to be transparent.
2. A paper-type display.
3. Optical camouflage suitable for visuo-haptic integrated display.

This work is supported by the Telecommunications Advancement Organization of Japan and JSPS's Research for the Future Program.



Projector

Visual Conductor

Collaborators
Senthil Kumar
Xiang Zhang
Joshua Gluckman
Eowyn Cenek
Bell Labs

Fred Bianchi
Richard Campbell
Bianchi & Smith LLC

A "live" conductor directs a complex electronic orchestra with natural expressions of hands and baton. No sensors or wires impede the conductor's movements, which are sensed with video cameras. The system detects beat events and gestures related to rhythm patterns and dynamics, and uses this information to control the tempo and volume. Participants in this interactive display conduct a large live-sounding orchestra. For ballet segments, the experience is enhanced by an animated dancer whose movements are synchronized to the music.

Dancing to Visual Conductor's beat.



Visual Conductor with Virtual Orchestra



Microworlds, Sirens, and Argonauts is a fantastic journey through multiscale microscopic worlds that grow and transform as users interact with them, revealing new patterns, structures, and sounds. It introduces the concept of “living narrative landscapes:” virtual spaces that allow users to successfully construct their own navigational maps and build their own representational models that can coexist with the narrative of the environments. Thus, the virtual space becomes a living narrative landscape as the Argonauts (users) navigate along time and space, taking part in the complex visual and aural behaviors of the environments.

The environments mirror reality by containing rich experiences at multiple scales. Sometimes, we do not perceive diminutive worlds that could have existed, now exist, or will exist, because of their very small resolution. As in nature, users can explore the essence of these structures and patterns.

“Attractors” help users navigate, like the Sirens in Jason’s famous Greek voyage. In this case, however, the Sirens facilitate navigation with their songs and magnetism. Their songs are spatialized so users can follow virtual “musical maps” that change as the melodies evolve with the users’ behavior. To improve navigation and interaction with virtual objects, the Sirens can modify various environment parameters such as sensor sensitivity and interocular distance, among others.

Collaborators


Music
Eve Tramullas

Soprano Voices
Eva Peng
Kristina Valcarce

Programming
Juan Francisco López
Yu Uny Cao

Spatial Audio Server (SAS), donated by Fraunhofer
Institut für Graphische Datenverarbeitung

Acknowledgements
The Labyrinth Project
Robotiker



Rebecca Allen
University of California, Los Angeles
rallen@arts.ucla.edu

Director
Rebecca Allen

Programmers
Eitan Mendelowitz
Loren McQuade
John Ying

World Designers
Daniel Shiplacoff
Damon Seeley
Jino Ok
Pete Conolly
Vanessa Zuloaga
Karen Yoo
Rico Magsipoc
Josh Nimoy

Sound Design
Mark Mothersbaugh
Mutato Musika


Additional Sound
Jay Flood

Emergence

Emergence is a PC-based, real-time 3D software system that explores the role of human presence in a world of artificial life. A unique interface system utilizes voice input and a haptic device for tactile feedback. Novel forms of communication between human participants and artificial life forms include symbolic and expressive sounds, gestures, and movements.

A proprietary 3D engine handles rendering and display of 3D, texture-mapped characters and environments, and a physics-based behavior system that enables complex behaviors and interactions between all objects in the environment. A high-level behavior scripting language allows for specification of behaviors and relationships between characters. Sounds are linked to objects and characters to enhance the sense of life and space.


This work is partially funded by a research grant from Intel Corporation.



Emergence

This hemispheric, see-through display reveals images, sounds, and movement in a head-mounted system. Its hemispheric virtual immersion gives users a transparent sensation of being covered with water without getting wet.

Water



Karrie Karahalios
Fernanda Viegas
Massachusetts Institute of Technology
kkarahal@media.mit.edu
fviegas@media.mit.edu

Collaborator
Judith Donath


VisiPhone

Visiphone is a communication object that opens a visual and auditory portal through space by visualizing the sounds flowing between two places. A continuous audio connection between two distant locations brings the inhabitants closer, allowing them to talk informally and easily. Yet using audio alone has disadvantages. It is difficult in a noisy environment, to know if one's voice has carried or if others are speaking at the other end. Furthermore, long periods of silence make it easy to forget the device, which then takes on a quality of covert surveillance.

VisiPhone's graphical rendering of the audio brings greater continuity and expressiveness to this connection. It portrays the existence of the connection even in moments of silence and it expresses the dynamics and inflections of conversation originating in both locations.

This rendering is evocative rather than technical; our goal in building the VisiPhone has been to create an aesthetic object that enhances sociable awareness.

More information about our research can be found at:
www.media.mit.edu/~fviegas/visiphone.html




Two different experimental VisiPhone interfaces: podium display (left) and dome display (right).

The abstract graphics convey incoming as well as outgoing audio messages on both ends of the connection.

In this familiar maze game, a marble rolls on a flat surface, its direction determined by two knobs that control the level on perpendicular axes. The surface is crisscrossed with a network of passages separated by shallow barriers and is accented here and there with the dreaded holes that swallow the marble and end the game. The conventional version of this game, which sits on a table top or precariously on your lap, prompts the question: What are the ideal dimensions and controls for engaging in this activity?

A 3D model of the maze is projected onto a human-scale interactive projection floor with an imaginary pivot point at the center. The model tilts, seemingly under the weight of the players, according to where they stand on the game surface. As the projected surface tilts, the marble moves through the maze and appears to obey the laws of gravity.

Human-scaled interactive systems succeed when a tight feedback loop is established between the actions of the player and computer-generated images and sounds. The metaField Maze achieves this by providing a fast-paced, continuous activity that demands skillful initiative. A game strategy is developed intuitively, and the player's entire body is used to express it by moving frenetically over the full surface without any specific orientation. This ambi-directional, kinetic quality hints at the elusive feeling of being inside a computer application and is enhanced by the slight tilting of the computer-generated model. The uncanny effect of challenging the player's sense of balance further contributes to a heightened sense of immersion.




Maze

HyperMask: Virtual Reactive Faces for Storytelling


HyperMask projects an animated face onto a physical mask that is worn by an actor. As the mask moves within a prescribed area (the stage), its position and orientation are detected by a camera, and the computed projected image moves accordingly. If the orientation of the mask changes, the projected image changes with respect to the viewpoint of the audience. The lips of the projected face are automatically synchronized in real time with the voice of the actor, who also controls the face's expressions.

As a theatrical tool, HyperMask enables a new style of storytelling. In the Millennium Motel, a self-contained system in a linen cart projects onto the mask worn. The actor pushes the cart and portrays a chambermaid ("Millie") who tells amusing stories set in the motel.

Mask




The prototype HyperMask system, as used in a live performance of Mujina, a Japanese ghost story.



"...TV has, some feel, introduced a kind of rigor mortis into the body politic."
Marshall McLuhan (Understanding Media, 1964)

Viewers confront a living room environment that contains a video display playing live-broadcast programs. The television imagery is overlaid with cross hairs within a circle, giving the impression that viewers are separated from the programs by a viewing scope. When viewers move within the installation, the audio and video fade out, and the cross hairs recede into a black screen followed by text that requests viewers to remain still. The television imagery and sound resume only after all viewers within the installation have remained motionless for at least five seconds. Each time the television program is switched off in response to viewers' movements, different text messages are provided on the screen, such as "Please Remain Still" or "Just Relax."



Digital Cloning System

The Digital Clone is a 3D, realistic “human” digital character that appears on screen with the look and feel of a live actor. A state-of-the-art facial tracking and animation system tracks the motion of an actor on a live set (without a cumbersome motion capture suit) and drives the movement of a 3D digital character. The Digital Clone is then composited into the live action.



Digital W.C. Fields



Digital Marlene Dietrich

Cloning

Collaborators

Christopher R. Wren
Alex P. Pentland

Ron MacNeil
Glorianna Davenport

Jeffrey R. Bender
Teresa D. Hernandez

MIT Media Lab

City

This immersive, interactive, and dynamically growing 3D web browser fetches and displays URLs to form skyscrapers and alleys of text and images that participants visit as if they were exploring an urban landscape of information. The system starts with a wireframe floor map of a chosen city or area. As participants follows paths, a virtual 3D Web-based world of information is built dynamically. *City of News* takes advantage of human abilities to remember the surrounding 3D spatial layout, helps participants recall and group information. And it invites them to create mental associations between information and geography as if they were living in a customized memory palace.

City of News is presented with two technologies:

1. A projected map of the chosen location. Participants walk on the map and trigger dynamic growth of a 3D Web World by walking along areas of interest. The sensing system is an untethered, wireless, real-time computer vision system that tracks position and gesture.
2. A wearable computer that allows participants to physically wander through the Millennium Motel. Based on their location, the wearable computer uses a combination of micro display, wireless spread spectrum, and embedded RISC processors to deliver text, graphics, sound, and video information.

The entire system is fabricated using soft modular packaging bonded with flexible interconnections. The electronic systems are highly integrated into a comfortable lining for use with a vest or jacket. They convert the wearer into a mobile internet node.





Informal. Surprising. Adventurous.
Risky. Fascinating. Popular.



Sketches & Applications

Animation



- 194 A Ghostly Figure Rising Out of an Evil, Dark Bog:
The Making of The Wraith from "The Mummy"
- 195 Cloth Animation for Star Wars: Episode I
"The Phantom Menace"
- 196 Creature Wrangling and Enveloping for Star Wars:
Episode I "The Phantom Menace"
- 197 Creating a Digital World from Scratch: The Launch
of the First Union Bank Advertising Campaign
- 198 Creating Digital Corpses for "The Mummy"
- 199 Creature Modeling and Facial Animation on Star Wars:
Episode I "The Phantom Menace"
- 200 Deep Canvas in Disney's Tarzan
- 201 Digital Cars
- 202 Directing 3D Animated Characters for Advertising:
Turning Marketing Strategy Into Storytelling Strategy
- 203 The Haunting
- 204 The Making of the Painted World: "What Dreams May
Come"
- 205 Multiple Creatures Choreography on Star Wars:
Episode I "The Phantom Menace"
- 206 Technical Animation Issues for the Battle Droids
of Star Wars: Episode I "The Phantom Menace"
- 207 Viewpainting Models for Star Wars: Episode I
"The Phantom Menace"

Richard Kidd
Cinesite Visual Effects

Art, Design, and Multimedia

- 208 A 3D Natural Emulation Design Approach to Virtual Communities
- 209 The Application of Non-Periodic Tiling Patterns in the Creation of Artistic Images
- 210 Computational Expressionism: A Model for Drawing with Computation
- 211 Explorations of New Visual Systems
- 212 Hyper-3D Paintings in QuickTime VR: Wunderkammer and Hyperaesthesia
- 213 MSA's Attractors: Navigational Aids for Virtual Environments
- 214 The Mutable Cursor: Using the Cursor as a Descriptive and Directive Device in Digital Interactive Stories
- 215 Nami
- 216 Nicholson NY/Pequot Interactives
- 217 Passion Spaces Based on the Synesthesia Phenomenon
- 218 Phene-: Creating a Digital Chimera
- 219 SaltoArte: Explorations in Spatial Interactive Multimedia
- 220 Setup of the Konsum Art.Server
- 221 Virtual Music Reproduction
- 222 VisiPhone

Technical


- 223 3D Facial Reconstruction and Visualization of Ancient Egyptian Mummies Using Spiral CT Data
- 224 3D Gait Reconstruction Using Two-Camera Markerless Video
- 225 3D Imaging for Rapid Response on Remote Sites
- 226 3D Physics-Based Brush Model for Painting
- 227 3D Rendering Effects for 2D Animation
- 228 A Level-Set Approach for the Metamorphosis of Solid Models
- 229 An Interface for Transcribing American Sign Language
- 230 Animating Bird Flight Using Aerodynamics
- 231 Animating Expressivity Through Effort Elements
- 232 Application of Computer Graphics for Design and Delivery of Conformal Radiation Therapy
- 233 Asynchronous, Adaptive, Rigid Body Simulation
- 234 Automatic Recognition and Mapping of Constraints for Motion Retargetting
- 235 Capturing the Motions of Actors in Movies
- 236 Color Super-Histograms for Video Representation: Preliminary Research and Findings
- 237 Declarative Behaviors for Virtual Creatures
- 238 Incremental Delaunay Triangulation
- 239 Dynamic Texture: Physically Based 2D Animation
- 240 Enhancing the Efficiency and Versatility of Directly Manipulated Free-Form Deformation
- 241 Fast Polygon Mesh Querying by Example
- 242 Filtered Noise and The Fourth Dimension
- 243 Freeform Curve Generation by Recursive Subdivision of Polygonal Strip Complexes
- 244 Handheld Interactions: Tailoring Interfaces for Single-Purpose Devices
- 245 Head-Mounted Projector for Projection of Virtual Environments on Ubiquitous Object-Oriented Retroreflective Screens in Real Environment
- 246 The Holodeck Interactive Ray Cache
- 247 Image Moment-Based Stroke Placement
- 248 Image Re-Composer

- 249 Image-Based Modeling, Rendering, and Lighting in Fiat Lux
- 250 Image-Based Techniques for Object Removal
- 251 Interactive CSG
- 252 Interactive Haptic Modeling of Tensegrities and Network Structures
- 253 Interactive Rendering with Arbitrary BRDFs Using Separable Approximations
- 254 LiveWeb: Visualizing Live User Activities on the Web
- 255 Methods for Preventing Cloth Self-Intersection
- 256 Modeling HIV
- 257 The Morphological Cross-Dissolve
- 258 Multi-Dimensional Quaternion Interpolation
- 259 Multifluid Finite Volume Navier-Stokes Solutions for Realistic Fluid Animation
- 260 Oblique Projector Rendering on Planar Surfaces for a Tracked User
- 261 Occlusion Culling with Optimized Hierarchical Buffering
- 262 OpenGL Texture-Mapping With Very Large Datasets and Multi-Resolution Tiles
- 263 Phong Shading at Gouraud Speed
- 264 Physically Based Anatomic Modeling for Construction of Musculoskeletal Systems
- 265 Postprocess 2D Motion Blur for Cel Animation
- 266 Projecting Computer Graphics on Moving Surfaces: A Simple Calibration and Tracking Method
- 267 Prototype System of Mutual Telexistence
- 268 Quasi-Linear Z Buffer
- 269 Real-Time Shadows, Reflections, and Transparency Using a Z Buffer/Ray Tracer Hybrid
- 270 Real-Time and Physically Realistic Simulation of Global Deformation
- 271 Real-Time Principal Direction Line Drawings of Arbitrary 3D Surfaces
- 272 Real-Time Translation of Human Motion from Video to Animation
- 273 Rendering 3D Objects into Photographs Taken by Uncalibrated Perspective Cameras
- 274 Representation of the Tactile Surface Texture of an Object Using a Force Feedback System.
- 275 Shading and Shadow Casting in Image-Based Rendering Without Geometric Models
- 276 Shape Extraction for a Polygon Mesh
- 277 Speedlines: Depicting Motion in Motionless Pictures
- 278 Stereo Analyst: Visualizing Large Stereoscopic Imagery in Real-Time
- 279 Tangible Modeling System
- 280 Tracking and Modifying Human Motion with Dynamic Simulation
- 281 Video Embodiment - MovieSpiral: Towards Intuitive/Comprehensive Interfaces for Digital Video Interaction
- 282 Virtual Car
- 283 Volumetric Modeling of Artistic Techniques in Colored Pencil Drawing
- 284 Wet and Messy Fur
- 285 Which Way Is the Flow?
- 286 WorldBoard: Enabling a Global Augmented Reality Infrastructure

Introduction

Informal. Surprising. Adventurous. Risky. Fascinating. Popular.

Those are the words that SIGGRAPH conference attendees most often use to describe Sketches & Applications, the forum for works in progress, preliminary drafts of promising inquiries, case studies of proven solutions to tenacious problems, and "the making of" presentations on how today's visual effects technologies are becoming tomorrow's routine tools.



This is where all the disciplines within and related to computer graphics and interactive techniques converge to exchange insight, inspiration, and just plain facts. Every year since it began (as skeptically received Technical Sketches at SIGGRAPH 94), this program has generated more and more submissions and attracted larger and larger crowds.

The jury deliberated, reviewed, and debated for hundreds of hours before selecting 93 presentations in three categories:

Animation


Again this year, animation technology advanced dramatically, and those who use it discovered some surprising and very imaginative ways to apply the technology in feature films and advertising. Fourteen animation Sketches & Applications include sessions on ghosts, creature wrangling, digital corpses, skin movement, animated paintings, cloth modeling and animation, and digital cars. The turn-of-the-century trend in this category indicates that the SIGGRAPH community is moving rapidly toward seamless CG realities.

Art, Design, and Multimedia

These SIGGRAPH 99 Sketches & Applications demonstrate aggressive extension of the principles of computer graphics to new visual systems. The jury selected 15 outstanding examples of how artists, designers, and engineers are always looking beyond the horizon. Topics include: digital drawing and painting, virtual communities, digital choreography, and synesthesia.

Technical

This year's Technical Sketches range from possible solutions for practical production problems to amazing flights of intellectual exploration. The jury selected 64 presentations that illustrate the very healthy worldwide spirit of adventure in the computer graphics community. Topics include: expressive animation, asynchronous simulation, virtual behaviors, dynamic textures, interactive rendering, mutual telexistence, recursive subdivision, new methods of quaternion interpolation and texture mapping, haptic surfaces, and a 3D physical model of HIV.





Adventurous

Acknowledgements

On behalf of the Sketches & Applications Committee, I thank everyone who submitted their work for consideration. Your contributions to the international SIGGRAPH community are deeply appreciated. In fact, without your enthusiasm and contributions, the community would not exist.

Many thanks also, to the members of the Sketches & Applications Committee and the SIGGRAPH 99 Conference Committee for their generous support, encouragement, and sense of humor. For their energy, commitment, and competence, I am indebted to Vicki Schaefer and Carrie Ewert, who managed and coordinated the submission, selection, and publication process. And I extend enthusiastic thanks to the other special people who made this program possible: Warren Waggenspack, Nancy Reynolds, Jill Smolin, Robin Myran, Carolyn Roemheld, and my colleagues at Cinesite Visual Effects.



Sketches & Applications Committee

Committee

Chair

Richard Kidd
Cinesite Visual Effects

Sketches & Applications Committee

Zsolt Krajcsik
Disney Feature Animation

Ken Musgrave
MetaCreations

Dena Slothower
Pratt Institute

Sketches & Applications Jury

Tom Appolloni
Harris Corporation

Curtis Edwards
Disney Feature Animation

Andrew Glassner
Microsoft Corporation

Madge Gleeson
Western Washington University

Michael Gleicher
University of Wisconsin

Steve Goldberg
Disney Feature Animation

Rex Grignon
PDI

John Hart
Washington State University

Jacquelyn Martino
Philips Research, USA

Marcus Mitchell
Digital Domain

Maureen Nappi
New York University

Aaron Pfau
Industrial Light & Magic

Michelle Robinson
Disney Feature Animation

Kathleen Ruiz
Rensselaer Polytechnic Institute

Brian Wyvill
Imagis GRAVIR/IMAG

A Ghostly Figure Rising Out of an Evil, Dark Bog: The Making of “The Wraith” from “The Mummy”

In the remake of the classic Universal Picture, “The Mummy,” director Steve Sommers requested a spectral apparition of a living person rising out of a bog and floating across an Egyptian burial chamber to rest on Imhotep’s lover to bring her back to life. This character became known as “The Wraith” during the post-production period at Industrial Light & Magic.

The animation of “The Wraith” was a collaboration of work completed by a team of character animators and technical directors (TDs) in a rather unusual way, due to ILM’s structure. Typically, the work completed by a TD and a character animator is done consecutively, where the character animator hands off the animation to the TD, who then places it in the scene. For “The Wraith,” the TD and the character animator worked cooperatively in creating the final animated look.

The principal animation was developed by creating a series of blended shapes using ILM’s in-house character animation software, called Cari (short for caricature). On “The Mummy,” the character animator developed the main animation and story for the entire scene, developing the motion for communication of the spirit-like ghost drifting across the set. Along with Cari, Softimage was used by the character animator for spline path animation, articulated chains, and animation by expressions.

The animation was then imported into Alias|Wavefront’s Maya software package, where the TD added cloth-like softbody animation on top of the shape animation that had been created by the character animator.

A softbody is a geometry in which the control vertices (CVs) are replaced by particles, which are dynamic objects influenced by forces like gravity, wind, and inertia. The original geometry and its CVs are the target geometry of the softbody and its particles. Each particle is attracted to the corresponding CV of the original object. The challenge with softbodies, and dynamics in general, is to keep a high level of control over them. Very often the dynamic elements in the scene are totally at the mercy of the forces that act on that scene. Particles seem to have a life of their own, which is partly what we wanted, because they can create natural-looking motion. But how does one keep the softbody particles under control, while keeping their natural motion intact?

The motion of the cloth-like material was controlled during the course of the scene using goal weights. At certain critical points, we needed to maintain the integrity of the character animator’s intent, so our goal weights needed to be keyframe animated.

When the elements were exported to Pixar’s RenderMan for rendering, two procedural turbulence displacement shaders were applied, adding an additional watery feeling. Opacity maps were animated by turbulence to create moving fingers on the edges. Environment maps were also applied to enhance the reflective quality of the surface. The secular component was enhanced by applying an iridescent shader.

The animation of the character animator, after it was imported into Maya, enabled us to use the velocity of the vertices to propel the particles, ultimately giving “The Wraith” an additional directional smoky quality. These particles were rendered separately and later composited in.

For this sequence, the director gave us a great deal of creative freedom to come up with an interesting effect: “Just make it look cool.” Every day, we would create a new character animation and a new simulation, which were reviewed in dailies by the effects supervisor, character animators, and TDs. We went through many iterations to achieve a look that worked.



Cloth Animation for Star Wars: Episode I “The Phantom Menace”

Tim McLaughlin
John Anderson
Industrial Light & Magic

In Star Wars: Episode I “The Phantom Menace,” George Lucas populated his alien worlds with fantastic creatures, many of whom wore clothing. This provided a new challenge for ILM. We had created clothing for other shows, such as the Martian robes in “Mars Attacks,” but nothing that required the realism and on-screen scrutiny that the clothing for the creatures in Episode I would face.



To meet the challenge, the research and development team went to work in early 1997 and identified the key needs of the characters in the movie and the scope of the problems to be solved. The main focus of the effort was on Jar Jar. As a fully digital co-star of the film, Jar Jar was required to perform a wide range of actions in over 400 shots. His clothing consists of a short leather vest resting on his shoulders and a leather “sweater” wrapped around his waist. Additionally, Jar Jar’s loose, floppy ears, while not considered clothing per se, were treated as if they were cloth.

There were two key challenges in designing the procedural animation software for the clothing: development of appropriate representations for the physical properties of the large range of materials that were needed for the film, and development of a set of controls that were to be used to define the performance aspects of the clothing. These controls allowed the artists to intermix traditional keyframe animation and procedural animation techniques in the same shot, and offered a wide range of creative control.

Success with Jar Jar’s cloth applications encouraged the crew to use the same technology on over 40 other models featured in nearly 800 different shots. For each shot, the motion of the cloth was treated as an aspect of the character animation. Initial parameters for the cloth, such as stiffness and damping, were developed by a creature developer for each character.

Once the appropriate “look” of the cloth was identified, the files defining it were accessed from the model database whenever an artist was importing the model itself. The animators were able to run the simulation as a stand-alone process or as part of a batch render. For shots requiring special actions, the artists were able to dial between keyframed or fully simulated actions, as needed. Objects in the environment were added as collision bodies, wind-blown or gravity altered, as required by the performance of the shot.

Application of the cloth technology ranged from full clothing, such as that worn by Darth Maul and Boss Nass, to hats, jackets, vests, and even to reins, antennae, and tassels. The clothing software was so easy to use, and so versatile, that the most difficult task was creative: deciding what the cloth should do and how it should look.



Star Wars

Creature Wrangling and Enveloping for Star Wars: Episode I "The Phantom Menace"

Meeting the challenges presented by the large number of complex digital creatures needed for Star Wars: Episode I "The Phantom Menace" required the full attention of a handful of creature developers equipped with specialized tools. With over 225 digital models, including nearly 100 photorealistic creatures, we needed to find new, efficient ways to both define the enveloping, or skin movement, of our models, and make sure that the models and their related data files were easily accessible to all artists on the show.



Beginning in the fall of 1996 and working through the spring of 1999, a crew of up to eight creature developers was busy full time managing the zoo of creatures for Episode I. The primary responsibility for this group was to ensure that the models retained the appearance of being organic as they were moved through their actions by the animators. This process, which we call "enveloping" at ILM, involves weighting the modeled surface to its animation controls. When done correctly, the skin moves in a natural manner according to the creature's motion and conveys the impression that the creature has an internal structure of bone, muscle, and fat under a supple skin.

Our new proprietary enveloping software, called Carienv, provided ways of defining the envelopes that were scalable to the performance needs of the creature, from very broad basic envelopes for background creatures to the very fine level of control needed for the hero models. Using Carienv, we were able to quickly scrub through animation on high-resolution geometry and build accurately detailed and robust envelopes. With some creatures, we had ample preproduction time to research relevant real-world creature movement and develop envelopes with proper skin and muscle motion. However, as the blocking of shots often changed, our enveloping system had to provide efficient methods for adjustment. Many times, we had to add complexity to existing envelopes, refit envelopes for creatures with changeable geometry resolutions, and in some cases refit one model type's envelope for use on another.

The secondary responsibility of the creature developer team on Episode I was to make the menagerie of digital creatures accessible to the team of around 40 animators and 70 technical directors putting together shots on the show. Caricature, the parent program of Carienv, was originally developed as a facial animation tool for "Dragonheart." It has since evolved into the tool for assembling all of the geometric information that goes into a shot. This includes animation data, mostly in the form of various Softimage scene files, high-resolution geometry data, envelope data for skinning, facial-animation data, cloth-simulation data, and texture data.

For Episode I, we added features to Caricature that enabled any artist to independently bootstrap a shot and setup and deliver fully textured and roughly comped thumbnail renders, without the assistance of technically skilled support staff. This autonomy enabled animators to work independently of the technical directors, who performed accurate checks of geometry elements before starting a full render.

Because we were often obligated to start animation on shots before final creature development was complete, the creature developers worked to ensure that the various pieces of data being used by each artist were compatible. Changes to models, whether animation, model, envelope, or texture related, were first tested on appropriate scenes. Once they were proven adequate for a wider range of shots, those changes were made available through a central database. For other artists, incorporating the new data was most often only a matter of using graphical file updating features built into the Caricature interface.

For Episode I, we were faced with the task of skinning and managing a larger and more disparate group of models than ILM had ever worked on before. These models had to be malleable to meet a wide range of performance requirements, and they had to be accessible to a crew of well over 100 artists. Through software developments and new techniques, we were able to scale the work required to develop each creature to fit its role in the show while at the same time coordinate quick alterations to models over a wide range of shots.

Creating a Digital World from Scratch: The Launch of the First Union Bank Advertising Campaign

Mary Beth Haggerty
Tim Stevenson
Industrial Light & Magic

Combining several different techniques for a desired effect is a staple of computer graphics production. In Industrial Light & Magic's work for First Union Bank, we employed several different cutting-edge and traditional techniques to achieve our final image.

We were confronted with the problem of creating a financial world consisting of familiar economic phrases and objects. Buildings would be shaped in the form of the words "Sell Now." The textures on buildings would be currency. Sculptural detail would consist of coins, dollar signs, and other elements. It was to become a visually robust world of financial information that, while connected with our current reality, also embodied the unpredictable world of finance.

Here is a breakdown of the opening shot, where we are introduced to this financial world. The shot combines a plate filmed on location, bluescreen people, pyrotechnic smoke, computer graphic people, and cars controlled with a particle system, digital matte backgrounds, computer graphic buildings, and painted buildings on cards. These techniques are combined into a seamless reality.

We filmed the plate and bluescreen people with a motion-controlled camera for consistency. We then imported this camera into the computer. A mock-up scene was created with computer-graphic mock-up models of the buildings to be represented in the shot. We then began constructing the network of roads and walkways that would connect the various buildings and regions in the world we had created. We delivered this city with no textures for the digital matte painters to use as reference. Each foreground building was painted separately. The reference frame for each structure was based on when the building's face was closest to perpendicular to the camera's direction of view. The background digital matte buildings were grouped and painted in a similar fashion. The paintings were then projected onto cards lined up at the reference frame and rendered in Softimage. The 3D computer-graphic buildings were modeled, and projection textures were applied for computational efficiency. Having 3D models in the scene was necessary to cue lighting changes. For the cars and people, we took advantage of the ability of Pixar's RenderMan (version 3.7) to load in different resolutions of data based on screen space. The path animation of the cars and people was done in Alias|Wavefront's Dynamation. The humans had animated cycles applied to these paths for realistic motion.

Intelligent use of different techniques can enhance the look of a production. Through the application of stage, computer graphics, and digital paintings, we were able to achieve images with a visual depth not often achieved with a single approach.

Digital World

Creating Digital Corpses for “The Mummy”

“The Mummy”, Universal’s modern remake of the classic horror film tells the story of Imhotep, a cursed Egyptian High Priest.

ILM’s Effects Supervisor John Berton, was given the task of bringing the title character to life in various stages of reincarnation, from walking skeleton to full human form and all the steps along the way. This key challenge required digital recreation of an actor, inside and out, that had the look and performance of real human bone, tissue, skin, and body mechanics.



Model Supervisor Jim Doherty, Viewpaint Lead Catherine Craig, and Computer Graphics Supervisor Mike Bauer began the process of creating the digital monster. Reference materials created under the guidance of Art Director Alex Laurant became the bible for the look of the mummy in his various stages of decay. The geometry, built in three main stages, was with six viewpoint texture sets ranging from minimal rotted flesh over bone to a complete construction of Arnold Vosloo, the actor playing the undead monster.

A complete human model was constructed in five interdependent layers: bone, flesh, decayed skin, human skin, and clothing. Stage one had the most variance with rotted skin and withered muscles. Stage two was the next generation and took the creature closer to its final human state. The third major model was sculpted to match the actor and allowed creation of the final stages of Imhotep’s transformation.

While the sculpts of the three models differ greatly, they are identical in surface name, orientation, resolution, center, and group association. Keeping these elements matched allowed transformations from one model into another. It also was a time-saving issue. The texture set-up was shared between the models. The surface enveloping for successive models started from the previous model. And the set-up of procedural elements was identical.

Because of the complexity of the creature, the geometry was too dense to get the final sculpt with the model alone. Large level displacement mapping provided the “secondary modeling” used in conjunction with color, opacity, bump, and specular maps.

Animation Supervisor Dennis Turner and Lead Creature Developer Rick Grandy created the animation control system for “The Mummy”. Because of the realism needed in the creature’s movements, motion capture was used to sample Vosloo’s performance. A single animation model incorporated both keyframe control and motion-capture input. It allowed the animation team to refine the motion capture while adding enhancements to it. That same animation model was used for every CG mummy in the film, from Imhotep to the soldiers.

The mummy is the first creature at ILM to simulate the relationship between the skin and its underlying muscle structure. The system was developed out of the need to visualize animated tissue under the surface, which was to be made evident through opacity maps. The skin needed to behave as real skin, and it was decided that the traditional skin technique, a process referred to as “enveloping,” would only be the beginning of what was needed. This process puts the animation directly onto the hires surface geometry and becomes the basis for the flesh simulations. Enveloping is a process of assigning a value to each surface control vertice for the amount of effect it receives from each control in the animation model.

The dynamic simulations on the creature were divided into two categories, cloth and flesh. The cloth engine simulated pieces of hanging flesh and cloth that needed to show the effects of inertia and motion through space, such as bandages and cobwebs. The flesh simulation engine was developed to deform the outer skin surface to the underlying bone and muscle geometry. This involved creation of a secondary animation model that controlled individual muscles and allowed for independent and realistic motion.

Creature Modeling and Facial Animation on Star Wars: Episode I “The Phantom Menace”

Geoff Campbell
Cary Phillips
Industrial Light & Magic

The making of Star Wars: Episode I “The Phantom Menace” involved an unprecedented amount of digital creature modeling. A team of 10 creature modelers created six major talking creatures, including Jar Jar Binks, who is one of the central characters in the story, together with 65 other life-like creature models. The duties of the modelers at Industrial Light & Magic involved building the static models, either from scanned maquettes or by sculpting from scratch, building libraries of facial shapes and expressions for the talking creatures, and working with the envelopers responsible for the flexible skin models to sculpt corrective adjustments at the last minute when something didn’t look right. The models created for Episode I presented a variety of challenges, from experimenting with early concept designs to quickly putting together new creatures added on the fly as the show progressed.

ILM has a variety of proprietary in-house modeling tools, including software to build surface models directly from scanned data for creatures that were originally designed and approved in clay. A surface sculpting tool, called Isculpt, provides a very direct way of manipulating surfaces by pushing and pulling in ways analogous to modeling with clay. Isculpt makes it easy for modelers to make creative changes to a model after it has been built, whether it originated as a scanned model or was built from scratch in the computer.

ILM’s facial animation system, called Cari, is based on shape animation, and the process of creating a talking creature begins with the modeler sculpting shapes that define facial movements and expressions. Cari is designed to handle creatures of very great geometric complexity and visual realism, and to give direct creative control to the artists, so that the process of describing the movements of the face and skin is largely a creative one, not a technical one.

Jar Jar Binks was the first creature created, and he received the most attention from the beginning. There were no maquettes built for him. Instead, the computer model was based only on early-concept artwork, and the director wanted to experiment with different looks before agreeing to a design. In fact, the original design evolved dramatically, due to difficulties in structuring the mouth so that the creature could talk convincingly.

A very different example was the two-headed Pod Race Announcer, which was originally planned as a practical effect. Only late in the process was it decided that the heads would be animated digitally, so the head models had to be constructed very quickly by altering other available models of similar structure.

The digital Yoda that appears at the end of the film represented a far extreme of the modeling experience, as it wasn’t until the shot was in production that the director decided to include Yoda. At that point, it was too late to arrange to film the practical Yoda, so we built a digital model and augmented it with some simple facial expressions in a matter of several days.

Another important role of the modelers on Episode I was to correct imperfections in the movement of the skin, both generically and shot-by-shot. ILM uses a variety of techniques, including soft body dynamics, for animating skin, and each presents ample opportunity for the skin to move in ways that are not aesthetically appropriate. A major advantage of our skin animation system is that it allows for after-the-fact improvements that can be made by the artist at the last minute before a shot is finalised.

Deep Canvas in Disney's Tarzan

What if you could paint a painting, then have the brushstrokes themselves come alive and move around?

This is what Deep Canvas is: an animating painting. Rather than texture-mapping, where the end result of the painting process is wrapped onto a surface, this process instead animates the events that make up a painting: the brushstrokes themselves. The process requires an entirely different type of renderer, one that actually repaints the entire painting for every frame.

Created in response to an unusually difficult problem (animating 10 minutes of painterly looking, densely lush jungle for the Disney movie "Tarzan," with a relatively small crew), the process leans heavily on the skills of the traditionally trained artists who work at the Disney Feature Animation Studio.

Early on, it was determined that an "automatic" approach would likely fail to capture the qualities that a good painter, with a lifetime of practice, can intuitively apply to the painting process. Since Disney has on hand dozens of the best painters, it seemed only natural to figure out a way that those artists could best contribute to these 3D scenes.

About a year before, Barb Meier had developed a painterly rendering technique that showed that brushstrokes could be rendered individually to look like a successful "painting in motion." Designed as an automatic process, it was "hard-wired" to create a Monet style of impressionism. Arguably, it could have been tweaked or rewritten to mimic any painting style. But we felt that an even more interesting challenge was to create an approach whereby the style was entirely in the hands of the painter, not the tool. Having seen what good painters can do with nothing more than some colorful goo and a stick with a bunch of hairs on the end, we couldn't wait to see what they could do with a stylus and some intelligent software.

First, we had to create painter-friendly software. Fortunately, we didn't need all the gizmos that a full-featured commercial program would need; we simplified the toolset down to the bare minimum necessary to satisfy the artists. As the artists paint, their brushstrokes are recorded for reuse.

When the time comes to render the 3D scene, the finished painting is disregarded entirely. Instead, each brushstroke is given information as to where it was intended to be in 3D space. Then we begin the "repainting" process, where the renderer, in effect, steps through each "painting event" (stroke) and determines where it should be in camera-space for that particular frame. It then paints that stroke into a buffer.

The cumulative result of all these painting events is called, of course, a painting, for the same reasons that the original is called a painting. The only difference is that in this painting, the brushstrokes might be in a different place than they were when the artist originally painted them. The completed sequence of paintings is an animation; just as a sequence of drawings is an animation.

Because of the unusual nature of the Deep Canvas process, there had to be an extremely tight relationship between the artistically adept people and the technically adept people. In fact, we found that the process went more smoothly the more the skillsets overlapped.

The use of CG cars at Digital Domain was first employed as a technique in the Plymouth "Neon" 1996 automotive advertising campaign, which presented a number of playful Neons bouncing off of an unseen trampoline. Conventional wisdom had the cars being shot practically, and then manipulated digitally. But while working for Director Terry Windell of A Band Apart, Visual Effects Supervisor Fred Raimondi came to the conclusion that the most artistic and cost-effective way to achieve this "effect" was to build and animate the cars digitally.

Raimondi's thought process went something like this: "What does computer graphics do best? The answer: shiny metallic things. Ding! A car is a big shiny metallic thing, so it should be natural!" With this in mind, he brought CG Supervisor Eric Barba and Digital Domain's Lightwave team onboard to recreate the photo-real look that was crucial to making the spot work. Relying on extremely high-resolution models, the team created all the textures, down to the safety-glass decal on the windows, that allow the viewer to distinguish reality from fiction, and, in this case, practical from CG. According to Raimondi, "lighting the cars became the key element in the process. When real cars are lit, they use a Fisher box, which actually isn't as much lighting the car as it is reflecting a big white card onto the car. So what we had the artists do was, instead of putting light sources over the car, we had them put a big white card over the car. We then lit that card and, like magic, we had a great CG car!"



The basic concept of CG cars took its next step into the mainstream with Dodge "Time" and "Time 2." Using a technique similar to that employed in the "Neon" spot, Director Terry Windell, this time paired with Visual Effects Supervisor Ray Giarratana, had Eric Barba and his Lightwave team create a Dodge Viper and then, for "Time 2," two additional digital vehicles.

General Motors' Phil Guarascio, General Manager, group VP-marketing and his advertising managers saw the success and cost benefits of the "Neon" and "Time" campaigns and in 1997, convened a seminar at Digital Domain for GM executives to determine how GM could take advantage of the latest advances in digital imaging. Of particular interest to GM were situations where practical photography was unavoidable.

This was the case in two spots for General Motors: Chevy Blazer "Roads" and "Obstacles," both of which were created in 1997 and then digitally updated in 1998 to reflect model-year modifications. GM, with Director Eric Saarinen and Plum Productions, and Visual Effects Supervisor Fred Raimondi, was able to realize the cost benefits while receiving original and exciting spots. The spots incorporated original tracking data from existing practical photography shot on location. This allowed new "updated" features to be "added" to the pre-existing campaign, which had been highly regarded the earlier year.

With the 1999 Pontiac campaign for GM and creative agency DMB & B ("Metal City" and "Metal Desert"), Director Ray Giarratana sought to create a highly stylized and entirely CG environment with a photo-real CG car. In particular, the goal of this campaign was to create digital camera moves intended to "mime" the characteristics of practical photography. In so doing, Giarratana, with Digital Director of Photography Eric Barba and CG Supervisor Wayne England, succeeded in pushing the animation to the next level – tricking the viewer into believing that a practical vehicle had been shot and composited into a digital universe.

This latest series of commercials is, in many respects, the culmination of the efforts of Digital Domain to bring CG cars to life, and, perhaps most importantly, they represent the increased flexibility that digital imaging provides in creation of advertising images.

Directing 3D Animated Characters for Advertising: Turning Marketing Strategy Into Storytelling Strategy

A spot is a 30-second movie. While the constraints of a commercial can seem overwhelming, remember how restrictive many classical forms are. Sonnets and Haiku, for instance. A tremendous amount of great work has been created by treating constraints as liberating, so why not take a similarly optimistic and enabling approach to making spots?

To successfully direct a commercial, you must turn marketing strategy into storytelling strategy. The availability of 3D animated characters provides a powerful and creative solution to this challenge. The first step is to reduce the marketing strategy or the message of the campaign to a singular emotion. Then write the story to evoke that emotion from the audience.

Don't be fooled when a corporate marketing strategy seems to be more of an idea than an emotion. A story with an idea but without emotion always fails. The story must convey the emotion and then carry the idea along for the ride.

A good story must draw the audience. Typically with a sympathetic protagonist. Well-designed 3D characters can provide wonderful and stylized characters for your stories. Design can build empathic protagonists, heroes, and villains, according to the needs of your story.

Use the language of cinema to persuade your audience to identify with your characters. Hitchcock persistently demonstrated that the most effective way to build strong psychological audience identification with a character is to use the POV sequence. Take advantage of it!

Visual storytelling is, simply, editing. To edit is to speak film language. If you don't speak it yet, hire someone who does or take the time to learn it.

To create the most dynamic and powerful shot sequences and to do so while working for (and hopefully with) clients who are often untrained and unwise in the crafts of editing, filmmaking, and animation, I suggest the following approach:

1. From the first day of the project to the last day, look at sequences edited to time (30 seconds, on occasion). Many a gorgeous storyboard simply won't cut, so on the first day of the production, scan in the agency concept boards and cut them to time. This cut will immediately demonstrate the weaknesses (and strengths) of those agency boards.
2. Show this cut to the agency team, so they can see with their own eyes what does work and what doesn't work in their board.
3. Now rewrite the story as necessary. Generate a new storyboard, scan the frames, and edit them to time. A scratch track should be used at this stage, if possible. Study this "board-o-matic," analyze it, and then modify the story or storyboard as needed to make the story stronger (to more strongly evoke the emotion of choice).
4. Once this "board-o-matic" is working and approved, direct a small animation team to create a 3D animatic with appropriate lenses and camera moves. This 3D animatic will finally let you know if the shot sequences will work in what will be their final form: 3D computer animation. Change timings, lenses, framings, or anything else as needed.
5. Is the story told? Is the emotion conveyed? If so, now (and only now) are you ready to enter into full production with your production team. You have a story, and you have the shot sequences to tell that story.

This method can be used for live-action, especially effects sequences, to great advantage and for great fiscal savings. So, while you exploit the mobility and weightlessness of your virtual camera, don't forget to take full advantage of the power of the 100-year-old language of cinema. After all, for better or worse, it's become the universal language of our time.



The Orkin Man strikes a manly pose while starring in *WAR ON PESTS*.



The Honey Bear, played by Dom Deluise, yells for help in *SPYGLY*.

This sketch analyzes two scenes from “The Haunting.” In one scene, the heroine combs her hair in front of a mirror. She watches in horror as her hair takes on a life of its own, culminating in an elaborate braid before she shakes free of the ghostly encounter. The second scene involves the sleeping heroine. As she drifts into sleep in a giant antique bed, a ghost child enters through the window, causing a CG silk curtain to billow and rise. The ghost child continues her journey by traveling under the satin sheet of the bed and up to the pillow where her silhouetted form dissipates.

The CG hair had to be completely photorealistic, even when shown in close-up and inter-cut directly with the actor’s real hair. A strictly procedural-based method of growing hair was not feasible, since the hairs had to grow in an orderly yet chaotic manner (uniform but unique at the same time). For maximum control over the modeling, the process started with a foam head with armature wires inserted into the scalp. These wires were digitized and represented by splines in the computer. Since modeling and animating each hair would have been overly laborious, software was applied to a reduced number of hair splines as guides for growing additional hairs to densely populate the head. The hair had to conform to the physics of the reality. Not only did it have to respond to gravity and inertial forces, but it also had to be combed. The hair had to transcend the realities of action and reaction, fall under supernatural control, gather itself together, and twist itself into a bun.

The silk sheet shot involved creation of CG fabric that takes on the shape of a child as it moves from the bottom to the top of the bed. The bed sheet utilized pseudo-ray traced renders to create displacement for geometric objects. These renders, together with standard key-frame animation and manipulation of a digital puppet, produced dynamic ripples along the object, creating the effect of a child moving through the satiny sheet, leaving folds and gathers in its wake.

Unlike other attempts to deform geometry, Tippett Studio’s dynamic particle simulation (using a custom-spring algorithm) does not produce stretching, so the sheet doesn’t look rubbery. It has all the natural movement one would expect from a sheer material like satin. The goal was to show a child in a sheet, not a child under a sheet.



The Making of the Painted World: “What Dreams May Come”

The painted world of “What Dreams May Come,” is transformed over 8.5 minutes of live-action photography into moving painted imagery in the style of 19th century painters, such as Casper David Friedrich and Claude Monet. “What Dreams May Come” required unique analysis of how artists and computers perceive and interpret the world.

Vision

In order to present the style of painting that the director Vincent Ward felt most appropriate to the emotional pitch of a shot, our digital artists learned to view the live-action footage using an eye similar to that of a fine artist working in oils. Additionally, the shots reflect the emotional state of the character, who was starting to realize his ability to transfer his environment to the painted world.

Process

The live-action plates were analyzed for motion to retain the feel of the natural world, such as wind blowing through flowers and trees or water ripples over a lake. The plates were then graded as a source of color palette for the paint particles, based on a style of painting influenced by 19th-century artists. Areas of the plate were recognized and isolated as layers to be applied with particular paintstroke styles, much as a painter will fill and layer details on the canvas.

The live-action plates were re-constructed in 3D space. We could then “grow” plants, flowers, and grass, and replace trees as elements to be composited back into the painted plate. We also added 3D effects such as light beams, cloud shadows, and waterfalls to either the live-action plate or the matte painting to enhance the other-worldly light and atmosphere inherent in the nature of the painted world.

Technical Challenges

The painted world sequences required untreated actors who freely move about and interact with the environment. The painted world also required the spatial relationship of an immersive 3D environment, with the effect of a 2D application of Paint. The viewer would have to register this as a painting yet also believe it to be a 3D world. We had to retain the motion of the live-action plates and the natural world. Motion-control rigs could not be used during filming in the wilds of Montana. The end result had to include character interaction with the environment, yet their rendering was not to be “painted.” To meet this challenge, software called Motion Paint was developed in-house during the production. This toolset allowed us to match motion and spatial changes within the scene and attach 3D objects (such as wet paintstrokes) to the motion data. We were able to develop fast interaction between the artistic demands and technological solutions throughout the production.

As the artistic ideas were developing, we developed the software to meet these challenges. For example, one shot of grey static sky shouted out for a transformative Van Gogh swirling sky effect at the point in the story sequence where the character is starting to realize his influence on his environment. We were able to add additional software features that were subsequently used throughout the entire sequence.

New Vision

“What Dreams May Come” presented the opportunity to utilize 3D animation and reconstruction in a realm far from the “creature feature” of recent years. The technologies developed are of the lineage of visual effects that films like “2001” realized, where subjective immersive environments are key to the cinematic experience.

Multiple Creatures Choreography on Star Wars: Episode I “The Phantom Menace”

Marjolaine Tremblay
Hiromi Ono
Industrial Light & Magic

Boom! A battle droid has just been destroyed, but other droids keep moving forward. That is how our battle with the particle begins.

Creating the Naboo ground battle on Star Wars: Episode I “The Phantom Menace” was not an easy task. Organization and efficiency were our best allies. Working in Softimage, a group of character animators created a series of animation cycles needed for the battle while a group of technical directors (TDs) created and rendered particle choreographies in Alias|Wavefront’s Maya.

The complexity of each of the cycles varied. Character animators often had to combine motion-capture data with keyframe animation, run a series of cloth simulations, and create shape animation for each of the creatures or droids.

Each cycle needed to have a high level of detail. A cycled creature might be placed right next to a hero character or far from the camera. We could not play our animation to a certain camera angle like a normal shot because the same cycles were used in various shots with different camera angles. This meant that each and every cycle had to be working from the back, the sides, or the front. It elevated the amount of work and detail necessary in a cycle, pushing the animators to perfect their animation techniques. Combining two or more characters was also a challenge, since they needed to be treated as a single unit.

Using a proprietary motion database worked to the character animator’s advantage. The tool helped us manage and store the ever-growing number of motion cycles. It allowed us to follow an easy naming and numbering convention for each of the files needed for each individual creature cycle.

When an animation cycle was finalized and checked in the motion database, the TD took over. The animation cycle data was converted into renderable hires, midres, and lores rib files, as well as a more simplified polygon mesh called cube data, which was used to represent the creatures in Maya.


Based on the creature’s distance from camera, the latest RenderMan version let us render the creatures using the different levels of rib files.

During “Return of the Jedi: Special Edition,” a new technique allowing thousands of human figures to be choreographed with particles was in its infancy. As it evolved, we wrote a plug-in for Maya that displayed cube data, allowing us to manage more creatures in a growing number of sequences on a show like Episode I. The plug-in directly draws the geometry using Open GL display lists, which allows us to check the massive number of creatures that were animated in a shot. Several plug-ins and expression scripts were also written to enhance control.


In Maya, each loaded cube was assigned to a particle. For each of the particles, an animation cycle and a frame range was set using a particle rule. We created several generic rules for the particles that let the TD procedurally choreograph the particles. We positioned the particles on the ground and animated them forward, sideways, or turning. The velocity was determined by the rules and by Maya fields and collisions, but each cycle had a distance curve that restricted the magnitude of the displacement, keeping the creatures from sliding on the ground.

When a choreography was approved, we saved all the necessary information for rendering, such as materials, colors, and cycle data. To win the battle and keep total control of the crowd behavior, we used a procedurally choreographed-animation-cycles method instead of behavioral simulation techniques.

Finally, all the technologies began to work together, allowing us to animate and render an unprecedented number of detailed creatures with elaborate and varied motions. At last, the scene with all its creatures came together, and the grandest digital battle ever filmed took shape before our eyes.



Technical Animation Issues for the Battle Droids of Star Wars: Episode I "The Phantom Menace"



The battle droids in Star Wars: Episode I "The Phantom Menace" play a large role throughout the film. They are seen invading the city of Theed, engaging in action-packed combat sequences with Jedi Knights armed with light sabers, and battling an army of Gungans on the grassy plains of Naboo. As a result, two distinct factors were evident in the decision to use motion-capture technology in the early preproduction phase:

1. They were, by design, very humanoid in nature and, as such, required a hyper-realistic range of motions.
2. They appear in large numbers making the duplication of basic actions and motions a necessary production need in order to complete the shot.

Motion capture technology was utilized for the motions of the Battle Droids. There were over 200 selected samples of motion-capture data of the Battle Droids. The majority of the captured motions were designed to be generic movements such as run cycles, walk cycles, marches, shooting weapons, dodging movements, retreating movements, and dying. There were, however, numerous shots requiring a specific action such as walking down a staircase on a practical set or pushing through the Gungan battle shields.

For practical reasons, we needed the resulting motion on the same controls that the animators would regularly use. In this way the motion can be altered, blended, or discarded for any of the shots in the film. To meet this need, software was developed that converted the data onto any constrained animation hierarchy. This allowed for the conversion of motion capture data directly into the animation control systems. The converted motion information appeared as if it had been animated by an artist but had the controls available to manipulate the captured data.

With this technology, we were also able to use motion capture for various Gungans and distant pod race attendees even though the majority of the Gungan animation, including the Jar Jar Binks character, were keyframed.

In the film, the Battle Droids are no match for the powerful Jedi Knights. Wielding light sabers, the Jedi Knights are able to slice through the droids like a hot knife through butter. Ultimately, there were over 25 shots in the film that depicted 49 sliced droids. To achieve this effect, several versions of the sliced droids were created.

Each Battle Droid consisted of at least 204 dynamic elements, including 50 rigid bodies and 154 dynamic constraints. The actual slicing of the droids was achieved by using a rigid-body system created specifically for Episode 1.

The action of a droid was converted and read into the interactive dynamic system. The basic skeletal "bones" were converted into rigid bodies. At that point, dynamic pin was used to represent a ball and socket connection (three degrees of freedom) and two pins created a hinge-like connection (one degree of freedom). The geometry of the droid became the collision model that would then hit the ground or help to keep parts of the droid's body from penetrating other body parts. Other dynamic elements, like springs, helped to create an illusion that the droid's joints were stiff and had some resistance to external forces, thus altering their poses.

The Battle Droid sequences presented many technical and creative challenges for the motion-capture and physically based technology and specialists at Industrial Light & Magic. As with many of the techniques created for production needs, the images created using motion capture and dynamic simulations for Episode 1 have laid the groundwork for basic expectations that will become perfected as the ability to integrate these techniques into production grows.

Viewpainting Models for Star Wars: Episode I “The Phantom Menace”

Jean Bolte
David Benson
Industrial Light & Magic

One of the most significant aspects of successfully combining computer-generated models with live action has to do with the quality of the paint on the model. Computer model painting involves not only color, but also detailed sculpting of wrinkles and aging as well as any additional texture, such as fur. At Industrial Light & Magic, this process was first used on the dinosaurs in “Jurassic Park.” With the development of an in-house paint software called Viewpaint, not only wrinkles and fur, but also feathers, rust, metal patina, decay, and corrosion have been incorporated into the models. The stage has been set for just about any look, from photorealistic to stylized, to become a possibility.

With Star Wars: Episode I “The Phantom Menace,” facing a number of new and different challenges became the rule, not the exception. With over 200 different characters, creatures, props, and hard-surface models (some aircraft had as many as 5,000 individual pieces), the number of shots completed was unprecedented in any effects film to date. Clearly, handling the sheer size of the project, in a relatively short time, became the biggest challenge for ILM’s Viewpaint team. Working closely with the artists, who in turn made suggestions based on their experiences with the tool, the software team came up with a number of upgrades that provided a dramatic increase in our efficiency and productivity.

One of the challenges we faced was the introduction of CG characters wearing clothing, a first for ILM. As painters, we needed to pay close attention to the subtleties of material: how wrinkles fold and bend, how shadows of folded material are impacted and so on. Our software team wrote features into the package to facilitate our need. For example, wrinkle maps were created for Jar Jar’s shirt and trousers, which were painted in sequence then applied as displacement maps, somewhat like replacement animation. In order to see how the paint would stretch, where the point of impact occurred in crash sequences, etc., we wanted to see our paint jobs in motion. Again, working closely with the Viewpaint artists, the software team added the features that allowed us to accurately assess the painterly movement of fabric on the CG characters.

The character Jar Jar Binks is a good example of one of our creature paint jobs. Because he is a major character with an enormous amount of screen time, he came under much scrutiny during the development phase in pre-production. He was painted with color, bump, and specular maps. Additionally, to make sure that he would hold up well during extreme close-ups, we used the smallest paintbrush possible. During this presentation, Jar Jar is examined in stages, from the first concept art to the finished paint job, showing all the layers of paint that were necessary in the creation of realistic skin.

After we demonstrate the painting progression on one lead character, we show a brief montage of all the creatures and hard-surface models, including interiors, that ILM’s Viewpaint team was responsible for in the completion of Star Wars: Episode I “The Phantom Menace.”



Models

A 3D Natural Emulation Design Approach to Virtual Communities

The design goal for OnLive's Internet-based Virtual Community System was to develop avatars and virtual communities where the participants sense a tele-presence, so they are really there in the virtual space with other people. This collective sense of "being there" does not happen over the phone or with teleconferencing; it is a new and emerging phenomenon, unique to 3D virtual communities. While this group presence paradigm is a simple idea, the design and technical issues needed to begin to achieve this on Internet-based, consumer-PC platforms are complex. This design approach relies heavily on the following immersion-based techniques:

- 3D distanced, attenuated voice and sound with stereo "hearing."
- A 3D navigation scheme that strives to be as comfortable as walking around.
- An immersive first-person user interface with a human vision camera angle.
- Individualized 3D head avatars that breathe, have emotions, and lip sync.
- 3D space design that is geared toward human social interaction.

These techniques, which borrow from disciplines such as group dynamics, facial animation, architectural design, virtual reality, and cognitive sciences, allow the system to draw from the natural social neural programming inherent in all of us rather than creating artificial, social-enabling user interface mechanisms. The main goal of all of these techniques is to support multi-participant communication and socialization.

Community Comes from Communication: 3D Voice with 3D Navigation

The structural process of a community, whether real or virtual, is communication, and the most natural human form of communication is verbal. Verbal communication has both the explicit and the implicit message encoded in it.

We designed 3D spatial, multi-participant voice-with-distance attenuation and stereo positioning. Avatars closest to you are heard the loudest; those to your right, louder from your right speaker, and so on. Using this approach, the user interface becomes as simple as navigating towards the avatars you want to talk to and away from those you no longer want to talk to, just like you would at a real cocktail party. By using spatial sound with 3D navigation, natural group dynamics situations occur. Several small circular conversational groups of three to six avatars form and dynamically reform depending on individual and group social preferences.

Avatar Design: Binding the Pair – You Are Your Avatar.

Given the finite CPU/polygon/bandwidth resources, we invested them first in face-based avatars. The body, with its hand gestures and body language, is secondary for human communication and can be added as our resource limitations improve. The goal for us is what we call "binding the pair," binding the real person at the computer with a virtual avatar in cyberspace so the user experiences this feeling of tele-presence, of really being there. You cannot believably bind a person with an inanimate object or a texture-mapped photograph that does not emote. We tried to achieve "life" and believability with avatars that have autonomous blinking and facial movements (for example, "breathing"), that lip sync to their voices, and that can display (at user control) a range of emotions.

We now have some early positive indications that this technique is working, because it has been noticed that users make "eye contact" with each other. They turn toward the speaking avatar and can feel uncomfortable when another avatar comes too close and "invades their personal space." This last point was very encouraging considering our goal of "binding the pair." If someone in real life comes too close to you, you feel an uncomfortableness along with a physical tightening of your stomach muscles. This same sensation happens in the OnLive worlds, which shows that users perceive at some level that they are really there with other people, and that avatars are perceived as beings, not as objects being manipulated by other users on their home computers.

Evolving Issues: Getting Beyond Environment and Into Community Creation.

Even if we achieve all these goals (which we are still far from doing), we would end up with users feeling they are their avatar, and that they are really in a place with other avatar beings, collectively existing in the same virtual place. However, this merely provides infrastructure (the air, streets and buildings of a community), not a community. As we tackle the infrastructure issues and welcome thousands of people at our sites every day, we are just beginning to grapple with the social community issues. This requires the involvement of social engineers, cognitive psychologists, event producers, reporters, and sociologists, to begin to understand the nature and requirements of a virtual community.

Acknowledgements: Rod MacGregor, Henry Nash,
Dave Owens, Ali Ebtakar, Stasia McGehee, James
Grunke



The Application of Non-Periodic Tiling Patterns in the Creation of Artistic Images

Kenneth A. Huff
ken@itgoesboing.com

The physical world contains infinite examples of patterns of repetition, from the honeycomb of a beehive to the radial and spiral patterns found in a sunflower. While these patterns contain recognizable structure, they also have imperfections or a level of randomness. Computer-generated imagery often contains similar patterns, but frequently without the subtle imperfections found in nature, thus alerting our minds to the synthetic aspects of the image. The use of a non-periodic tiling pattern allows for an appealing structure along with natural randomness.

In 1704, a Dominican priest, Sebastien Truchet, painstakingly analyzed patterns by systematically placing multiple tiles, each divided by a diagonal line, into two different colored sections. Truchet's purpose was... "to be able to form pleasing designs and patterns by the arrangement of ... tiles." While Truchet was primarily concerned with symmetry and periodic repetition, these simple tiles may also be used as the basis for non-periodic patterns.

A pattern of continuous closures can be formed by using a grid of tiles, each consisting of a square divided into three sections by two quarter circles placed on diagonally opposite corners, and randomly orienting each tile constrained to 90-degree increments. The pattern is formed by the quarter circles. As is often the case with patterns in nature, these patterns contain a recognizable structure along with random qualities.

Additional levels of complexity can be produced by combining, or layering, a number of these patterns. The interference between layers can produce many unique patterns, each maintaining structure while remaining non-periodic.

Using a wide combination of commercially available software, the author has used this non-periodic pattern in various works, sometimes as a primary element, often to add an additional level of subtle detail.

Patterns




Figure 1
"99.1," 1999. The background pattern of this image is made up of four layers of a duplicated continuous curve. While similar in appearance, the pattern is not based on a non-periodic tiling pattern, but rather was created freeform.




Figure 2
"99.4," 1999. The geometry in this image is based on a non-periodic tiling pattern. Two grids of tiles were layered to form the basis for the profile curves used to create the geometric forms.




Figure 3
The development of the pattern used in "99.4".
1) A custom typeface was developed for two characters, one for each tile. A single character is highlighted in the illustration. A list of random characters was generated to form a grid. 2) Spacing of the characters was adjusted to remove gaps. 3) The grid was duplicated and offset from the original. 4) and 5) Additional manipulation produced the final 2D pattern.

Computational Expressionism: A Model for Drawing with Computation

Computers have introduced new concepts into the artist's vocabulary (interactivity, dynamism, and emergence) that shape the development of new forms of visual expression. In particular, the computational medium suggests a new approach to the act of drawing. Instead of duplicating the methods and materials we know from traditional media, we seek to develop a different perspective on visual thinking, one that involves a more active participation in the higher-level design of drawing tools.

The work consists of 24 Java-based drawing programs, each constructed and used by the author to create specific series of drawings. We define the "computational line" as a sophisticated drawing element controlled by various parameters of a gesture, such as speed, direction, position, or order. Complex shapes can be created with very simple gestures. A computational line has three attributes:

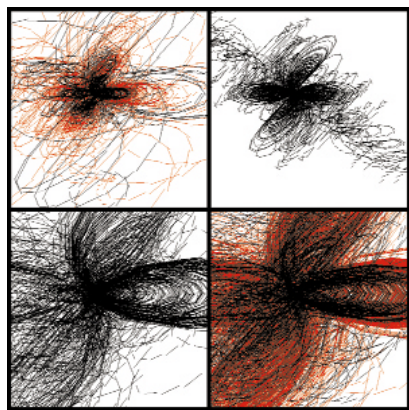

1. Physical appearance (a shape, a pattern, a representation of a mathematical algorithm, or a color).
2. Individual behavior or dynamism, some transformation over time (a shift in color, position, shape, or other attribute).
3. Behavior in its interaction with the other lines on the canvas. The line can sense the presence, proximity, and topology of others and respond in pre-determined but unpredictable ways.

Computational expressionism presents a model for drawing that combines higher-level conceptual design with real-time gestural input. It is a two-fold process, at two distinct levels of interaction with the computer. The artist programs the appearance and behavior of computational lines and then draws with these by dragging a mouse or controlling another input device. The process of computational sketching consists of writing the code, compiling the program, running it, and making some sketches on the digital canvas. The drawings thus produce the next iteration of the code, and the algorithms are refined and varied. The iterative bipartite process of programming and drawing generates an individualized algorithmic style for a particular drawing.

The computer differs from other media in that it is both a constructive and a projective medium. Computation invites us to represent, but also to interpret the world around us, through experimentation and abstraction. In this work, the artist develops individual interpretations of the meaning of gestures and constructs tools to represent them in the final composition. This conceptual interpretation of gestures is a very personal part of the process and a key element of the model of drawing with computation.

Because computational lines are no longer literal representations of a gesture, the aesthetic qualities of an algorithm or of a dynamic interaction generally outweigh representational potential. Thus our work in computational drawing has been almost exclusively non-representational. The quality of lines is a strong aesthetic theme, together with the expressive potential of gesture and the textures and tones achieved through a layering of lines. The compositions emerge from the interaction between the algorithms and behaviors of the computational lines and the gestures used in drawing. The resulting aesthetic takes advantage of the medium's unique characteristics (interactivity, dynamism, and emergence) and highlights the tension between repetition and variation, regularity and irregularity, mathematical order, and gestural disorder.

Expressionism




Perspective systems are designed to construct pictures that, when viewed, produce in the trained viewer the experience of depicted objects that match perceivable objects. Our capacities to see are constrained by the perspective system that we use (that is, by our way of depicting what we see). Mathematicians showed that geometries could be based on axioms other than Euclidean. This altered the notion of geometry by revealing its abstract postulational character and removing it from any connection with the structure of empirical intuition. Spaces were constructed on the basis of non-Euclidean axioms, revealing behaviors closer to our sensations than our perception.

This sketch investigates how computers and new media may extend the designer's perception and imagination. It presents a series of experimental mathematical functions that demonstrate some of these models or mappings for a variety of values for the parameters. The functions address geometric mappings as well as numerical models of projection, and their interest lies in the dynamic nature of continuous computer processing (real-time movement).

Some of these experiments investigate the implementation of art or design theories. For example, what would it look like to move inside a cubist world? Other experiments explore the effect of non-Euclidean theories in the exploration of visual systems. For example, how would an inverted perspective representation behave in a hyperbolic world? Most of the experiments can be combined together in rather unusual ways, such as, for example, a cubist world with supremacist shapes seen through a pointillist filter in real-time animation.

The computer system that was developed for this sketch is called "zhapes." It is a Java-based 3D-visualization system that resides at cda.ucla.edu/caad/java/x/formProj2/formB.html, where it can be downloaded for explorations.



Visual

Hyper-3D Paintings in QuickTime VR: Wunderkammer and Hyperaesthesia

The Wunderkammer and Hyperaesthesia interactive panorama series explore culture and consciousness. These are photographic/painted constructions based on the idea that the computer, a "universal solvent," relates to early studio/museum/expositions that integrated art and technology. Composed of interchangeable virtual worlds, these series present a dialogue between the internal psychology of each individual viewer and external physical or social realities, juxtaposing actual-seeming places with archaeological, painterly images. Although the panoramas seem to present a photographic realism, everything has been constructed or revised with digital techniques. Meanings are developed through viewers' responses. The works may be characterized as the extension of painting into four dimensions.

The idea is to create an environmental metaphor that is not real, a resonant framework, world: artist: artwork: viewer.

Antecedents include:

- Early design of user/viewer experience passed through multigenerational digital technology evolving into user interface design.
- Early 3D installation and 3D paintings transformed by the computer into virtual worlds existing as data displayed as patterns of light.
- Early photography and painting conflated by digital processes into digital imaging.
- Historical exhibits – early, mid, late 1980s in the East Village, an outpost for artists experimenting with art and technologies-migrating from physical to WWW space.

Traditional paintings are built up layer by layer, each successive layer superseding and covering what came before. There is a regret for what is lost, a nagging interest in the dynamic permutations of the process, the archaeology of a painting. With digital image creation, an artist can save each stage of the work, and display it as an animation or a series of stills. Extending this temporal process along a spatial axis lets the artist give a viewer the ability to explore and recreate this entire process of creation via a temporal/spatial record of the work in three dimensions using technology such as QuickTime VR.

The ability to view any digital image all the way down to individual pixels lets artists create images on a microscopic level, as well as the photographic hyperrealism available at the conventional, "zoomed out" view. QuickTime VR allows zooming into alternate, microscopic realities, enabling works that can simultaneously contain abstraction, realism, and surrealism.

New technologies tend toward work utilizing collage, montage, and vision in motion. Just as word processing altered writing and editing, collapsing it into one process, and desktop publishing gives every computer user the freedom of the press, digital imaging and authoring transform possibilities for artists. As the technologies of photography and film collided with and influenced modern art at the beginning of the century, digital processes are affecting current artists as they create work that will give shape to the next.

These paintings are presented as QuickTime VR panoramas within Macromedia Director movies that incorporate interactivity and sound. The presentation looks at finished "Hyper-3D" paintings and then demonstrates the imaging processes and the authoring techniques used to create 3D panoramas from 2D imagery.



Wunderkammer series: Initial view of four linked QuickTime VR panoramas.

An interesting and difficult challenge in designing a virtual environment is to attract the user's attention to an object or area with rich interactivity. If they are not guided, users could easily miss these special objects and areas. They would not be fully engaged by the virtual world we designed for them. Disoriented and frustrated, they would miss the magic of the virtual world.

We solve this problem in our environment, called Microworlds, Sirens, and Argonauts (MSA). In MSA, we use the mathematical concept of "attractors," which act like the Sirens in Jason's famous Greek voyage. The user, like an Argonaut, is drawn to special locations by the magnetism of the attractor's "song." The attractor leads the Argonauts to a perfect position, where they can view an interesting perspective of an object or reach the specific part of an object we want users to touch or pass through.

Attractors Characteristics

MSA uses spatialized sound objects to allow the user to find attractor regions. When users enter an area of attraction, they are pulled to a specific position called the final position. To achieve this, the path between the user's entry position and the final position is built on the fly. At this point, users lose navigational control until they reach the final position. Because the path is built on the fly, and the entry point of the attractor region varies between flights, there will never be two identical paths to the final position. Every "trip" is a unique experience.

Interocular distance and sensitivity of the 3D mouse are carefully controlled for each virtual object that uses attractors. Both parameters are modified at the final position. This greatly improves navigation. Otherwise, if the sensitivity is not adjusted, the user's movements in the virtual environment may be either too fast or too slow to move around and explore some objects. If the interocular distance is not controlled, it may be so big with respect to small objects that when users get closer to them, they suffer from diplopia. Attractors solve both problems.

When the user reaches the final position, the attractor is de-activated and will be re-activated when the user passes through a given boundary. In addition, the attractor's magnetic area decreases each time the user visits it until it reaches a minimum size. This avoids the problem of the user being pulled to the same attractor each time.

Attractor Variables


For each attractor we can define the following variables:

- Origin
- Area of attraction and its minimum size
- Border of re-activation
- User's final position
- Final interocular distance
- Final 3D mouse sensitivity
- Spatialized sound area


Attractors let Argonauts navigate freely, explore the virtual 3D space, and discover new emotions and thoughts. They facilitate the user's navigation and interaction within the environment. They help users orient in virtual space by providing navigational sound clues. Users can follow the spatialized sounds and make virtual "musical maps."

Acknowledgements

This project has been partially funded by The Labyrinth Project, a three-year research initiative funded by the Annenberg Center for Communication. The authors wish to acknowledge Vibeke Sorensen for her support and inspiration, Iguana Robotics Inc. for its valuable input to this research, FIGD for the donation of SAS (Spatial Audio Server), Robotiker for the use of their computers and support, WorldToolkit and Vrex for their support.



Attractor scheme



Virtual object

The Mutable Cursor: Using the Cursor as a Descriptive and Directive Device in Digital Interactive Stories

Background

The game industry is ever striving for a more convincing illusion of reality. In Sony's 2 March 1999 announcement of the new PlayStation they ask you to: "Imagine walking into the screen and experiencing a movie in real-time." However, creating highly realistic environments is only a small part of producing this type of experience.

The problem of real-time, first-person, interactive stories has been the subject of much research. Most productions have felt "more like wandering around a movie set than watching a movie.", The problem is that story telling requires an exactness of timing to guide the reactions of the audience. In interactive situations, timing is often dictated by the actions of the player. What previous research has clearly established is that in order for an effective and meaningful story plot to take place, the player needs some subtle guidance that is either invisible or that takes place in the context of the story.

In addition to the problem of guidance, there is also the problem of the protagonist's character. In most adventure-story computer games, the player usually takes the part of the protagonist, whose character is not richly developed, as it would normally be in other media.

The Technique

The cursor is often no more than a clumsy necessity in adventure games. The graphical representation tends to be both out of context visually and inadequate for conveying story information. So in this technique, the cursor is put to use to both communicate information about a story character and direct the player's attention to appropriate places in the environment. It is hoped that this technique of guidance will both enrich the character and provide a form of guidance that will reduce the time a player has to search for important items, clues or information; the task of searching adds little or nothing to an experience.

Firstly, it was decided that the cursor should no longer be a graphical representation, but blend with the scene. So it becomes a way of changing the appearance of what it alights on. The way the cursor changes an object reflects the character's relationship to it. For example, if an object is especially desirable to a character, it might be seen as brightly lit, perhaps glowing slightly. If it is over a large unimportant object such as a wall, the cursor forms a circular area that is brighter than the surrounding environment, so the cursor is always visible. As the cursor moves onto an object, it gradually mutates to the shape of the object and changes its appearance, and as it is moved away, it returns to a circle of brightness.

Secondly, to direct attention with the cursor and describe the character's relationship to the objects in a scene, it was decided to give the cursor a level of attraction or repulsion to the objects. For example, it slips over some and lingers on others. The level of slippage or stickiness can be defined and attached to an object. This can be altered as a character's situation or attitude changes.

The Scene

This technique is presented in a scene entitled "A Change of Image," in which the player inhabits a character out on an important shopping trip. During this experience, it is hoped that without any other information except the bias of the cursor, the player will come to understand an important aspect of the character's personality, as it is revealed during the course of the trip. The models for this scene were created in 3D Studio Max, and the code for the cursor was written in Visual C++.

Future Directions

In 1998, a method of placing the player in the first-person point of view of a particular story character was developed.³ This technique reveals to the player the thoughts in a character's mind as the player examines a scene. This creates game protagonists that have a history and thought processes of their own. The next step will be to combine this with the mutable cursor. A further step will be creation of a scene containing multiple characters that each have their own agendas and inhabit the same space.

References

1. Sony Computer Entertainment Announces the Design of the Next Generation PlayStation System, 2 March 1999.
2. Platt, C. (1995). Interactive entertainment: Who writes it? Who reads it? Who needs it? www.wired.com/wired/archive/3.09/interactive_pr.html
3. Tallyn E., Meech J. Developing the first person point of view: using characters as a sensory lens. Proceedings of SIGGRAPH 98.



Nami is a decentralized community of identical “orbs,” each of which can display a spectrum of colored light, respond to touch, and wirelessly communicate with its neighbors. When a user activates the spread of color within Nami by touching a single orb, the selected orb responds with expression of a new color. The new state is broadcast to neighboring orbs, prompting them to assume the color and forward the message. In this way, waves of colored light move throughout the distributed network and create visible patterns of behavior. There is no “right” or “wrong” way to arrange Nami, as messages are relayed through whatever pathways are made possible by the placement of the individual orbs.¹ The resulting combination of distributed networking, touch interface, and mutable colors of light provides a novel system for graphical animation in physical objects.

Computer Graphics in Sculpture

Computer-generated animation becomes possible when an artist can exercise dynamic control over a large quantity of individual pixels. Although a single changing pixel is uninteresting, increasing numbers of pixels will manifest recognizable patterns. Certain recursive algorithms can cause pixels to express life-like characteristics known as emergent behaviors, reminiscent of the spread of gossip within a community or the ripple of waves across a pond. Despite the powerful implications of graphic art on the computer, screen-based work is limited to what is expressible with a monitor or projection.


Nami is a preliminary attempt to broaden the scope of computer graphics through integration of color animation into a community of physical pixels, or “phyxels.” Phyxels can be imagined as game pieces, building blocks, or other sculptural objects that can express a variety of colors in response to user input and network structure. While the ubiquitous “beige box” format of a computer discourages the union of physical and virtual form, phyxels enable the artist to design both. Like pixels, a phyxel is most interesting when treated as a unit from which an overall structure can be built; but whereas pixels can exhibit only virtual qualities, phyxels have both sculptural form and animated features.

First Implementation of Nami

The introduction of graphical behaviors into a decentralized network of physical objects presented several challenges to the Nami project. Each orb must be able to respond to the user, exhibit the desired behavior, and communicate with other orbs. Capacitive touch-sensing³ seemed to be the most intuitive interface for the Nami orbs, because it enables the user to control an orb’s state by simply cupping the domed top. Each orb has a microcontroller to manage user interactions and communication. A separate microcontroller modulates three LEDs (one each of red, green, and blue) to create different mixtures of colored light. Considerable effort was devoted to creating a gentle and soothing blending of color, so that the orbs would have an organic aesthetic suitable to the wavelike behaviors from whence Nami derives its name. The physical appearance of the orbs is similarly designed to have an undulating form, reflective of the color waves that they carry.

Future Research

In the first iteration of Nami, orbs were designed for wireless communication over a limited distance within a plane. Future research will investigate alternate designs, including connectivity in three dimensions. The current Nami relies on fixed-format messaging for network communication; the next generation will employ “mobile agents” to provide messages that are executable software objects. This architecture will enhance design flexibility and make Nami a useful test bed for visualization of distributed networking algorithms. Several improvements are being made to the infrared communication and touch sensors, and an inductive charging scheme is under development. Forthcoming versions of Nami will also include an electronic artist’s palette, which is intended to enhance user control by providing a method for specific color and pattern selection.



A touch event initiates the spread of colored light throughout a community of orbs. (photo by Robert D. Poor)



Nami at the Tokyo ToyFair, March 1999. (photo by Robert D. Poor)

References

1. Robert Poor. Hyphos - A self-organizing wireless network, Master’s Thesis, MIT Media Laboratory. 1997, www.media.mit.edu/~r/projects/hyphos/
2. Harold Abelson, Thomas F. Knight, Gerald Jay Sussman. Amorphous computing, 1996, www-swiss.ai.mit.edu/~switz/amorphous/index.html
3. Rehmi Post. Personal communications.
4. Nelson Minar, Kwindia Hultman Kramer, Pattie Maes. Cooperating mobile agents for dynamic network routing, Chapter 12, Software Agents for Future Communications Systems, Springer-Verlag, 1999, ISBN 3-540-65578-6.

Nicholson NY/Pequot Interactives

In 1993, Nicholson NY was selected by the Mashantucket Pequot Tribal Nation to design and produce a series of interactive programs for their museum's permanent exhibits. There are six programs, accessible at 23 stations, on topics ranging from the geological history and natural resources of Southern New England to the social and political history of the Mashantucket Pequots. Touchscreens allow users to easily navigate an enormous amount of information in the form of 3D animations, documentary footage, traditional media, and even hand-painted cell animations, presented in broadcast-quality digital video. The interactives are accurate down to the native tree bark depicted in the 3D models, based on archeological data and US Geological Survey maps – realism that might be lost on some users, but not to the scholars who use the museum for research.

The Pequots, driven to the edge of extinction, were without a pictorial reference, so all designs had to be approved by historians as well as the tribe. Each of the programs contains from one to 2.5 hours of content, but they are also designed to provide casual users with informative nuggets of information that will enhance their understanding of the exhibit.

"Explore the Fort" uses 3D modeling to recreate a 17th-century fort discovered on the reservation in 1991. Users follow pre-rendered animation paths to navigate through the fort. Artifacts found at the site can be found in corresponding virtual locations in the model. Analysis and interpretation is provided in documentary video segments.

"Algonquin Languages"

No native speakers of Pequot survive, so this interactive was built around six stories told by native speakers in the surviving languages closest to Pequot. Each story is accompanied by animation and is subtitled in both English and its corresponding native language. The program also includes a vocabulary section, where users can compare cross-referenced words and phrases in four languages.



Interactives

Human recognition and human communications are classified as being either logical or emotional. Passion spaces here are compared with conceptual spaces in the analogy of logos vs. pathos.

Synesthesia Phenomenon

The painter Wassily Kandinsky was said to possess synesthesia. He said that when he viewed some colors, the colors would also become audible. Moreover, the reverse was true. He could see some sounds as colors. In addition, the French poet Arthur Rimbaud created a poem relating images to colors.

Synesthesia (a fusion of stimuli) involves reactions to a combination of senses. One example of this is the phenomenon of hearing a sound and feeling a color (colored hearing), which is a fusion of sound stimuli and visual stimuli. Yellow is usually assigned to the sound of a child's high-toned voice.

Generation of Passion Spaces by Media Exchange

What kind of world do people with synesthesia perceive? We investigated the generation of spaces by imagining their world. Since relationships are thought to be strong among colors, sounds, and feelings, we tested color image creation from the Japanese poems of Ogura Hyakunin Isshu, whose poems are taught in Japanese schools to encourage understanding of emotions or passions, both from the overt contents and between-the-lines meanings. We then focused on an anthology of 100 poems by 100 different poets. However, this study was not about the correspondence of simple sounds and colors, and it is not like the flashback of concrete scenes imagined from the meanings of the words in the poems. Our goal is to generate spaces to reflect the passions we experience.

We generated passion spaces from Japanese poems in two steps:

1. A table was created of the corresponding relationships between phonemes and colors. For example, a vowel was made to correspond to red, another vowel was made to correspond to black, etc. In addition, correspondences were established between differences in brightness for differences in the sounds of voiced consonants. Here, the relationships between phonemes and colors were determined by the senses of the authors themselves.
2. The phonemes that compose 31-syllable Japanese poems were made to correspond to image elements, and a multidimensional image element array was constructed. We used a 2D array so that the image became rectangular and formed a square matrix. The image elements corresponding to the phonemes according to the time series of the reading were colored from the upper left to the right in the 2D array. After that, various types of image processing (low-pass filtering, etc.) were carried out and displayed.

This same method of connecting words to emotional spaces can be applied to music. For example, if a quarter note is the shortest musical note in a musical score, it can be made to correspond to image elements as the shortest part of the score. Color images can easily be made to correspond to musical notes, and even people who cannot read complex musical scores are able to feel the similarities. In this way, it is possible to instantly judge similarities between pieces of music without the need for a performance.



No. 9 poem by Lady Ono no Komachi

*Color of the flower
 Has already faded away,
 While in idle thoughts
 My life passes vainly by,
 As I watch the long rains fall.*

No. 61 poem by Lady Ise no Osuke

*Eight-fold cherry flowers
 That at Nara—ancient seat
 Of our state—have bloomed,
 In our nine-fold palace court
 Shed their sweet perfume today.*

* English translations used by permission of the University of Virginia Library's Japanese Text Initiative: etext.lib.virginia.edu/japanese/hyakunin

Figure 1. Example of color image creation from poems of Ogura Hyakunin Isshu.

Phene-: Creating a Digital Chimera

My research and creative practice explore the intersection among digital, biomedical, and linguistic modes of bodily representation. In my most recent installations, I place my viewers in situations that force them to take on unfamiliar roles to gain new perspectives on the act "performed." In Phene-, I invite the viewer to impersonate a scientist and use tools associated with the laboratory to interact with the art.


In Phene-, the metaphor of the laboratory "gone amuck" pervades the space. Approaching the installation, the viewer hears garbled sounds coming from a microscope. Signs direct the viewer to don gloves and inspect a group of slides. In attempting to make sense of the samples, the viewer turns the knob to focus the microscope. In this process of twisting, the viewer focuses not only the image but also the sound. The slide is "named" by the computer voice as it comes into view.

Behind the microscope, a wall displays items that allude to the process of capturing, cataloguing, and identifying specimens: dirty scalpels, unwashed flasks, dissection trays, etc. Notes and drawings pinned to the wall portray a bizarre array of hybrid organisms. Also, many of the containers reveal surprising contents. For example, several petri dishes appear to contain routine cultures. However, on closer inspection, colonies of bacteria reveal "intelligent" behavior as they mass into simple word forms such as "CAT."

The focal point of the installation lies behind the exterior wall. A sign directs viewers to examine the "chimera" on the opposite side. Rounding the corner and entering a darkened space, the viewer immediately senses moisture and a dank odor. Here lies a large petri dish illuminated by the light from a projector mounted in the ceiling. Moving closer, the viewer notes that the specimen is composed of a rapidly changing animation layered atop a mass of biotic material. The podium that houses the dish is littered with all sorts of tools: eyedroppers, misters, magnifying glasses, etc. A fan and a shelf with two vaporizer units occupy the far corners of the space.

Viewers interact with the chimerical specimen via processes of feeding, watering, magnifying, etc. In crouching to pick up the magnifying glass, the interactor trips a switch that creates sound. The auditory information consists of amplified breathing noises. As the interactor positions the lens atop the specimen, small dots of light trigger an array of animated creatures to immediately materialize beneath the lens. "Watering" the chimerical specimen alters the look of the dish as well: a light misting of the surface causes the digital organisms to multiply so that the entire space darkens as a result.

Phene- lures the visitor into a fecund space resplendent with the mysteries of the cycling of life. Respiration is a key feature of "aliveness." Thus, the installation uses the literal sounds of human breathing as well as mechanical "breath" in the form of the random turning on and off of the fan. The microbial growth in the dish itself was originally cultured from the artist's breath. In Phene-, the biological process of procreation is juxtaposed with the process of artistic creation. The dish represents the heart of the installation, a dark, womb-like arena for the growth of both virtual and microbial life forms.



SaltoArte is the fruit of an ongoing collaboration between the Harvard Design School and the Busch-Reisinger Museum. It was part of the museum's exhibition titled "Multiple Configurations, Presenting the Contemporary Portfolio", 13 May - 1 August 1999. This exhibition focused on four possible ways to display a contemporary art portfolio in its entirety within the constraints of the museum. Each display necessitated a different type of engagement on the part of the visitor.


As part of this effort, the authors collaborated with the curator to develop an interactive display of the portfolio SaltoArte, a multiple edition of visual works by various European and American artists affiliated with Fluxus, a conceptual art movement of the 1960s.

The objective of this collaboration is to enable visitors to virtually handle art, thereby avoiding any harm to the objects. Visitors can select any number of objects, view them individually with details, as well as turn pages, unfold posters, and open envelopes. Using technology such as Java and VRML, we hope to empower museum audiences to activate Fluxus works as their makers' intended.

The concept of the project is based on two related notions:

1. We want to extend the tangible qualities of the artifacts as they exist in time and space. Whereas the traditional language of 2D multimedia helps us convey information about the artifacts, it typically fails to convey a sense of their physical presence and material qualities. To address this issue, we are building a seamless environment where the descriptive and informative aspects at which 2D multimedia environments excel are augmented by the kind of spatial experience afforded by virtual environments.
2. We need to communicate a variety of visual and auditory information that documents the artifacts in their historic and artistic context. This involves the use of textual and aural descriptions linked to related images and references. To seamlessly merge the textual, figurative, and spatial dimensions of the project, we propose a strategy that blurs the distinction between 2D and 3D by intertwining conventions taken from both realms. The artifacts are scattered in 3D space, but they are also laid out as a seemingly 2D image array. They can be accessed either by means of hypertext media conventions (by passing the cursor over keyword links) or by direct manipulation (by clicking on the artifacts themselves).

Subtle manipulations of lighting and viewpoint position are then used to reveal the 3D nature of the scene. 3D conventions include the ability to directly manipulate the orientation of the object and/or the distance between eye and artifact. Maintaining a level of fluidity between conventions involves development of fine-grained control systems, which enable museum visitors to handle the artifacts easily and intuitively with a virtual pointer and without extraneous visual control structures, such as dialog boxes and buttons.



Setup of the Konsum Art.Server

The Konsum Art.Server is a conceptual Web3D artwork. It evaluates the aesthetics of different operating systems as cultural code and syntax. Certain visual and acoustic codes have already trickled into the cultural consciousness, where they have become socially accepted. They can be read and decoded. These systems include, for example, the codes of Joystick Nations, video games, and computer-game culture that we learned as kids in the arcades.

This online world is also a cyber-psychological tool. Different machine aesthetics offer individual users the possibility to identify themselves as (or with) a certain operating system or brand appearance, for example as a blue Unix shell. So the postmodern identity, defracted, fluid with an unsecure kernel, can be brought back to certain kernels and associated with these kernels of the machine code at the moment of login. This is not an attempt to restore the identity as one proofable unit but to give it an appearance for the moment in a playful environment.


The Konsum Art.Server is a browser experiment to connect tech and arts communities. Users can consider abstract VRML objects as their personal, non-humanoid VRML online avatars. The display always shows an actual freeze frame of your login shell: simple, but symbolically effective. These objects really display information in the 3D Web – information that is not displayed under ordinary conditions, but is in the net beyond the Web. Trace route commands, ns-lookups, and scripts imbedded in the VRML can be extended, but basically all commands that you also could make as user of a Linux system can be made available. This may be the linking point between the arts community and the emerging community of X3D, and between developer communities.

Web3D is providing tools for representing 3D data online. In relation to the 3D environment, the question of dynamic self representation becomes urgent. Aesthetic avatars based on concrete online information are another option. By displaying multiple identities (aka a variety of "nixes shells"), from IRIX, Sun OS, Linux, Free/Net BSDees, etc. directly into the 3D environment as aesthetic experiences, the underlying layer of the net is brought back into Web3D, and at the same time a cultural syntax is expressed through the surface of a specific login shell.

SUPERFEM Datavatars - Datasituationism as Interface

In the SUPERFEM performances, VRML Datavatars create situations as a social metadesign to enable users to experience conditions and netparameters in addition to the use of the machine prosthesis.

The structural aspect of the Internet is displayed. All navigation through objects is generated in real-time. By entering 3D objects, different textures of your actual movements on the net (time spent online), or log files of your chosen destinations, are displayed. Time is also displayed by sound samples that express the volume and speed/slowness of data flow.



Stereo sound has been the state of the art for musical recordings for more than 30 years. Consumer media like tapes, records, and CDs attempt sound fidelity by using two channels for audio playback. However, the illusion of "being there" is limited via only two channels.

This proposed application takes advantage of 3D graphical and acoustical output. It follows the idea of a new interactive, immersive, and semi-artificial way of experiencing music. All parameters, like location and orientation of the listener and sources as well as the visual and acoustical appearance of the virtual environment, can be modified using a 3D application (see Figure 1).

Required data files like scene descriptions, images, or live video of the musicians, as well as their audio tracks, are stored on a CD-ROM or a DVD. The prototype implementation consists of graphical and acoustical subsystems. The network SAS (Spatial Audio Server) has been developed for acoustical rendering of multiple audio sources in 3D environments. Thus, events can be invoked, like changing the acoustical environment (room acoustics), playing back sources, or updating their positions (see Figure 2).

In order to place sound sources in virtual environments, it is necessary to record non-reverberant signals of the sound material. Another important requirement is the acoustical separation of the different sources. Pop and rock music is usually produced by creating multitrack recordings. For acquisition of music material for a classical interactive production, a different approach has to be followed: In the prototype "VirtueString," the four instruments of a string quartet (Haydn, Op. 20 No. 5) were recorded individually in an anechoic chamber to get total separation without any room acoustics. In order to get synchronized material, first a stereo prerecording of the whole quartet was produced in the conventional way. This was later used as a pilot track for the single-instrument recordings in the anechoic chamber.

For the consumer market, a CD or DVD implementation of the software provides low-cost access to a virtual interactive music environment. Another marketing option is a variable "music-minus-one" playback system. "Minus-one" means half-playback recording of a certain music piece where one voice is muted. This allows musicians or music students to play with a complete virtual symphony orchestra.

The application is currently enhanced by an audio-visual environment with three different rooms of various sizes, geometry, and acoustical properties. Further investigations focus on how to achieve an adequate spatial mapping between graphics and sound under various geometric and technical constraints.




Figure 1
The audio-visual presentation is rendered in real-time.

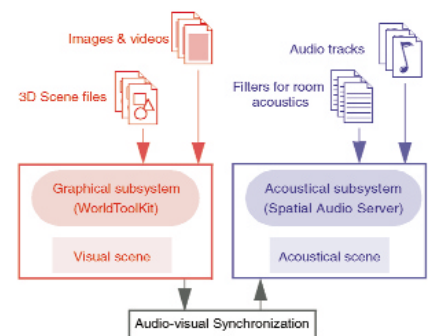


Figure 2
Graphics and sound are processed by two linked subsystems.

Music

VisiPhone

VisiPhone is a communication object that opens a graphical as well as an audio portal between two spaces. It is designed to provide a continuous, ubiquitous connection between these two different locations. The goal is to create an aesthetically intriguing object that enables users to perceive conversational patterns that are present but not obvious in traditional communication interfaces. Through this experimental medium, we are exploring the social and aesthetic aspects of sound visualization.

Function

The VisiPhone project began with the notion of building a virtual portal between two domestic spaces, one that would allow the inhabitants of two separate homes to communicate easily and to be intuitively aware of each other's presence. VisiPhone's graphics express the dynamics of the conversations originating at both locations by providing visual feedback that one's voice has carried sufficiently and indicating the presence of those on the other end. It portrays the existence of the connection even in moments of silence, thus removing the surveillance-like aura of the audio-only system.

Today, most communication devices such as the telephone or the computer are fairly utilitarian in design. Yet as networked communication and permanent connectivity become the norm, we believe that their capabilities will spread to a variety of interfaces and become integrated into a variety of objects and environments. VisiPhone is an exploration of the aesthetic and social possibilities we can achieve by melding computer graphics and communications.

Form

For a piece such as VisiPhone, form is function. It must be attractive and intriguing enough to claim a central place in a kitchen or a living room.

We are currently experimenting with two versions of the display unit: a dome and a podium pedestal. The dome shape breaks away from the flatness of current digital displays; its placement on the pedestal puts it into the category of sculptural object (and conceals the projector). The dome version is a smaller, more intimate interface whereas the podium is a larger unit meant to be used by a group of people. On both displays, the graphics are designed to convey a sense of rhythm and activity.

The abstract graphics rendered on the display surface portray the audio originating at each end of the connection. The source of the audio is revealed through the use of color (blue tones for incoming audio and orange tones for outgoing audio). Volume is mapped onto the size and color of the graphics: the louder the volume the bigger and more saturated the circle.



Whenever speakers at different ends of the connection talk at the same time, the graphics for both audio signals are rendered as overlapping colored circles. The different colors that represent each speaker blend, portraying the mixing of the audio signal.

VisiPhone displays both the real-time audio and a brief history of the conversation. The real-time rendering is shown in the center of the display and then spirals out toward the edge of the display. The circles fade as they approach the end of the spiral, enhancing the notion of evolving time.

Technical Description

A two-way VisiPhone interaction requires two VisiPhone interfaces. Although the dome and the podium differ in their physical appearances, they share the same computing and projecting hardware. Each VisiPhone consists of a set of speakers with embedded microphone, a PC with full-duplex audio card and GL compatible video card, a scan-converter, a projector, and a projection display.

Audio input comes from microphones; the input signal is transformed into a video and audio composite. The output is a graphical display and a set of attached speakers. The video from each corresponding PC is rear-projected onto the display surface, thereby transforming the VisiPhone apparatus into a dynamic interactive object.



3D Facial Reconstruction and Visualization of Ancient Egyptian Mummies Using Spiral CT Data

Maurizio Forte
Consiglio Nazionale delle Ricerche
forte@nserv.icmat.mlib.cnr.it


Giuseppe Attardi
Marilina Betru
Silvano Imboden
Francesco Mallegni
Università di Pisa

Roberto Gori
Antonella Guidazzoli
CINECA-VISIT

The problem of rebuilding a face from human remains has been, until now, especially relevant in the ambit of forensic sciences, where it is obviously oriented toward identification of otherwise unrecognizable corpses, but its potential interest to archaeologists and anthropologists is not negligible. We present here the preliminary results of joint research by the Università di Pisa, the Visualization Laboratory of CINECA, and the CNR-ITABC (Institute of Technologies Applied to Cultural Heritage, National Research Council, Rome) whose aim is reconstructing, through spiral computed tomography data and virtual modeling techniques (in our case with VTK software), 3D models of the possible physiognomy of ancient Egyptian mummies. This work is carried out through a multidisciplinary approach, involving several specialties: image processing, anthropology, Egyptology, and computing archaeology.

Main project tasks are:

1. Anthropological and Egyptological analysis of the head.
2. Spiral CT of the head.
3. Reconstruction of a 3D model of the skull generated from CT data processing.
4. Reconstruction of soft tissues.
5. Application of textures fitting the somatic features.



3D facial reconstruction and visualization of ancient Egyptian mummies

Visualization

3D Gait Reconstruction Using Two-Camera Markerless Video

Gait analysis is a valuable tool in evaluation of neuromuscular disorders, athletic injuries, and prosthetic joint replacements. Our goal is to reduce deployment of technology in acquisition of video used for gait analysis without sacrificing accuracy. We envision the following application context for this research: A patient with an abnormal gait has his or her leg measured and a typical gait sequence digitized at a central clinic. The patient then undergoes surgical correction and begins post-operative therapy with a local physician. After some time, that practitioner sends a video of the patient's walk to the central clinic for processing. Before-and-after sequences may reveal improvement in the gait and help determine further treatment.

Unsynchronized markerless color video sequences were taken at the OSU Gait Lab from the front and side of a person walking, and divided into frames. The subject wore form-fitting leg tights of different colors.



Leg contours and axis, skeleton, and NURB surface

Leg Silhouette Contour Extraction

The side view images were deinterlaced and then expanded, yielding two field images per frame image. Edge detection was performed on a gray-scale version of each image using a combined Sobel operator. Strong edge pixels on each scan line of the images were used as starting points for a directed search for leg segments in the color images. Segment extraction used a simple color predicate, since different colored leg coverings were worn. Pixels satisfying the color predicate for each leg were grouped into a run, and the endpoints of the run were labeled as silhouette points. The resulting well-connected set of points was smoothed using a B-spline fit (see Figure). Points on the leg axis were assumed to lie on the midpoints of the scan lines belonging to the segments.

Hidden Surface Reconstruction

In the side view, parts of one leg that were occluded by the other leg were recovered in two ways. If, on a given scan line, one of the contours of the leg was found, the other contour point was estimated from a template that contained the width of the legs in any scan line. If both contours for a leg were missing, the positions of the missing contour points were estimated using the closest known contour points above and below them.

Knee and Ankle Location

Each leg axis in the side views was searched for knee and ankle "points," which divided the axis into three contiguous parts such that the independent linear fits applied to each part produced minimum squared error. Knee positions with larger angles between the upper-leg and lower-leg segments were used to predict the knee positions in less-reliable frames. Knee and ankle positions in front views were estimated using the corresponding side view measurements.

3D Skeleton and NURB Surface Reconstruction

Using the camera configuration geometry, the 3D positions of the knee, ankle, and leg-top were obtained as the intersection, or point of nearest approach, of the projection rays defined by the camera centers and the estimated anatomical position in the front and side views. The resulting pair of 3D line axes for each leg was used to reconstruct skeleton and NURB surfaces (see Figure).

Conclusion

The knee angles in each frame of the reconstructed motion, and in the actual 3D motion obtained from a six-view simultaneous synchronized marker video were compared. The reconstructed motion followed the actual 3D motion over time but produced a much reduced bend magnitude. Future work will include improving the accuracy and robustness of leg contour extraction and using anatomical data to determine the leg axes.

Acknowledgments

The OSU Interdisciplinary Seed Grant Program provided initial funding. Thanks to the OSU Gait Lab for their facilities and personnel, and the OSU Department of Computer and Information Science for infrastructure support.

References

1. D. DeCarlo et al, Deformable Model-Based Shape and Motion Analysis from images Using Motion Residual Error, Proceedings ICCV98, pp. 113-119, 1998.
2. S. Kim et al, Two-dimensional Analysis and Prediction of Human Knee Joint, Biomedical Sciences Instrumentation, 29, pp. 33-46, 1993.

Non-contact measurement techniques like those based on structured light have found widespread use in industrial metrology and reverse engineering. Laser-based technologies attempt to structure the environment, through artificial light projection. The results are dense range maps extracted from visible surfaces that are rather featureless to the naked eye.

In the field of heritage recording, onsite and complete 3D documentation is necessary in situations where objects and/or environments can't be moved or their access is restricted because of natural causes (for example, earthquake or tourism-related degradation). In this situation, 3D acquisition must be performed in a rapid-response timeframe, from several hours to maybe a day. In addition, many applications do not allow any alterations to the object and surroundings to suit the vision system (for example, by placing markings or changing the reflectivity of the surface with powder or abrasion).

We have put together a self-contained 3D imaging system that is capable of satisfying many demanding onsite 3D documentation tasks. Special attention is placed on accuracy, as it is critical for obtaining high-quality reconstruction of 3D models from range imagery.

Compact 3D Range Camera

The range camera (Biris, Figure 1) was developed to work in difficult environments where reliability, robustness, and ease of maintenance are as important as accuracy (see URLs). The main components of Biris are a solid-state laser diode, a CCD or CMOS camera, a camera lens, and a mask with two small apertures. The mask replaces the iris of the camera lens (Bi-iris). This imaging technique, when combined with advanced signal processing algorithms, allows Biris to become very tolerant to ambient light illumination (sunlight, for example). The Biris camera is mounted on a pan unit and tripod, and is connected to a PCI card installed in a portable PC. Reconstruction of a 3D object is achieved by acquiring a sufficient number of range images to cover the surface of that object. The views must have enough overlap between them to find their relative spatial position and orientation.

Results from Remote Sites in Italy, 1997-1998

Figures 1 and 2 show a number of 3D models acquired in Italy over a two-year period. The models include a marble sculpture of the Madonna col Bambino by G. Pisano, a bronze bas-relief from Donatello, and some architectural elements from the old Abbey of Pomposa. The range images for these projects were acquired in three separate one-day sessions. The models were generated using Polyworks on a PC. A number of sites have been documented using either Biris or National Research Council Canada's other 3D technologies. The sites, in Israel, Florence, NASA facilities, Boston, and Canada, are described at these URLs:

www.vit.iit.nrc.ca/blais/bi_home.html
www.vit.iit.nrc.ca/beraldin/Web_italia/Biris_in_ITALY.html
www.vit.iit.nrc.ca/Pages_Html/English/Research.html
www.vit.iit.nrc.ca/Pages_Html/English/Applications.html
www.innovmetric.com/

Rapid 3D



Figure 1
 Top and center: portable 3D imaging system includes a Biris camera (standoff: 30 cm, range: 200 cm, data rate: 15,360 3D samples per sec, weight: 900 gr.), a pan unit and a tripod. Bottom: outdoor application at the Abbey of Pomposa (circa 850) in Italy 1998.

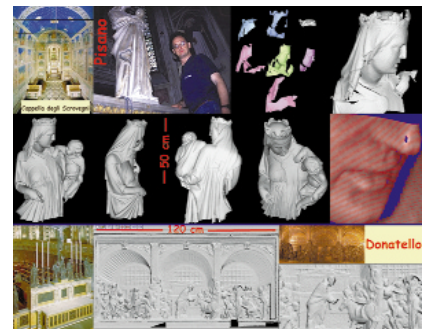


Figure 2
 Top and center: Madonna col Bambino, G. Pisano (circa 1305) showing the seven distinct color-coded images used for the face and model of the top portion using 62 views for a total of 1,340,000 polygons. Bottom: bronze bas-relief of Donatello (circa 1446-1450), model with about 580,000 polygons, Italy 1997.




Figure 1 (a)




Figure 1 (b)




Figure 1 (c)



Figure 2 (a)




Figure 2 (b)



Figure 2 (c)

3D Physics-Based Brush Model for Painting

Though a great deal of free and commercial painting software already exists, we feel that the current software does not yet allow users to draw free and varied brush strokes. In the real world, a painter can express various strokes without “mode changes,” using only one brush.

One major problem is that brushes in current major painting tools are based on the trace of a disc,¹ which limits the painter’s ability to control the stroke shape. Other approaches for producing more natural-looking strokes exist. However, techniques that widen the center curve^{2,3} have limitations similar to those mentioned above for popular paint programs. Morphing with a painted texture can express rich strokes, but the method is not intuitively controllable.⁴

Modeling

As an artist paints, the hair of the paint brush bends and changes shape, and so does its footprint. These changes allow various rich and expressive strokes. To represent natural strokes on a computer, we designed a 3D brush hair model.

The physical model is shown in Figure 1(a). When the position and posture of the brush handle are given, the joint angle is derived by calculating the minimum energy of the tuft model. This energy is a summation of the bend energy of the joint, the potential and kinetic energy of the tuft mass, and the frictional energy between the tuft mass and the paper.

The curve of the tuft spine is defined as a Bézier curve with three control points in Figure 1(b). Because the locations of the control points are determined from direct input of the brush stick posture and simulation results, the shape of the spine curve is reasonably natural and can be controlled easily.

Finally, the tuft itself is defined as a set of discs whose centers are on the spine curve. The footprint of the tuft is calculated by detection of the intersection of the discs and the xy plane in Figure 1(c). The drip shape of the footprint changes size and also bends according to stick motion. In order to apply paint to paper, the model also contains a simple liquid movement model.

Painting Results

The images in Figure 2 were painted with our brush model. A pen-type input device was used for stroke input. The paint model used is based on Kubelka-Munk theory.⁵ Rendering is completed in real-time on a 450MHz Pentium-based Linux system and is fully interactive.

Figure 2(a) consists of only four strokes. However, each stroke is very expressive, and together they succeed in creating a meaningful artistic picture. In Figure 2(b), the sharp curve at the bottom right of the second character shows off the features of our brush model well. When the motion of the brush handle curves, the hair tip traces a different route than the handle. This effect is quite difficult to create with general disc tracing methods. Figure 2(c) is an oriental style flower painting. It shows various strokes and novel expression as a computer graphics image.

Conclusion

The 3D physics-based brush model widens the possible expression of stroke inputs. Because the hair tuft of the brush changes dynamically, stroke starting point, ending point, and curves are very expressive. Most importantly, the brush allows natural input in a manner very similar to real painting with a pen-type input device.

References

1. Painter, www.metacreations.com/products
2. Steve Strassmann. Hairy Brushes, *Computer Graphics*, 20(4), pp 225-232. 1986.
3. Oinglian Guo. Generating Realistic Calligraphy Words, *IEICE*, E78A(11), pp 1556-1558, November 1996.
4. Siu Chi Hsu and Irene H.H.Lee. Drawing and Animation Using Skeletal Strokes, *SIGGRAPH 94 Proceedings*, pp 109-118, 1994.
5. Von Paul Kubelka and Franz Munka, *Ein Beitrag Zur Optik der Farbanstriche*, *Zeitschrift Fur Technische Physik*, pp 593-601. 1931.

Absence of shading and texture is a problem of traditionally animated film. It inspired soul-searching and strenuous effort well before shaded, textured 3D computer animation arose. An interesting approach to this problem is to use the 2D shapes drawn by the animator to define 3D surfaces, which can be rendered by conventional 3D graphics algorithms. This method was outlined in a previous paper,¹ which describes an "inflation" algorithm for creating smooth 3D surfaces from 2D shapes.

These "inflations" are a necessary component of the automatic airbrushing, specular ornaments, and gobo shadow effects to be described. It is worth noting, though outside the scope of this presentation, that such surfaces arise in many contexts. Implicitization of a 2D matte can be performed using the signed distance function, but smoother implicitizations are possible in 2D and 3D.² Interiors of such functions are smooth "inflations" of the shapes. These surfaces, derived from drawings, are very useful for achieving a wide range of continuous-tone animated effects. They are not, however, accurate models of the characters the drawings depict.

To achieve the fidelity to the props and characters required for such effects as texture mapping, accurate 3D models are necessary. Previous efforts in this direction³ require animators both to draw a sequence and animate a corresponding sequence in 3D. By tracking drawings and conforming models to them, we eliminate the 3D animation step. Instead, the 3D animation defined by the model and the drawings can be further edited in 3D after the tracking is performed.



Specular reflections on the ornaments are a reflection mapping of the torch flame (from "Prince of Egypt").



3D dress model automatically tracks and conforms to the drawing of the dress, permitting texture-mapping.

Effects


References

1. Lance Williams. Shading in Two Dimensions, Graphics Interface '91, pp.143-151, Morgan-Kaufman Publishers, June 3-7, 1991.
2. Luiz Velho and Jonas Gomes. Approximate Conversion of Parametric to Implicit Surfaces, Computer Graphics Forum, vol. 15, #5, pp. 327-388, Elsevier Science Publishers, 1996.
3. W.T. Correa, R.J. Jensen, C.E. Thayer, and A. Finkelstein. Texture Mapping for Cel Animation, Computer Graphics Proceedings, Annual Conference Series, 1998, ACM SIGGRAPH, pp.435-446.

A Level-Set Approach for the Metamorphosis of Solid Models

We present a new approach to 3D model morphing that employs an active deformable surface that starts at the source shape and smoothly changes into the target shape. We represent the deformable surface as a level set (iso-surface) of a discretely sampled scalar function of three dimensions.

The sampling produces a regularly spaced rectilinear volume dataset. Such "level-set" models have been shown to mimic conventional parametric deformable surface models by encoding surface movements as changes in the grayscale values of the volume. A well-founded mathematical structure leads to a set of procedures that describe how voxel values can be manipulated to create deformations in the level sets. The result is a voxel-based modeling technology that offers several advantages over previous methods, including support for a wide range of user input, no need for reparameterization, flexible topology changes, and results with sub-voxel accuracy.



As with previous volumetric model morphing algorithms, our approach consists of two central steps: a "warping" step that utilizes a user-defined coordinate transformation that deforms the source model into approximately the same shape as the target model, followed by a "blending" step that completes the morph by automatically moving the level-set model towards the target model's surface. Since the warping stage of volumetric morphing has been satisfactorily studied by others, our work primarily focuses on developing a superior method for the blending stage, one which does not require significant user input in order to produce an acceptable morphing result. Our level-set approach will produce a reasonable morph with no user input. Users may then apply an increasing amount of input until the desired morphing sequence is achieved.

In general, an animator controls the morph by defining how the models overlap. The level-set model only moves in those areas where the target and source surfaces are separated from each other. The portion of the level-set model that is outside the destination model will shrink, while the portion inside the destination model will expand to become the final shape. As long as the models initially overlap, the source model is guaranteed to smoothly deform into the target model.

The level-set approach to 3D model morphing provides several advantages over other 3D morphing algorithms. Level-set models are active; their underlying motion is defined algorithmically and not defined interactively. The amount of user input required to produce a morph is directly proportional to the amount of control the animator wishes to impose on the process. The animator may allow the system to automatically generate the morph or employ a full 3D warping to describe the morph, with the level sets simply providing a small fine-tuning of the surface. Any intermediate amount of user input will produce a reasonable morph, as long as the two models overlap. Since level-set models do not rely on any kind of parameterization, they do not suffer the problems of parametric surfaces (for example, limits on their shape and the need for reparameterization after large shape or topological changes). This lack of parameterization, along with no direct representation of topological structure, allow level-set models to easily change topology while morphing.

The model can "split" into pieces to form multiple objects. Conversely, several disjoint objects may come together to make a single object. Finally, both the scan conversion and deformation stages of our morphing approach produce models with user-defined sub-voxel accuracy. This allows the user to specify the time/quality trade-off, and generate superior results as compared to previous work.

Deafness is not only a barrier of sound but also a barrier of language. American Sign Language (ASL) is a natural language used by members of the North American deaf community and is the third or fourth most widely used language in the United States. While ASL shares vocabulary with English, there is no simple word-for-word translation. Because of the differences between the two languages, most native ASL signers read English at the third- or fourth-grade level. This is why closed captioning on television is a good first effort at making spoken English more accessible to the deaf population but does not represent a completely satisfactory solution.

At present, deaf people rely on sign language interpreters for access to English, which is an awkward solution at best. A digital sign language interpreter, which translates spoken English into ASL, would help bridge the gulf between deaf and hearing worlds. We are currently developing a database of ASL that will be used as the lexical database for machine translation. It includes position, orientation, and shape of the hands as well as motions that comprise a sign. Just as important as the raw geometric information is the linguistic information. Certain geometric aspects of a sign may be changed based on its usage, and without linguistic information it is impossible to synthesize grammatically correct sentences. Thus, the geometric information is somewhat similar to a keyframe, but certain aspects must be treated as parameters during synthesis.

Problem

The largest task is gathering the data once the database scheme has been created. Gathering data involves transcribing ASL, and the people who know it best are ASL signers. While motion capture initially presents an appealing option, it cannot record the linguistic aspects of a sign. It is analogous to scanning in a printed sentence and attempting to use the raw bitmap instead of extracting the characters. Further, the motion and position of a sign may be modified depending on its context. However, using a general animation package to transcribe a sign also poses problems. Learning such a package requires a significant amount of time. Few members of the deaf community are willing to invest a large amount of time in training before beginning the transcription process.

Our Approach

To reduce the large time investment, we have customized an animation system (3DS Max) to accommodate ASL. The controls are couched in terms that are more familiar to ASL signers than such conventional animation operations as rotations and translations. The hierarchical system allows a user to specify a sign in terms of hand shape, location, and orientation. Figure 1 is a screen capture of the subsystem that specifies hand shape. Data encoded by users are stored and recalled from a Microsoft Access database. A Visual Basic ActiveX control communicates between the graphics package and the database.

Initial usability studies have been quite promising. On average, it takes 10 minutes for native ASL signers to learn enough about the system to input hand shapes, and transcribing a handshape takes an average of 82 seconds.




Figure 1

The interesting work by Haumann & Hodgins¹ addresses hovering flight and concentrates on the control aspect of flight. In our aerodynamic model for the forward flight of a bird with two wing segments, the emphasis is more on simulating the bird's undulating trajectory and the pitching motions of the bird's body due to the wing motions rather than simulating the wing motion itself.

We generate the wing shape and motions such as beating and gliding by judiciously varying the sweep, dihedral, and twist as functions of time.² The aerodynamic model can be applied to any other scheme that generates these as functions of time.

Based on wing shape, orientation, velocity, and motions, we calculate the aerodynamic force acting on the wing. The force distribution on the wings is predicted by an aerodynamic model that employs a combination of analytical and numerical techniques. The trajectory of the bird is then determined using principles of flight mechanics.³

The aerodynamic model is applicable to many flight phases and can be employed to animate imaginary characters as well. Interactive computational speeds are achieved due to assumptions of negligible induced velocity and quasi-steadiness.

We can also control the trajectory of the bird to pass through specified targets by direct or indirect control of the pitching moment. Direct control of pitching moment involves specification of the pitching moment of the bird through a p-d controller, which alters the body axis angle. The body axis angle alters the average angle of attack on the wing, thereby controlling the trajectory of the bird. The indirect specification of the pitching moment involves driving a control variable such as the wing sweep of the secondaries by means of a p-d controller. Changes in the control parameter alter the pitching moment in such a way as to propel the bird toward the target. While the direct approach is less costly, the pitching motions of the bird are not correlated to any observable wing motions. The indirect approach rectifies this deficiency, but it is about thrice as costly as the direct approach.

Our Web site contains a technical report⁴ that presents the details of the methodologies and three animation sequences that illustrate the methodologies. A video sequence of a flying character utilizing the indirect approach to generate the flight trajectory was presented in the SIBGRAPI 98 Computer Animation Festival.⁵ An image from this sequence is shown in the figure.

The main advantage of our approach lies in the interactive speed of the models, with which the effects of changes in parameters can be assessed.

Acknowledgements

The authors thank S.K. Lim, L.L. Goh, and B.H. Low for animating the flight sequences.

www2.cg.it.ntu.edu.sg:8000/~iagwww/BirdFlight/index.html



Figure 1

References

1. D. Haumann and J.K. Hodgins. The Control of Hovering Flight for Computer Animation, Nadia Magnenat Thalmann and D. Thalmann, editors, Creating and Animating the Virtual World, Computer Animation Series, pages 3-19, Springer-Verlag, 1992.
2. A. Azuma. The Biokinetics of Flying and Swimming, Springer-Verlag, Tokyo, 1992.
3. C.D. Perkins and R.E. Hage. Airplane Performance, Stability and Control, Wiley, New York, 1949.
4. B. Ramakrishnananda and K.C. Wong. Animating Bird Flight Using Aerodynamics, Technical Report CGIT-IAG-TR98-4, Centre for Graphics and Imaging Technology. Nanyang Technological University, Singapore, Aug 1998.
5. K.C. Wong, S.K. Lim, L.L. Goh, B.H. Low and B. Ramakrishnananda. "Litter Bug," In SIBGRAPI '98 Video Festival, 1998 International Symposium on Computer Graphics, Image Processing and Vision, Rio de Janeiro, Brazil, Oct 1998.

Flight

Current methods for figure and character animation involve a tradeoff between the level of realism captured in the movements and the ease of generating the animations. At one end of the spectrum, computer-animated features and movie special effects display extremely realistic individuals with personalized movement characteristics. However, they require teams of animators using labor-intensive systems where every movement must be explicitly specified. At the other end of the spectrum, behavioral animation systems can automatically generate multiple characters that interact with each other and their environments by merely specifying a set of initial conditions, but the various characters often move in a fairly mechanical manner and share very similar motions. We introduce a motion-control paradigm that combines the advantages of both: the ability to specify a wide range of natural-looking movements with minimal user overhead.

Our approach involves simplifying generation of expressive movements by proceduralizing the qualitative aspects of movement to hide its non-intuitive, quantitative aspects and provide the user with easy-to-use, interactive control through textual and spatial descriptors.^{1,2} To do this, we needed a language for describing expressive movements and a translation of that language into quantitative movement parameters.

We found that the Effort component of Laban Movement Analysis^{2,3,4} met our requirements as a means for describing movements, has a solid foundation based on observation and analysis of humans performing a wide range of movements, and is being used in a growing number of fields. Effort is the part of Rudolf Laban's theories on movement that describes the qualitative aspects of movement using textual descriptors along four motion factors: space, weight, time, and flow.

The extremes of each motion factor give the eight Effort Elements: indirect, direct, light, strong, sustained, sudden, free, and bound.

We built an empirical model of Effort using kinematic movement parameters. The movement parameters fall into three general categories: those that affect the arm trajectory, those that affect timing, and flourishes. The arm trajectory is defined by two parameters: path curvature and interpolation space. The path curvature determines the straightness or roundness of path segments between keypoints. The interpolation is performed in end-effector position, joint angle, or elbow position space. The timing of movements is controlled by the number of frames between keypoints and parameters to a timing function (inflection time, time exponent, start, and end velocities), which define velocity curves, anticipation, and overshoot. Flourishes are miscellaneous parameters that add to the expressiveness of the movements and include: squash and stretch, breath, wrist bend, arm twist, displacement magnitude, and limb volume.

We describe EMOTE (Expressive MOTion Engine), a 3D animation control module that demonstrates this approach. EMOTE uses inverse kinematics with spatial movement requirements specified through end-effector positions. It lets a user express the essence of the movement by setting the four Effort motion factors. EMOTE uses these Effort settings to determine values for the kinematic movement parameters, specifies an animation that follows the defined position sequence, and displays the selected Effort qualities.

Since our translation process is computationally inexpensive, EMOTE provides interactive editing and real-time motion generation. EMOTE produces a wide range of expressive movements with an easy-to-use interface that is more intuitive than joint angle interpolation curves or physical parameters. Preliminary evaluations have shown that both experienced and first-time EMOTE users can create simple expressive character animations in a shorter amount of time at a quality that is the same or better than animations built using a popular commercial animation package. Embedded in a real-time simulation system, EMOTE provides procedural control for editing basic movements based on a character's motivation and intent.

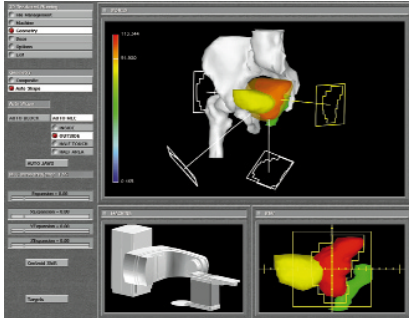
References

1. D. Chi. A Motion Control Scheme for Animating Expressive Arm Movements, PhD thesis, University of Pennsylvania, 1999.
2. I. Bartenieff and D. Lewis. *Body Movement: Coping With the Environment*, Gordon and Breach Science Publishers, New York, 1980.
3. C. Dell. *A Primer for Movement Description: Using Effort-Shape and Supplementary Concepts*, Dance Notation Bureau, Inc., New York, 1970.
4. C. L. Moore and K. Yamamoto. *Beyond Words: Movement Observation and Analysis*, Gordon and Breach Science Publishers, New York, 1988.

Application of Computer Graphics for Design and Delivery of Conformal Radiation Therapy


Radiation therapy, one of the most widely used forms of cancer treatment, affects the lives of over 600,000 patients annually. The challenge of planning radiotherapy treatments is to tailor delivery of radiation to conform to the tumor shape and avoid as much healthy tissue as possible. Visualization techniques exploiting 3D computer graphics and medical image data are now used to help accomplish this. By manipulating patient-specific 3D anatomic data and models, the clinician can better appreciate the complex spatial relationships among tumor, healthy tissues, and beam trajectory.

Visualization techniques are also used to help guide the delivery of the planned treatments. Together, such techniques permit more precise planning and delivery of beam arrangements, shapes, and strengths, making it possible to safely escalate radiation doses in many tumor sites well beyond previously defined levels of "tolerance" and achieve improved local control and better cure rates.



Visualization for radiation therapy planning. a) Beam trajectories and dose isosurface (transparent orange) relative to surface models of patient anatomy. b) Beam's-eye view of the target (red) and normal tissues (bladder in yellow and rectum in green). c) Machine model to help the planner detect impractical beam orientations.

232



Visualization for treatment delivery verification. a) A workstation at the treatment unit displays a realistic model of the machine and its components. A solid-surface representation of the machine b) and the beam-shaping collimator c) show actual machine positions, and the wireframe models indicate the intended move.

Therapy

Impulse-based algorithms for rigid body simulation consist of an outer loop that repeatedly steps forward in time, integrates trajectories, and identifies and processes collisions.¹ Algorithms use various methods to choose time-steps and detect and resolve collisions. Current algorithms update all objects at every step, so that:

- Individual objects are updated more frequently than necessary. An object that is far from any other will be updated at the same frequency as a set of closely packed objects.
- State integration and bounding box computation must be performed for every object at every step. Global data structures (for example, records of object overlaps) must be updated at every step.

We have developed an event-driven simulation algorithm that addresses these inefficiencies. Our algorithm considers different objects at different times (it is asynchronous) and adaptively adjusts the time at which a given object is considered based on its recent history and expected future behavior (unlike the method developed by Kim, Guibas, and Shin,² which cannot adapt and requires exact collision detection).

The algorithm stores state (position, velocity etc., and time) and an associated bounding volume for each object. Global data structures are: a priority queue for events, sorted lists of bounding box extents for each dimension, and a set of hash tables storing bounding box overlaps in 3D and each dimension.

There are three types of events: start, update, and pairwise. The simulation begins by inserting a start event for each moving object. An outer loop then repeatedly deletes the earliest event in the queue and processes it. The core operation is the Predict procedure, which predicts what may happen to a given object within a time interval. A spatio-temporal bounding volume that contains the object throughout the time interval is computed, and its correct position in the sorted lists is located (starting at its former position). During the re-sort, overlapping volumes and hence potential collisions are detected, and pairwise events are scheduled. Events are processed as follows:

Start: The Predict procedure is called with the object and a small initial time interval. The object's first update event is then scheduled for the end of the interval.

Update: The object's state is integrated up to the event time, and the Predict procedure is called with an interval of twice the time since the last event. Another update event is scheduled for the end of the interval.

Pairwise: The two objects involved are integrated to the current time then checked for collision. If not colliding, another pairwise event is scheduled based on a new predicted collision time. If colliding, the collision is resolved, and the update event process above is applied to each of the two objects.

Running time is determined by the total number of events and the cost of each event. The speedup we obtain depends on the simulation, but general observations indicate that the maximum possible speedup is proportional to n , the number of objects in the simulation, and experimental evidence demonstrates that this speedup is attainable. Both our algorithm and synchronous algorithms are driven by the need to perform collision tests, so the number of iterations of the event processing loop will be similar. Our algorithm processes additional update events, but adapts to keep the number low.

The advantage of the algorithm comes from a reduced cost for each event:

- Collision processing cost: will be the same for each algorithm, because both process the same sequence of collisions.
- Object integration cost: our asynchronous algorithm only integrates one or two objects per step, but the integration step-lengths are generally longer. The maximum available speedup is proportional to n .
- Overlap processing cost: Our asynchronous algorithm updates the bounding volume for at most two objects, rather than all n , and the re-sort process takes time proportional to the number of overlap changes for those objects, which is likely to be a small constant because the size of the bound used to find overlaps is adapted. A speedup of n results.

We compared the performance of our algorithm and a generic synchronous algorithm on a traditional example: simulating a number of cubes in a fixed-size box with each cube started in a random state. Each simulation ran for 500 seconds of simulated time with objects under the influence of gravity but frictionless, perfectly elastic collisions. Three different sets of initial conditions were simulated for each number of objects. The averaged results are presented in Table 1 and Figure 1. Our algorithm processes 5–20 percent more events (due to update events). However, the cost of each event is linear in n for the traditional algorithm, and sub-linear for our algorithm. Overall, for this particular scenario, the asynchronous algorithm demonstrates a speedup approximately linear in the number of objects.

n	Synchronous			Asynchronous			speedup
	t (s)	#events	t/event (ms)	t (s)	#events	t/event (ms)	
8	58.80	403255	0.1458	16.45	450857	0.0365	3.57
15	200.45	805232	0.2477	36.78	937317	0.0392	5.45
27	748.80	1680996	0.4455	81.20	1987819	0.0408	9.22
42	2123.54	3148716	0.6744	160.69	3680741	0.0437	13.21
64	6357.87	5083977	1.047	317.13	8928818	0.0458	20.08
91	15628.0	10758042	1.453	572.43	11923404	0.0460	27.30
125	37493.2	19098906	1.963	999.69	20122471	0.0497	37.50

Table 1: Results from timing tests comparing our asynchronous algorithm with a generic synchronous algorithm. n is the number of moving objects in the simulation and t is the total execution time. Times were obtained on a 200MHz Intel Pentium Pro.




Figure 1: A plot of the speedups from table 1. Our new algorithm demonstrates a linear improvement for this scenario.

Automatic Recognition and Mapping of Constraints for Motion Retargetting

In this sketch, we address the problem of automatically controlling¹ virtual human models of various sizes and degrees of freedom that imitate the action of a motion-captured subject. We abstract information about significant events and spatial constraints from 3D motion-captured data. A subject, called the primary agent, is motion captured executing a specific action. We are mainly interested in movements involving interactions with other objects (for example, drinking from a cup, digging with a shovel, etc.). The imitators, referred to as secondary agents, try to execute the same action by interacting with the objects in a similar fashion. As the primary and secondary agents may be of different sizes, the motion imitation will be similar but not exact. We assume that maintaining spatial constraints for hands and eyes is more important than maintaining a similar trajectory or motion style between the constraints.

Gleicher² has addressed the problem of mapping constrained motions to different-sized agents using space-time constraints. But that technique cannot be solved in real-time, does not consider constraints on the figure's gaze direction, and is unable to automatically recognize the occurrence of the start and end times of constraints.

Zero crossings of the second derivative indicate significant changes in the motion and correspond to the local maxima of velocity. The start and end times of a spatial/visual constraint can be automatically established by the co-occurrence of the zero crossings of the raw data and the proximity of the end effector with the object. To reduce computation time, we tag only some specific sensors (monitoring sensors) for zero-crossing tracking and some specific objects (tag objects) for computing spatial proximity.


We begin by generating motions for the primary agent from motion-captured data. Then, we automatically recognize and map the constraints to a second agent in three stages. In the first stage, Recognition, a constraint is automatically set or deleted based on the proximity of a monitoring sensor and a tag object at a zero-crossing frame of the sensor. In the second stage of Location Determination, the new constraint location for the secondary agent is determined based on the type of the tag object. Finally, in the third stage, Motion Generation, inverse kinematics is used to solve for the joint angles at the relevant zero-crossing frames, and joint-angle interpolation techniques are used for the in-between frames.

To maintain a similar action style (frame-wise variations in the angular velocity), we compute the normalized angular distance, moved in joint space along a trajectory by the primary agent, at each frame. This style factor used as the interpolation factor causes the velocity profiles in the actions of the two agents to be the same and also retains the zero crossings of the original.

Capturing and maintaining visual attention is very important for movement realism in the secondary agent. Without it, actions appear unnatural even if all the other constraints are correctly satisfied. For visual constraints, we use a head-top sensor to track the focal points of attention. At zero crossings of the acceleration of the head sensor, we check for intersections of the line of sight with the tagged objects. The type of the tag object determines the exact location of the visual constraint for the secondary agent.

This new motion-capture tool could work with and augment other motion-editing software. It was successfully tested to automatically recognize and map actions of an adult to various virtual models including children.

Drink-from-mug action mapped from a six-foot-tall adult model (top) to a 4.5-foot-tall child model (bottom).



References

1. Rama Bindiganavale and Norman Badler. Motion Abstraction and Mapping with Spatial Constraints, In *Modeling and Motion Capture Techniques for Virtual Environments*, International Workshop, CAPTECH '98, pp 70-82, Geneva, Switzerland, November 1998. Springer-Verlag.
2. Michael Gleicher. Retargetting Motion to New Characters, In *Computer Graphics Proceedings*, pp. 33-42. SIGGRAPH, ACM, 1998.


Image-based tracking of the human body in 3D motion has become an important technology for generating the motions of characters in computer animations and TV games. This sketch proposes an image-based method for capturing the motion of actors from movies. Assuming that an articulated model of an actor and the 3D pose at a few keyframes (for example, a starting frame, intermediate frames, and final frame), are given, this method can calculate a pose sequence of the actor from an image sequence. The pose sequence is obtained from iterative minimizing of an error function that is composed of some constraints to be satisfied by the human motion.

Using our proposed method, we show examples of capturing actors' motions from old movies. The first row in the figure depicts a tracking result of Scarlett (Vivien Leigh) slapping in "Gone with the Wind," 1939, obtained by overlapping a wire-frame representation of her 3D model onto the images. The second row in the figure depicts a tracking result of the monster (Boris Karloff) in "Frankenstein," 1931.

This tracking captures 3D motions of the actors. Therefore, we can faithfully reproduce their actions in a new movie. We produced a new short animation, "Scarlett is Slapping the Monster," based on the tracking results. The third row in the figure shows several images sampled from the new animation.

Most of the tracking methods proposed so far fit the model to the human body on a frame-by-frame basis. A disadvantage of this approach is that enormous computation resources are required in the process of analyzing the human body. In contrast, our method is a direct approach without the search process. It requires model fitting at least once (usually at the starting frame) to save computational resources and continues estimating and accumulating pose displacements for tracking.^{1,2,3}

The pose displacements (motion parameters) can be directly estimated from the image sequence by gradient-based techniques. However, the estimation errors of the pose displacements are also cumulative. Over the long term, tracking may result in a great discrepancy between the model and the human body. Our method can eliminate this discrepancy by propagating the correct pose given at a few keyframes. Even if a part of the human body is occluded at some frames, this propagation enables the motion in these frames to be estimated.



References

1. M. Yamamoto and K. Koshikawa. Human Motion Analysis Based On a Robot Arm Model, Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp.664-665, 1991.
2. M. Yamamoto, A. Sato, S. Kawada, T. Kondo and Y. Osaki. Incremental Tracking of Human Actions From Multiple Views, Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp.2-7, 1998.
3. C. Bregler and J. Malik. Tracking People With Twists and Exponential Maps, Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp.8-15, 1998.

Color Super-Histograms for Video Representation: Preliminary Research and Findings



Figure 1
 Extraction method for creating super-histograms. A super-histogram is the ordered set of family histograms. The system orders the sets with respect to the temporal length of the video segment they represent.




Figure 2
 An example of an extracted family histogram.



Figure 3
 Comparison of results: the three largest family histograms for four different TV broadcasts.

While working in the areas of keyframe extraction and commercial detection, we noticed that television programs tend to have strong visual consistencies within a specific program series. For example, we can easily identify a program as “Seinfeld” based on a very few frames of video. Obviously, this is due in part to our human ability to recognize key personalities who appear episode after episode. After watching very few episodes, we can also quite easily recognize the locations and set designs, or backgrounds, that recur during episodic programming. Since television content creators carefully craft these sets and locations, television programs develop an inherent visual consistency. With the current state of the art, some of these consistencies tend to be easier for computers to recognize than, say, the actor Jerry Seinfeld. Color is one such important aspect that content creators maintain through set and location design. This is also an aspect that computers are good at recognizing and classifying.

In our current research in progress, we hypothesize that classification of thematic visual elements is useful in overall recognition and classification of video content when we combine such classification with other methods. Specifically, we are developing a method that uses color information to represent video segments. We apply this method after applying video indexing and classification^{1,3,4,5,6} methods.

In order to classify video content based on color, we have devised a method for computing super-histograms, which we outline in Figure 1. Briefly, the method computes color histograms for individual shots and then merges the histograms into a single cumulative histogram called a family histogram based on a comparison measure (for example, Euclidean distance). As in Figure 2, given the keyframes K_1 , for a shot of 192 frames, and K_2 , representing 79 frames, we extract the two histograms H_1 and H_2 . We then merge H_1 and H_2 into family histogram FH_1 . We use a total of 271 frames to create the family histogram. This family histogram represents the color union of the two shots. However, if the histogram of a new frame is different from the previously constructed family histograms, we form a new family. In the end, we have a few temporally interleaved families of histograms to represent the entire television program. We order this set of families with respect to the length of the temporal segment of video that they represent. The ordered set of family histograms is called a super-histogram.

In the examples in Figure 3, we calculated the histogram difference, D , using L2 distance measure. In the formula below, H_c is the current histogram, H_p is the previous histogram and N is the number of color bins (9 in our example). The values obtained using this formula range from 0 to twice the maximum number of pixels per image. To obtain percentage of similarity, we normalize the values of D by dividing with the total number of pixels. The normalized values are between 0 and 1. Values close to 0 mean that images are similar while values close to 1 mean images are dissimilar.

$$D = \sqrt{\sum_{i=1}^N (H_c(i) - H_p(i))^2}$$

The super-histograms representing episodes of the same sitcom look strikingly similar. The histograms extracted from TV news are not similar to the histograms extracted from sitcoms. This method can be used in a variety of applications that need video classification and retrieval methods: video editing, studio archiving, digital library search and retrieval, consumer products, and Web crawling.

Color histogram indexing has traditionally been applied to image retrieval² and video segmentation.⁷ However, our preliminary results show that the histogram extraction and indexing method can be applied to video. We maintain that systematic analysis of thematic visual elements such as color can assist in defining richer characterization and description of video content. To this end, we are pursuing a method for representing video segments using cumulatively averaged, or super-histograms. In the immediate future, this extremely compact representation can be used for a variety of applications including program boundary detection and classification and search in large video archives. In the more distant future, we envision systems that use color as one of many visual elements to understand viewing habits and to reveal information about content at a thematic level.

References

1. M. Abdel-Mottaleb, N. Dimitrova, R. Desai and J. Martino. CONIVAS: Content-Based Image and Video Access System, Proc. of ACM Multimedia, pp 427-428, Boston 1996.
2. S. Krishnamachari and M. Abdel-Mottaleb. A Scalable Algorithm for Image Retrieval By Color, The Fifth International Conference on Theoretical, Experimental and Applied Image Processing, ICIP-98, Chicago, October 1998.
3. S-F. Chang, W. Chen, H. E. Meng, H. Sundaram and D. Zong. VideoQ: An Automated Content-based Video Search System Using Visual Cues, Proc. ACM Multimedia, pp. 313-324, Seattle, 1994.
4. M. Christel, T. Kanade, M. Mauldin, R. Reddy, M. Sirbu, S. Stevens and H. Wactlar. Informedia Digital Video Library, Comm. of the ACM, Vol. 38, No. 4, pp. 57-58, 1995.
5. T. McGee and N. Dimitrova. Parsing TV Programs For Identification and Removal of Non-Story Segments, SPIE Conference on Storage and Retrieval in Image and Video Databases, January, San Jose, 1999.
6. W. Niblack, X. Zhu, J.L. Hafner, T. Bruel, D. B. Ponceleon, D. Petkovic, M. Flickner, E. Upfal, S.I. Nin, S. Sull, B.E. Dom. Updates to the QBIC System, Proc. IS&T SPIE, Storage and Retrieval for Image and Video Databases VI, Volume 3312, pp. 150-161, San Jose, 1998.
7. S. Pfeiffer, R. Lienhart, S. Fischer, and W. Effelsberg. Abstracting Digital Movies Automatically, Proc. Journal on Visual Communications and Image Representation, Vol. 7, No. 4, pp. 345-353, 1996.

We present a high-level language (VRCC) for describing behaviors of reactive autonomous creatures in 3D virtual worlds. VRCC is a concurrent programming language connected to the VRML environment and based on the timed concurrent constraint (TCC) framework of V. Saraswat et al.

Declarative Behaviors Based on Constraints

The basic idea of TCC is that of synchronous programming exemplified by languages such as Esterel by G. Berry et al, programs run and respond to signals instantaneously. Time is decomposed in discrete time points generated by clocks outside the system, and programs are executed between time points. Thus, at each time point, concurrent agents are running simultaneously until quiescence, and then the program moves to the next time-point. This achieves the classic sense-think-act loop of autonomous agent systems. Basic actions performed by the agents are either posting a constraint to a shared store (Tell operation), suspending until some constraint is entailed by the store (Ask operation), performing some method (Call operation), or posting an action to be performed at the next time-point (Next operation). The Next operation allows actions to be delayed or executed repeatedly over time. Combined with the use of constraints and the event-driven behavior of the Ask operations, it provides a very simple and powerful mechanism to declaratively define behaviors. Constraints are used to state relations that have to be satisfied by the agent: minimal distance with regard to some other object (collision avoidance), maximal distance (following behavior), equidistance (flock or boid-like behavior), etc. It is, moreover, easy to combine such simple behaviors to build more complex ones.

Biologically Inspired Creatures

In order to design autonomous, life-like creatures that can autonomously navigate in the 3D world, we propose some simple behaviors derived from biologically inspired models of navigation. There is currently a growing interest in such models in both the artificial life community and the robotics community, and such models could obviously be applied to virtual agents as well. The creatures will have to react to a changing environment and avoid collision with moving obstacles.

Here, we only consider an agent that uses the taxon system for route navigation, with a very limited capacity for intelligence and no cognitive map. In our experiments, the creature traces a simple route toward a goal by avoiding obstacles. Therefore, the basic constraint is non-collision (enforcement of some minimal distance with regard to obstacles). But we have also considered some non-trivial navigation, as the creature does not know the location of the goal in advance but rather has to explore its environment, guided by a stimulus (for example, light or smell) toward the goal (food, for example).

In particular, we have investigated exploration guided by a stimulus using either temporal difference or spatial difference methods. Temporal differences consist of considering a single sensor (for example, the nose) and checking the intensity of the stimulus at every time point. If the stimulus is increasing, the agent continues in the same direction. Otherwise, the direction changes randomly. This behavior is exemplified, for instance, by the chemotaxis (reaction to a chemical stimulus) of *Caenorabditis Elegans*, a small soil nematode.

A more efficient strategy uses the spatial differences method, which requires two identical sensing organs placed at slightly different positions on the agent (for example, the two ears). The basic idea is simply to favor, at any time point, motion in the direction of the sensor that receives the stronger stimulus. This behavior gives very good results, and the creature moves directly toward the goal. When the goal is moved away, the agent reacts instantly to the new location.




Figure 1 Definitions.




Figure 2 Decremental Delaunay triangulation algorithm.

Incremental Delaunay triangulation is a well-studied problem. Removing points from a Delaunay triangulation or decremental Delaunay triangulation is less studied.^{1,2} This sketch presents theoretical results and an algorithm for decremental Delaunay triangulation.

Often, many points are removed, and it can be more efficient to remove many points at once than to remove them singly. Let $\{a,b,c\} \in \mathbb{R}^2$. Let $[a,b]$ and $[a,b,c]$ denote a line segment and triangle respectively. Let F denote the face operator, $F([a,b])=\{a,b\}$, $F([a,b,c])=\{[a,b],[b,c],[c,a]\}$, extended to collections of edges and triangles. Let M be a triangulation. The *boundary* of M , ∂M , is all those edges which are the edge of a unique triangle. If $T \subset M$, then the *sub-boundary* of T in M , δT is given by $\partial T \setminus \partial M$ (\setminus is subtraction). The *sup-interior* of T in M are the vertices given by $I(T)=FF(T) \setminus F(\delta T)$. Figure 1 gives examples of boundaries, a sub-boundary and a sup-interior. Let $V=\{v_1, \dots, v_n\} \subset \mathbb{R}^2$ be such that no four points lie on a common circle. Let $D(V)$ be the Delaunay triangulation of V . The following lemmas and corollary can be proven.

Lemma 1. Let $\{a,b\} \subset X \subset V$ with $[a,b] \in D(V)$. Then $[a,b] \in D(X)$.

Lemma 2. Let $T \subset D(V)$. Then for all $W \subset I(T)$, $D(V \setminus W) \subset D(V \setminus T) \cup D(W)$.

Lemmas 1 and 2 show decremental Delaunay triangulation is a local operation. Furthermore, when removing multiple points many edges and hence triangles remain. The unshaded triangles in Figure 1-(f) will be in the Delaunay triangulation, $D(V \setminus W)$. This is similar to results for incremental Delaunay triangulation.²

Lemma 3. Let V , T and W be as in Lemma 2. Let $X \subset V$ such that $X \cap W = \emptyset$ and $FF(T) \setminus W \subset X$. Let $S \subset D(X)$ be all those triangles outside of δT . Then $D(V \setminus W) = D(V) \setminus T \cup D(X) \setminus S$.

Definition. Let T be a triangulation. For $[W \subset FF(T)]$ let $Star(W) = \{t \in T \mid FF(t) \cap W \neq \emptyset\}$.

Corollary 4. Let $W \subset V$. Let $T = Star(W)$, let $X = FF(T) \setminus W$ and let $S \subset D(X)$ be all those triangles outside of δT . Then $D(V \setminus W) = D(V) \setminus T \cup D(X) \setminus S$.

Corollary 4 gives a method to compute a decremental Delaunay triangulation. We have shown that the Delaunay triangulation of the star is sufficient even when the star is not convex. Figure 2 illustrates Corollary 4 applied to the mesh in Figure 1.

There are similar results for constrained decremental Delaunay triangulation. To remove a constraining edge it is necessary to use the star of the edge, that is the star of the two end vertices of the constraining edge.

Decremental and incremental Delaunay triangulation is used in our work and is an efficient means to maintain a mesh during adaptive visualization of large parameterized surfaces.

Acknowledgement

We would like to thank Andreas Hubeli for our first decremental Delaunay triangulation implementation.

References

1. M. Heller. Triangulation Algorithms For Adaptive Terrain Modeling, Proceedings of the 4th International Symposium on Spatial Data Handling, pages 163-174, July 1990.
2. T. Midtbo. Removing points from a Delaunay Triangulation, Proc. 6th Int. Symposium on Spatial Data Handling, pages 739-750, vol. 2, Advances in GIS Research, 1994.

The physically based approach is successful in creating realistic motion for 3D objects. Modeling complex 3D shapes is, however, a labor-intensive job, and their rendering is computationally expensive. This limits the application of physically based animation to top-end computer animation.

Coupling physically based techniques and image morphing techniques has the potential to yield new directions in computer animation. In this approach, complex 3D models are not required, so it can be applied even to 2D animation. This idea was first proposed by Overveld.¹ Unfortunately, the paper was rather conceptual, without strong experimental support, and little attention has been paid to this approach.

This sketch demonstrates successful examples of this coupling approach. We successfully synthesized the stochastic motion of plants under the influence of wind, in which 2D textures are realistically animated based on dynamic simulation.

Dynamics

Modal analysis was adopted for dynamic simulation because it has been successfully applied to the motion synthesis of swaying trees.² In modal analysis, deformation of branches is described by a second-order differential equation, which can be solved, for example, by Euler's method.

The applied force is calculated from a stochastic wind field synthesized by the Fourier method.² Since the applied force is periodic, the equation has a periodic solution. Periodic motion is convenient because it can be repeatedly used in long animation sequences.

To obtain the required periodic solution, we adopt the superposition method.³

Procedure

Input images can be created by using conventional painting tools. To separate the images into components (for example, trunks and branches), the image parts are created in different layers.

Skeletons

Skeletons are specified for each component. In the current implementation, we manually draw skeletons as line segments on the image via a GUI. However, since our images are already segmented, it is not difficult to automatically extract skeletons from images by using image-processing techniques. We also specify kinematic characteristics, such as natural frequencies, amplitudes, and damping time, which are then converted to mechanical properties.

Dynamic simulation


The dynamic simulation is performed by the modal analysis, and it yields the periodic motion of skeletons.

Morphing

Following the deformed skeletons, the corresponding textures are deformed by using piece-wisely linear mapping functions.

Result

We successfully applied the method to create the texture-based animation, "Gift for Nature." Although the plants were all modeled using 2D textures, the motion appears very realistic. Images to the right are example shots from the sequence.



References

1. C. van Overveld. A Technique for Motion Specification in Computer Animation, *The Visual Computer*, vol. 6, pp. 106-116, 1990.
2. M. Shinya, A. Fournier. Stochastic Motion - Motion Under the Influence of Wind, *Computer Graphics Forum (Proc. of Eurographics '92)*, vol 11, No.3, pp. C-119-128, 1992.
3. M. Shinya, T. Mori, N. Osumi. Periodic Motion Synthesis and Fourier Compression, *The Journal of Visualization and Computer Animation*, vol. 9, pp. 95-107, 1998.

Enhancing the Efficiency and Versatility of Directly Manipulated Free-Form Deformation

Directly manipulated free-form deformation (DMFFD)³ is a constraint-based space deformation technique that addresses the task of free-form solid modeling. The user is able to drag surface points directly and have the surrounding solid conform smoothly. This sketch focuses on improving the efficiency and versatility of DMFFD.

Free-form deformation: (FFD) warps the space surrounding an object and thereby transforms the object indirectly. The object is contained within a parallelepiped lattice of control points. Object points (x, y, z) are assigned local co-ordinates (u, v, w) within the lattice and their deformation is expressed as a weighted sum of control point displacements. For efficiency and local control reasons, we employ trivariate cubic B-spline tensor product basis functions and a world co-ordinate aligned lattice. FFD can be expressed in matrix notation as follows: $\Delta Q = B \Delta P$ (1)

Each row of ΔQ represents a change in the (x, y, z) co-ordinates of an object point. ΔP stores (by row) the changes in those control points that influence the object point; B holds the basis functions. A particular (i, j) entry of B is zero if the control point in row j of ΔP does not affect (because of local control) the object point in row i of ΔQ .

Direct Manipulation

Controlling deformations by moving lattice vertices, while producing sculpted results, tends to be cumbersome and counter-intuitive.³ Direct manipulation uses the motion of a collection of constraint points to dictate alterations in the lattice. This new lattice is then applied to the original object. With DMFFD a user specifies the motion of some constraint points (ΔQ) and the system of linear equations (eqn. 1) is solved to find the corresponding alteration in control points (ΔP). In our context, eqn. 1 is generally underdetermined, and we use a pseudo-inverse solution:

$$\Delta P = B^T(BB^T)^{-1} \Delta Q \quad (2)$$

Efficiency (A Compact Basis Matrix)

We considerably improve the space cost and efficiency of the pseudo-inverse calculation by exploiting the sparsity of B .² Each row of B holds 64 non-zero coefficients (weights for a contiguous 4 by 4 by 4 block of control points). The remaining row entries are zero filled. Instead of standard matrix multiplication to form $A = BB^T$ in eqn. 2, we calculate each $A_{(i,j)}$ entry by intersecting the lattice blocks of row i and j of B and summing the multiplied weights of this overlapping region. The remainder of eqn. 2 proceeds by modified Choleski Factorization. This technique represents an order of magnitude speedup (for example, 15.4 times speedup with 60 constraint points across 10 adjacent blocks) over the previously best approach.³

Versatility (Direct Manipulation of First Derivatives)

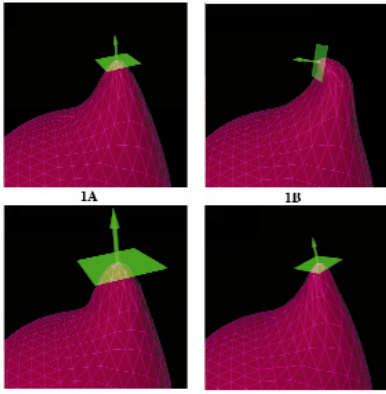
DMFFD provides direct manipulation of the position of constraint points. This can be extended (from work by Fowler¹) to encompass the first derivative frame. For first derivative manipulation at a single point (q) , eqn. 1 takes the form:

$$\begin{bmatrix} b & \emptyset \\ b_u & \Delta q_u \\ b_v & \Delta q_v \\ b_w & \Delta q_w \end{bmatrix} \Delta P = \begin{bmatrix} \emptyset \\ \Delta q_u \\ \Delta q_v \\ \Delta q_w \end{bmatrix} \quad (3)$$

b is a row of basis function coefficients evaluated at q . The subscripts denote partial derivatives with respect to u, v, w as appropriate. The first row ($b \Delta P = \emptyset$) constrains q to a stationary position. The partial derivatives can be extracted in a geometrically intuitive fashion from a rotation matrix (J) applied to the first derivative frame at a point, since: $[\Delta q_u^T \quad \Delta q_v^T \quad \Delta q_w^T] = J - I$ (4)
 This result follows from the contravariant transformation rule⁴ (since J is also the Jacobian matrix of FFD) and the simplicity of the mapping is due to the world co-ordinate alignment of the lattice.

Results

The figures above illustrate an initial object (1A) and the results achievable by tilting (1B), scaling (1C), and twisting (1D) the first derivative frame at a constraint point. Each of these deformations was executed in real-time on an SGI Indigo² 175MHz R10000. Previously, as many as 20 positional direct manipulations would be required to achieve the same results.



References

1. Fowler, B. Geometric Manipulation of Tensor Product Surfaces, Computer Graphics (1992 Symposium on Interactive 3D Graphics), pp. 101-108.
2. Gain, J. Virtual sculpting: An Investigation of Directly Manipulated Free-Form Deformation in a Virtual Environment, MSc. Thesis, Rhodes University.
3. Hsu, W., Hughes, J., and Kaufman, H. Direct Manipulation of Free-Form Deformations, Computer Graphics (SIGGRAPH 92 Proceedings), pp. 177-182.
4. Millman, R., and Parker, G. Elements of Differential Geometry, Prentice-Hall, Englewood Cliffs, NJ.

The reusability of 3D models has led to a proliferation in the number and size of 3D object repositories. These typically store objects in some polygon mesh file format and allow selection by means of categories and keywords. While keyword searches are useful in many circumstances, they suffer from the subjectivity, ambiguity, and resolution difficulties inherent in natural language. These problems become more acute as database sizes increase and may force users to sift through large numbers of returned thumbnail images.

One solution is to query the object database by example. The user constructs a rough low-resolution query model with the shape characteristics of the target object. This is then compared against entries in the database, and the closest matches are returned in order of similarity. The core of this method is the matching algorithm.

We advocate the use of wavelet decomposition¹ because this technique:

- Accepts objects at widely differing resolutions.
- Allows an overall reduction in fine detail by truncating wavelet coefficients, thereby mimicking the human focus on global shape characteristics.
- Has, as will be shown, fast search times.

Our implementation is an extension of wavelet-based image querying² to 3D polygon mesh objects. As shown in Figure 1, wavelets can be used to derive a “signature” from the original object in four steps: normalization, voxelization, wavelet decomposition, and signature extraction.

Polygon Mesh Objects

Our implementation imposes two restrictions on the input objects.

1. Only shape is considered, and surface details are discarded. In future work, this limitation could be overcome by a secondary image query on any texture maps.
2. A standard “natural” orientation and position are assumed. For instance, the positive y- and z-axes represent upwards and forwards respectively, and the object is centered on the origin.

Normalization

Given these considerations, normalization proceeds by uniformly scaling the object to fit within a unit cube, thereby preserving the aspect ratio and orientation.

Voxelization

This is accomplished using a 3D generalization of the standard polygon scan conversion algorithm.

Wavelet Decomposition

The non-standard Haar wavelet basis² is applied in each dimension to decompose the voxel array into a set of wavelet coefficients. The Haar basis was selected because of its simplicity and consequent efficiency. Inclusion and comparison of other wavelet basis functions are left as areas for future research.

Signature Extraction

In keeping with Jacobs et al.,¹ a compact signature is then extracted by truncating and quantizing the coefficient array. Only the sign of the larger coefficients is retained. One exception is the average intensity, which is stored without quantization.

The Signature File

Two operations are permitted on a signature: incorporation into and comparison against the signature file. Since adding signatures is not time critical and can be executed offline, the signature file is optimized for comparison operations by using Jacobs et al.¹ inverted structure.

Performance

The signature extraction process executes in less than a second for an object with 8000 polygons. Since query models are likely to have low polygon counts, this is not a time-critical component of the system. The table lists the execution times of a signature query for different database sizes (numbers of models). Timings were taken on a Pentium II 300 MHz with 96Mb RAM.

Database Size	10000	20000	30000	50000
Query Times	0.001s	0.01s	1.0s	2.5s

Figure 1




FIGURE 1. The operations and data structures involved in deriving a wavelet signature from a polygon mesh object.

References:

1. Stollnitz, E., De Rose, T., and Salesin, D. Wavelets For Computer Graphics: Theory and Applications, Morgan Kaufmann, San Francisco (1996).
2. Jacobs, C., Finkelstein, A., and Salesin, D. Fast Multiresolution Image Querying, Computer Graphics (SIGGRAPH 95 Proceedings), pp. 277-286.

Filtered noise¹ has become an indispensable tool for creating texture and animation effects. It is best generated by convolving a filter function with a sparse grid of pseudo-random numbers (PRNs).² If the grid of PRNs is "poorly packed," the number of grid points falling within the filter radius can vary wildly depending on the location of the sample point. This causes undesirable axis-aligned features.

In this sense, a cubic grid in space is poorly packed. In a 3D grid sampled with a filter of radius 2 grid spacings, the number of grid points used varies from 28 to 56.

We have improved noise quality by using a grid that is more densely and evenly packed than a simple, axis-aligned cubical arrangement. This works particularly well in four dimensions.

In two dimensions, the densest and "most even" packing of points is the triangular grid (Figure 1). This can be indexed using a skewed grid suggested by the blue lines.




Figure 1
Ideal 2D packing




Figure 2
Tight packing in 4D

Spheres do not pack well in 3D, but in 4D there is a very neat packing related to the 2D case. Figure 2 shows a 3D cubical grid of red spheres with an additional blue sphere at the centre of each cube. The blue spheres constitute a second, identical, interleaved grid. A grid in 3D space can be seen as a 3D slice of a corresponding grid in 4D. The blue spheres represent an adjacent slice, shifted along the red/blue axis. In Figure 2, the spheres are drawn undersize to show the structure. If the spheres have a unit diameter and the grid consists of unit cubes, then the (red) spheres touch their neighbors.

The spheres are really 3D sections of 4D hyperspheres, and the blue hyperspheres in the next layer each touch eight of the adjacent red hyperspheres. Each hypersphere touches a total of 24 adjacent hyperspheres: eight in its own layer and eight at the corners of a cube in the 3D layer "before" and "after." Like the 2D packed grid, this grid can be indexed neatly by means of skewed axes.

Using this structure and a filter radius of 1.3, between 25 and 40 grid points affect any one noise value. We achieve a cheaper and better-quality 3D noise by taking a slice from a 4D structure than by working directly in 3D.

Is this structure the best possible? Surprisingly little has been proved about densest packing of hyperspheres. Conway and Sloane³ have published a useful set of plausible hypotheses. Our structure fits with these.

An interesting application of 4D noise is the generation of a simulated underwater lighting effect. A 3D slice of the noise is used to apply lighting to every surface in the scene. The slice is then moved through the fourth dimension to produce the appearance of realistic caustics changing with time. Figure 3 shows a frame from "Dragon Gate."⁴ The waterfall, mist, splashes, pebbles, and underwater lighting are modeled and animated using 4D noise.

Implementation details are available at:
atlas.otago.ac.nz:800/~geoff/graphics/Noise.html




Figure 3
"Dragon Gate:" Leaping Fish

References:

1. Perlin, K. An Image Synthesizer, Computer Graphics (SIGGRAPH 85 Conference Proceedings), Vol. 19, No. 3, July 1985, pp. 287-296
2. Lewis, J.P. Algorithms For Solid Noise Synthesis, Computer Graphics (SIGGRAPH 89 Conference Proceedings), Vol. 23, No. 3, July 1989, 263-270
3. J.H. Conway and N.J.A. Sloane. Why Are All The Best Sphere Packings In Low Dimensions? Discrete and Computational Geometry, 13, 1995, 383-403
4. "Dragon Gate," Directed by Geoff Wyvill, University of Otago, 1999

Recursive subdivision is an active area of research in computer graphics, modeling, character animation, and multiresolution.

A subdivision surface can simply be defined by a couple (P,R) where P is a mesh of arbitrary topology and R is a set of rules. The idea is to apply R to P to generate another mesh to which R is applied again and so on. At the limit, the mesh converges to a smooth surface.^{1,2}

A polygonal strip complex is a sequence of n -sided polygons (or panels) such that every two adjacent panels share one edge only. The set of midpoints of the shared edges form the vertices of a control polygon called mid-polygon of the strip.

A free-form curve can be defined by a couple (Q,R) where Q is a polygonal strip complex and R is a set of rules applied to Q . At the limit, the strip will converge to a curve whose control polygon is difficult to determine. However, if the panels of the strip are symmetric, the limit curve will be the piecewise B-spline curve of the mid-polygon of the strip Q . A panel is symmetric if its vertices are symmetric about the Chebychev points of the edge joining the midpoints of its shared edges. A straightforward application is that if a symmetric strip Q is included in the mesh defining a subdivision surface S , the limit curve of Q is interpolated by the limit surface S with $G1$ continuity across.³

If two strips share one panel f (intersecting strips), their limit curves will intersect at the centroid of f , and the curves are tangent to it. If the shared edges of f are opposite (adjacent), the curves meet with $C1$ ($C0$) continuity (Figure 1).

As applications, given an arbitrary set of intersecting strips, a subdivision surface can be generated through their intersecting limit curves,⁴ (Figure 2). Alternatively, given a mesh P with tagged control polygons, our method constructs corresponding strips on P or its subsequently refined mesh, and the limit surface will interpolate the curves of the tagged polygons. Another application is that the shape of a subdivision surface can be controlled across its interpolated curves by repositioning the vertices of the panels around its Chebychev points.

Finally, a surface can be trimmed along an interpolated curve by cutting along the mid-polygon of the polygonal strip.

Our future work incorporates extending the quadratic approach to higher-order surfaces.⁵




Figure 1



Figure 2

Curve

References

1. E. Catmull and J. Clark. Recursively Generated B-Spline Surfaces On Arbitrary Topological Meshes, *CAD*, 10(6):350-355, 1978.
2. Doo, H. and Sabin, M. Behavior of Recursive Division Surfaces, *CAD*, 10(6):356-349, 1978
3. A. Nasri. Curve Interpolation In Recursively Generated B-Spline Surfaces Over Arbitrary Topology, *CAGD*, 14(1), 1997.
4. A. Nasri. Interpolating Meshes of Curves By Subdivision Surfaces, TR-10/97, American University of Beirut, 1997.
5. A. Nasri. Recursive Subdivision of Polygonal Complexes and Its Applications, in *CAGD*, submitted for publication.

Handheld Interactions: Tailoring Interfaces for Single-Purpose Devices

One of the world's largest financial firms, and the clear technology leader on the floor of the New York Stock Exchange currently has a simple application running on a handheld computer (the Casio Cassiopeia). It allows brokers to send digital-ink "looks" (descriptions of market conditions for specific stocks) directly to traders and salespeople. We are completing design and implementation of a huge step beyond this: actual posting and execution of orders.

Orders are directions from traders and salespeople to buy or sell stocks and other financial instruments. The handheld device will receive orders via a wireless link and allow brokers to execute one or more trades that fill the order. This extremely time-critical and human relationship-centered process is currently carried out with hasty scribbles on scraps of paper. Replacing the paper in this scenario is a very difficult task, since paper is concrete, easy to use, carries lots of information in both explicit and implicit forms, and rarely needs rebooting. We have developed innovative new interface widgets combining the benefits of text and object interactions, and synthesized them into a consistent overall design.

This project combines traditional UI design, applied interaction research, and physical interface design similar to industrial design. This sketch outlines our physical/graphic user interface design process, including the following steps: needs analysis, onsite immersive observation, designing with the user, paper and interactive prototypes, pilot release, user testing and observation, and production release. Each step is illustrated with examples from specific projects and physical artifacts.

Interfaces

Head-Mounted Projector for Projection of Virtual Environments on Ubiquitous Object-Oriented Retroreflective Screens in Real Environment

Masahiko Inami
Naoki Kawakami
Dairoku Sekiguchi
Yasuyuki Yanagida
Taro Maeda
Susumu Tachi
Kunihiko Mabuchi
The University of Tokyo
minami@star.t.u-tokyo.ac.jp
www.star.t.u-tokyo.ac.jp/

The head-mounted display (HMD),¹ or head-mounted projector (HMP), and other projection-based virtual reality systems such as the CAVE² are typical examples of virtual reality visual displays. Although quite useful, these two kinds of displays have some demerits. The disadvantage of the former is the trade-off of high-resolution and wide field of view. On the other hand, the latter systems impose a shadow of the user's body on the virtual environment and the interaction of the user's virtual body with the user's real body.

Design of an HMP with X'tal Vision

HMP is one of the traditional methods in the field of stereoscopic displays.³ However, the conventional implementation presents obstacles. The projector's small depth of focus is a serious problem, because it is the depth of focus that determines the limit of the movable area in which the user can observe a focused image on the screen. An additional problem is the uneven brightness of the projected image, because the normal screen diffuses incident rays. Hence, the shape of the screen is restricted to a plane, cylindrical, or spherical surface for easy compensation, or the brightness of the image should be calibrated, which requires complex image computation.

X'tal Vision (Crystal Vision), which uses a projector with a small iris and retroreflective material as its screen, was originally developed for an object-oriented display.⁴ With this system, the user observes a projected stereoscopic virtual environment on ubiquitous object-oriented retroreflective screens in the real environment. Because X'tal Vision is a simple but advantageous method, we speculated that applying X'tal Vision in the HMP method would be suitable for virtual/augmented-reality systems.

Figure 1 illustrates the principle of the display. A projector is arranged at the conjugate position of a user's eye, and an image is projected on a screen made of or painted or covered with retroreflective material. A small iris is placed in front of the projector to secure adequate depth of focus.

The retroreflector screen together with the small iris ensures that the user always sees images with accurate occlusion relationships. This means that if a real object has retroreflective material on it, it becomes a part of the virtual environment and disappears, replaced by a virtual object. On the other hand, if it does not have retroreflective material on it, it will occlude the virtual environment without any troublesome shadows on the virtual environment.

In the configuration of X'tal Vision, screen shapes are arbitrary, any shape is possible. This is due to the characteristics of the retroreflector and the small iris in the conjugate optical system. By using the same characteristics of X'tal Vision, binocular stereo vision becomes possible using only one screen with an arbitrary shape. This system should be mounted on the head of the user as an HMP. Two liquid crystal projectors mounted on a helmet project full-color NTSC-resolution images.

Applications

In the first application, we observed the virtual images of a skeleton on the retroreflective screen as if the patient's body had become transparent (Figure 2). The second application involves ubiquitous displays. Any material painted or covered with retroreflective materials can work as a display. We demonstrated a paper-type display (Figure 3). The third application is optical camouflage. We can make any object transparent, which interferes with images of a virtual object (Figure 4). These applications verify and demonstrate the effectiveness of an HMP with X'tal Vision.

References

1. Sutherland I. A Head-Mounted Three Dimensional Display, Proc. Fal Joint Computer Conference, AFIPS Conf. Proc., vol. 33, pp. 757-764, 1968.
2. Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE, Proc. of Computer Graphics, vol. 27, pp. 135-142, August 1993.
3. Ryugo Kijima and Takeo Ojika. Transition Between Environment and Workstation Environment With Projective Head-Mounted Display, Proc. Virtual Reality Annual International Symposium, pp. 130-137, March 1997.
4. Naoki Kawakami, Masahiko Inami, Yasuyuki Yanagida and Susumu Tachi. Object-Oriented Displays, Conference Abstracts and Applications (SIGGRAPH 98), p. 112, July 1998.




Figure 1
Principle of an X'tal Vision system.




Figure 2
Diagnosis-assisted system.



Figure 3
Prototype of paper-type display.



Figure 4
Example of Optical Camouflage for a real scene.

The Holodeck Interactive Ray Cache

Rendering with full global illumination enables designers to make intelligent choices regarding the materials, geometry, and lighting of complex objects and spaces. However, typical designers cannot wait all day for a final rendering of a candidate building or car design. But they want to look at their design, and walk around to see it from every angle. Conventional ray tracing doesn't facilitate this, because its view-dependent calculation must start over at each new eye point. A radiosity solution allows free movement, but only in diffuse environments, where there is much less to see. What we need is accurate visualization of realistic environments that is both interactive and progressive, employing parallel processing and graphics hardware to provide the fastest feedback possible to the designer.

In this sketch, we present a technique for interactive rendering of complex environments with global illumination and arbitrary reflectance functions as provided by the Radiance lighting simulation system¹ and hardware acceleration. Rendering is interactive in the sense that the user is allowed to move freely in and around designated regions while reasonable camera animation is presented, and the image is refined when the user is still. In the context of our system, Radiance serves as the sample generator.

We introduce a 4D holographic ray-caching data structure, called the holodeck due to its similarity in form and function to the Star Trek invention. This structure serves as the rendering target that allows us to quickly load all precomputed ray samples related to a specific view. No ray computations are ever wasted or lost, and dozens of ray tracing processes may be managed in parallel. The process that manages ray computation and holodeck caching is called the server.

We take advantage of graphics acceleration hardware with custom-tailored display drivers, which render a given set of ray samples into a coherent view. The screen representation is generally "2.5D" in the sense that it encodes some depth information for local motion and stereopsis, but maintains unary depth complexity over most of the image. The driver also manages user interaction and drives the progressive rendering.

The sample generator, server, and display driver work together to compute, cache, and display ray samples interactively as the user moves freely about a space.

The air traffic control tower shown in Image 1a was rendered on a 24-processor SGI Onyx2 with IR graphics. In a precalculation that achieved 96-percent linear speedup running 24 ray trace processes over two hours, we produced a 300-Mbyte holodeck file that contains as much detail as shown in Image 1a at every eye point within the interior section, covering most of the control room. Each new view takes three seconds to retrieve and display 100,000 precalculated samples, and given another 30 seconds, improves to the resolution shown in Image 1b. The most important benefit for design applications is accurate prediction of lighting and visibility as provided by our global illumination and tone mapping methods. These features enable the user to make important design decisions regarding tower equipment and window shades, which in this example, could ultimately affect air traffic safety.



Image 1a
An air traffic control tower cab displayed from a precalculated holodeck file (1a, left), then after 30 seconds of progressive ray tracing on 21 processors.



Image 1b
The Holodeck Interactive Ray Cache.

Reference

1. Ward, Greg. The RADIANCE Lighting Simulation and Rendering System, Computer Graphics (Proceedings SIGGRAPH 94), ACM, July 1994, pp. 459-472.

Painterly rendering algorithms synthesize images with a hand-crafted touch from a source image of a real scene. The current major scheme is based on Haeberli's method^{1,2,3} which paints brush strokes on a canvas successively. In our approach, each rectangular stroke is controlled by several attributes: color \mathbf{C} , location (x_c, y_c) , orientation θ , width w , and length l .

We propose an alternative method to determine these attributes, so that the stroke nicely approximates a local region of the source image.

All stroke attributes are computed with the following steps.

1. Preparation

Firstly, a tentative location is given for a stroke: Figure 1(a). The stroke color \mathbf{C} is sampled from the source image at the location: Figure 1(b). The source image is cropped with a square at the point: Figure 1(c). This cropped image \mathbf{C}_w is to be approximated by the stroke.

2. Color Difference Image Generation

A color difference image is a gray-scale image, obtained from \mathbf{C}_w and \mathbf{C} , Figure 1 (d). Its pixel value $I(x, y)$ is larger if the corresponding $\mathbf{C}_w(x, y)$ is closer to the stroke color \mathbf{C} . $I(x, y)$ is given by:

$$I(x, y) = f(d(\mathbf{C}, \mathbf{C}_w(x, y))).$$

Here, $d(\mathbf{C}_1, \mathbf{C}_2) \in R$ is the Cartesian distance of two colors, $\mathbf{C}_1, \mathbf{C}_2$, in the CIE- $L^*u^*v^*$ color space. $f(d) \in [0, 1]$ is the function that maps the color distance d to the pixel intensity of the color difference image. In our current implementation, $f(d)$ is:

$$f(d) = \begin{cases} (1 - (d/d_0)^2)^2 & \text{if } d \in [0, d_0] \\ 0 & \text{if } d \in [d_0, \infty) \end{cases},$$

where d_0 is a constant value which indicates the bound of 'similar color'.

3. Equivalent Rectangle Calculation

The image moment is defined on the gray-scale image, $I(x, y) \in [0, 1]$. The image moment of l th degree about the x -axis and m th degree about the y -axis is defined as:⁴

$$M_{lm} = \sum_x \sum_y x^l y^m I(x, y).$$

For $n = l + m$, M_{lm} is called the image moment of n th degree. The binary image of the equivalent rectangle has the same zeroth, first³ and second moments as the original gray-scale image: Figure 1(e). The stroke attributes, (x_c, y_c) , θ , w and l , are determined by the parameters of the equivalent rectangle: Figure 1(f)). The color difference image is approximated by this rectangle, which determines the location, orientation, and size of the stroke.

Results

One of the results is shown in Figure 2. Note that the stroke size varies according to the local complexity of the source image. It means the flat area is painted by larger strokes, while the details are painted with smaller ones.




Fig. 1 Steps of algorithm. (a) Source image. The white rectangle indicates the tentative location. (b) Stroke color. (c) Local region of the source image. (d) Color difference image. (e) Equivalent rectangle. (f) Rendered brush stroke.



Fig. 2 Result.

References

1. Paul Haeberli. Paint By Numbers: Abstract Image Representations, SIGGRAPH 90 Conference Proceedings, 1990.
2. Barbara J. Meier. Painterly Rendering for Animation, SIGGRAPH 96 Conference Proceedings, 1996.
3. Peter Litwinowicz. Processing Images and Video for an Impressionist Effect, SIGGRAPH 97 Conference Proceedings, 1997.
4. William T. Freeman, et al. Computer Vision for Interactive Computer Graphics, IEEE Computer Graphics and Applications, Vol. 18, No. 3, 1998.

Stroke

Image Re-Composer

In order to draw or take a good picture, it is important to carefully consider how to compose the picture. If a picture is drawn or taken without an understanding of the composition, the result will be a boring or mediocre picture. There are no absolute rules for such composition, but there is a great deal of understanding of what will occur in terms of visual literacy organization and orchestration. Unfortunately, novice painters or photographers who do not have enough such knowledge usually find it difficult to determine the appropriate composition of a picture currently being drawn or about to be shot.



Re-Composition Galley

Image Re-Composer is a post-production tool that decomposes a picture and regenerates it as a new improved picture. Since the tool allows the user to generate different pictures depending on specified compositions, the user can experiment with a variety of good compositions of the original picture. As a result, the user can learn how masters of the past maintained the visual balance of pictures.

Image Re-Composer consists of three modules:

1. Figure extractor. The figure extraction method extracts figures based on the characteristics of the V4 cortex in the human visual system, which plays an important role in figure-ground segmentation. The figure extractor first segments an input picture into several regions. It then discriminates the segmented regions into portions of the figures and a part of the ground. Finally, it shows the user the extraction results, which include figures in color and the ground in gray. If the user is not satisfied with the results, they can be corrected by selecting or de-selecting regions that are mis-discriminated by the extractor. Then, the user is asked to discriminate each object by grouping the extracted regions.

2. Composition analyzer, which extracts composition information, such as the location, size, and shape of a figure, based on the idea of the golden section. The golden section has been used throughout the history of art to make the most beautiful and ideal proportions of architectures or art work. The Dynamic Symmetry principle has been used to analyze the composition of paintings according to the golden section. This principle can find golden

section points in a rectangle whose aspect ratio is the golden proportion. Actually, the aspect ratio of canvas for paintings and printing paper for photographs is the golden proportion. Therefore, applying the Dynamic Symmetry principle to paintings or photos makes it possible to determine how objects in a picture are composed to maintain a good visual balance of the picture. The composition analyzer first applies the Dynamic Symmetry principle to the figures extracted by the extractor. It then composes abstracted figures that are constructed by base lines of the golden section. Finally, it extracts the composition information from the abstracted figures.

3. Composer, which recomposes a picture according to the composition information extracted by the composition analyzer. For the image recomposition, the composer asks the user to input three pictures: an original picture, a ground picture, and a guide picture. The ground picture is a picture onto which the system recomposes figures of the original picture. The guide picture is a picture with composition information, which provides recomposition guidelines for the system. The composer recomposes the user-selected objects in the original picture by adjusting the sizes and locations of the objects according to the guide picture.

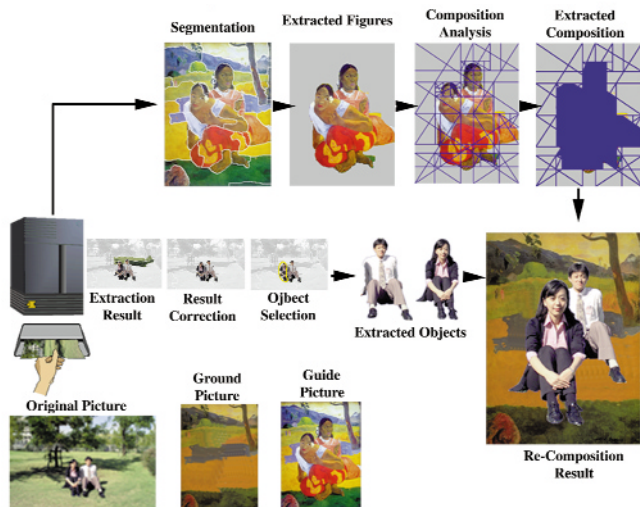


Image Re-Composer

This animation sketch presents how image-based modeling, rendering, and lighting were used to create the animation "Fiat Lux" (SIGGRAPH 99 Electronic Theater). The film features a variety of dynamic objects realistically rendered into real environments, including St. Peter's Basilica in Rome. The geometry, appearance, and illumination of the environments were acquired through digital photography and augmented with the synthetic objects to create the animation.

The Imagery and Story

"Fiat Lux" draws its imagery from the life of Galileo Galilei (1564-1642) and his conflict with the church. When he was 20, Galileo discovered the principle of the pendulum by observing a swinging chandelier while attending mass. This useful timing device quickly set into motion a series of other important scientific discoveries. As the first to observe the sky with a telescope, Galileo made a number of discoveries supporting the Copernican theory of the solar system. As this conflicted with Church doctrine, an elderly Galileo was summoned to Rome, where he was tried, convicted, forced to recant, and sentenced to house arrest for life. Though honorably buried in Florence, Galileo was not formally exonerated by the church until 1992. "Fiat Lux" presents an abstract interpretation of this story using artifacts and environments from science and religion.

The Technology

The objects in "Fiat Lux" are synthetic, but the environments and the lighting are real. The renderings are a computed simulation of what the scenes would actually look like with the synthetic objects added to the real environments. The techniques represent an alternative to traditional compositing methods, in which the lighting on the objects is specified manually.


The environments were acquired in Florence and Rome. The images in St. Peter's were taken within the span of an hour in accordance with our permissions. To record the full range of illumination, we used high-dynamic-range photography, in which a series of exposures with varying shutter speeds is combined into a single linear-response radiance image. Several scenes exhibited a dynamic range of over 100,000:1.

The appearance and illumination of each environment was recorded with a set of panoramic images and light-probe measurements. Each light-probe measurement was made by taking one or two telephoto radiance images of a two-inch mirrored ball placed on a tripod; each provided an omnidirectional illumination measurement at a particular point in space. Several radiance images were retouched using a special high-dynamic-range editing procedure and specially processed to diminish glare.

We constructed a basic 3D model of each environment using the Façade photogrammetric modeling system. The models allowed us to create virtual 3D camera moves using projective texture mapping and fix the origin of the captured illumination. The light probe images were used to create light sources of the correct intensity and location, thus replicating the illumination for each environment. The illumination was used to "un-light" the ground in each scene, allowing the synthetic objects to cast shadows and appear in reflections. The dynamic objects were animated either procedurally or by using the dynamic simulator in Maya 1.0. Renderings were created on a cluster of workstations using Greg Larson's RADIANCE global illumination system to simulate the photometric interaction of the objects and the environments. The final look of the film was achieved using a combination of blur, flare, and vignetting filters applied to the high-dynamic-range renderings.

Supported by Interval Research Corporation, the Digital Media Innovation program, and the Berkeley Millennium project.

Contributors: Christine Cheng, H.P. Duiker, Tal Garfinkel, Tim Hawkins, Jenny Huang, and Westley Sarokin



Galileo

References

1. Paul E. Debevec, Camillo J. Taylor, and Jitendra Malik. Modeling and Rendering Architecture from Photographs, SIGGRAPH 96, August 1996.
2. Paul Debevec and Jitendra Malik. Recovering High Dynamic Range Radiance Maps from Photographs, SIGGRAPH 97, August 1997.
3. Paul Debevec. Rendering Synthetic Objects Into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and Dynamic Range Photography, SIGGRAPH 98, July 1998.

Image-Based Techniques for Object Removal

In feature film effects, it is often necessary to remove an object from the background plate. The obvious solution is to remove the object from the actual set, and to shoot a second camera pass, called a clean plate. However, this is often not done or not possible. The camera move may not be repeatable, or the set may have been destroyed with the first take. This sketch describes a simple image-based rendering technique for producing a clean plate from the original background.



The algorithm consists of four steps: object identification, image depth extraction, image projection, and incremental rendering. In a full system, each of these portions can be rather complicated, but they all can be approximated through simpler methods to get the job done. Simply put, the algorithm makes a hole where the bad pixels are, then uses neighboring frames to incrementally fill the hole with good pixels.

The object identification phase can be as easy as a hand rotoscoped matte or as complicated as an optical flow system. The result needs to be a matte for the object to be removed. With these mattes, a hole is produced in each frame in the sequence, so that only background pixels remain.

In a complete 3D solution, the image depth would be determined for each pixel so that it can be projected into a different camera space in the next step. In the simple example shown, however, the depth can be approximated by assuming the ground is a flat plane. This also allows the following projection step to be performed as a 2D warp. In either case, the requirement is to create enough information to allow a future or past frame to be warped to the view of the current frame.

Image projection and incremental filling are done at the same time. For every frame to be cleaned, a frame is chosen that is incrementally far away in time. For this implementation, it was simply frame-1, frame+1, frame-2, frame+2, etc. Each of these frames is warped into the camera view of the current frame. In the full 3D solution, this involves taking the screen coordinates and known depths from the other camera, projecting into world space, then projecting back through the current camera to produce each pixel to render. In the simple 2D case, this is just a 2D-perspective matrix that can be determined from corresponding points in the image. Both methods will create a warped frame with a hole in it. But this hole will not correspond exactly to the hole that is being filled. This means valid pixels can be filled in, by comping the warped frame under the current frame. Future and past frames are accumulated in this manner until the hole is filled, or until some temporal threshold is reached.

In many cases, the object being removed was serving as a stand-in for a CG character, so the hole does not need to be entirely filled in. Some objects need to be removed completely, such as cables and rigs, and these narrow pieces usually fill in quickly. Note that there are antialiasing and lighting issues with this method, but this technique provides a simple way to obscure unwanted foreground material with very similar pixels in the desired background.

Object Removal

Constructive Solid Geometry has many features that make it attractive for conceptual design (for example, hierarchical modeling and intuitive operations). Existing techniques for interactive CSG¹ update the geometry of the entire model for every object manipulation. For larger models, such techniques cannot achieve the responsiveness required for interactive work. We use a space-subdivision technique to reduce the amount of computation by recomputing only those parts of the model that have changed since the last update.

Every primitive is polygonized in order to take advantage of graphics display hardware. This is done by dividing each primitive's surface up into a number of rectangular areas of space. We trace rays in a 3D grid inside these areas and find all intersection points of the primitive's algebraic surface with the grid. These points are polygonized by an implicit surface tiler² to generate a number of triangles inside every grid cell.

The modeling workspace is subdivided into an octree, typically three to four levels deep. Each node in the CSG graph contains an octree data structure that stores empty/full information. At the bottom level of the octree, if the object's surface is present, a "leaf node" stores a reference to every grid that intersects that area of space. Octrees that represent a CSG combination will reference grids that belong to a number of different primitives.

When a primitive object is created or moved, its octree is found by testing each section of the octree against the algebraic function. The octree is subdivided down to the bottom level, where each leaf node tests to see which grids overlap its area. References to these grids are stored, and if there are any changes from the previous octree then a "changed" flag is set and propagated up the octree structure, so that higher-level nodes are changed if any of their children are.

To perform a CSG operation between two objects, we traverse the changed sections of the octrees, looking for grids that overlap each other. For each cell in each grid, we find which cells it overlaps in each of the opposing grids that overlap its section of the octree (see Figure 1).

We use a triangle-splitting algorithm based on Laidlaw et al¹ to split the triangles in a cell against the opposing triangles, determining for each split triangle whether it lies inside or outside the opposing object. The CSG operation determines which triangle subset is inserted into the combined octree.

The split operation is $O(n*m)$ in the number of triangles, but the octree division of the workspace reduces the number of grids that must be tested, and the uniform subdivision of each grid reduces the number of cells whose triangles must be split against. This allows us to keep n and m down to 2-10 triangles in almost all cases, although a few cells that lie on corners of the model may have up to 40 triangles.

Complex cutting operations may be performed on models such as Figure 2 with approximately five updates per second (Pentium II - 450 MHz). This lets us build CSG objects interactively, using a direct-manipulation interface.

This system meets many of the requirements of conceptual design; it allows rapid prototyping of objects and gives immediate feedback on changes, while retaining the features that make CSG attractive for solid modeling.

www.cs.otago.ac.nz/postgrads/cbutcher/s99.




Figure 1
Each grid cell tests every overlapping grid to find cells that it overlaps. Triangles in the red grid cell are split against triangles from all the highlighted blue grid cells.

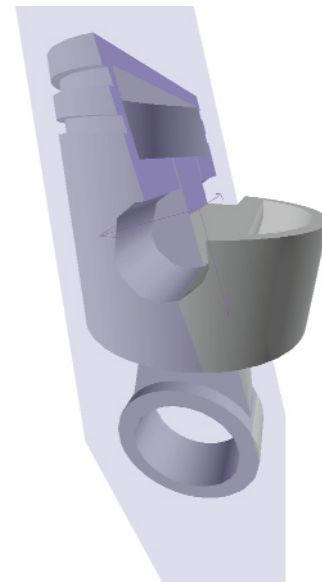


Figure 2
The cutting plane may be moved through this model in real-time, using the direct manipulation device. Approximately five frames per second are achieved during the manipulation.

References

1. D. Laidlaw, W. Trumbore and J. Hughes. Constructive Solid Geometry for Polyhedral Objects, Computer Graphics, Proceedings of SIGGRAPH 86, pages 161-170, August 1986.
2. P. Ning and J. Bloomenthal. An Evaluation of Implicit Surface Tilers, IEEE Computer Graphics and Applications, 13(6):33-41, Nov 1993.
3. A. Rappoport and S. Spitz. Interactive Boolean Operations for Conceptual Design of 3D Solids, Computer Graphics Proceedings, SIGGRAPH 97, pp. 269-278, August 1997.

M. Brady
K. Jung
H.T. Nguyen
Intel Corporation
brady@allover.com

R. Mullick
National Institutes of Health

W. Lawton
T. Poston
S. R. Ranjan
K. Schulz
S. Venkataraman
R. Viswanathan
Y. Yu
G. Zhu
National University of Singapore

R. Raghavan
John Hopkins University

Interactive Haptic Modeling of Tensegrities and Network Structures

We address handling deformable objects in the context of a reach-in environment¹ for dextrous work with virtual objects and tools. The user's interaction with the object is two-handed: the "workpiece" hand (commonly the left) moves the object rigidly in space with a tool (a stylus whose position and orientation are continuously monitored). The "detail" hand's stylus is part of a haptic device used to probe the structure and observe the response (visually and haptically).

Although our methods generalize readily, we limit ourselves here to tensegrity structures. They have many uses, in areas such as architecture, civil engineering, and biology, where they have been considered as simplified models of cell structure.² In addition, they are useful benchmarks for strategies for fast computation and haptic feedback. We have applied these techniques in the larger context of models of cell mechanics. The methods generalize directly to models of tissues, organs, or other deformable structures, and we anticipate a wide range of uses in the near future in medicine and education. The combination of working with full hand-eye coordination and dexterity in a 3D environment, and visualization and haptics, is very compelling for humans as tool users.

A tensegrity is a 3D framework of rigid rods and "rubber bands" attached at nodes. Consider a node, m , picked and displaced with the detail hand. On its initial selection, we equilibrate with the constraints that this node and one base rod are fixed. After this, we calculate a 3×3 linear-response matrix A where $f = A\Delta$, with f being the incremental force in response to the displacement Δ of node m . From this, we can calculate the change in displacement and orientation of all of the rods in response to the applied incremental displacement Δ . Displacing a rod (in both position, x , and orientation, ϕ) is handled similarly, except that the input is now a 6D vector, so linear-response matrices are 6×6 .

For a large structure, the above computation exceeds the one-millisecond delay allowed for smooth haptic interaction. We therefore apply two approximation techniques:

1. For small deformations, we compute and display the actual movement of the structure (so the result is visually correct) but compute the force felt with a constant linear response matrix A . This process may be refined by updating the linear-response matrices at reasonable intervals as the stylus is being displaced.
2. For larger deformations, we use a $3 \times 3 \times 3$ grid of points in space for each node to represent a discretization of the displacements of the node in space. The effective linear-response function is computed for the 27 values for each node. From the actual displacement of the node, an interpolated force is provided. This large precomputation provides a fairly accurate response.

A different technique gives a feel for the static structure of the object. We implement a force field, rather than computing surface contacts (which our linear elements lack).³ If the stylus tip x is within a 3D rectangular bounding box around the center line of a rod or string, we provide an attractive force toward the nearest point x' on the centerline. Its form is arbitrary, but we use $f = -a(d^2 - |x-x'|^2)(|x-x'|)$, where $a > 0$ is an adjustable constant, and $2d$ is the width of the bounding box. Where two rods or strings join, we interpolate between their centerlines.

In conclusion, we have computed the response of structures as accurately as compatible with haptic response (action and reaction are "instantaneous"), which is the fastest refresh rate required (every millisecond, compared with the visual, which needs refresh only every 30 milliseconds or so) for preserving the feeling of continuity that we have about the real world. We have done so in a fully integrated 3D virtual world for dextrous tool use.

References

1. T. Poston. The Virtual Workbench: A Path to Use of VR, CieMed Technical Report, Sept. 1993. Reprinted with added references: The Industrial Electronics Handbook, J. D. Irwin ed., CRC Press and IEEE Press, pp. 1390-3, 1997.
2. D. Ingber. The Architecture of Life, Sci. Amer., Jan. 1998.
3. R. W. Lawton, R. Raghavan, S. R. Ranjan, and R. R. Viswanathan. Tubes In Tubes: Catheter Navigation in Blood Vessels and Its Applications, Intl. J. of Solids and Structures, to appear.

Our aim is to render surfaces lit with point-light sources at interactive rates using arbitrary BRDFs, evaluated at per-pixel resolution on contemporary hardware. Neglecting surface position and wavelength, BRDFs are parameterized by at least four degrees of freedom.

Separable decompositions^{1,4} approximate (to arbitrary precision) a high-dimensional function f using a sum of products of lower-dimensional functions g_k and h_k :

$$f(x, y, z, w) \approx \sum_{k=1}^n g_k(x, y) h_k(z, w).$$

Under certain changes of variables many BRDFs are highly separable⁴ and even $N=1$ has proven to be visually adequate for many interesting BRDFs (see Figure 1). Because of the simplicity of the reconstruction process, we can perform the reconstruction at interactive rates using existing hardware support for texturing, compositing, and diffuse lighting.

Parameterization

The $\hat{\omega}_o$ x $\hat{\omega}_d$ parameterization is suitable for some kinds of BRDFs, but for common "glossy" BRDFs, we reparameterize them in terms of the halfvector \hat{h} and a vector \hat{d} , which is $\hat{\omega}_d$ represented relative to a new basis $(\hat{t}', \hat{b}', \hat{h})$ that is created by applying Gram-Schmidt orthonormalization to the tangent \hat{t} , the binormal \hat{b} , and \hat{h} . This parameterization aligns features of the BRDF to make it more separable.

Decomposition

A BRDF can be decomposed using the Singular Value Decomposition^{1,3} or using the Normalized Decomposition (ND) algorithm, which computes only a single-term approximation

$$f_p(\hat{h}, \hat{d}) \approx \tilde{f}_p(\hat{h}, \hat{d}) = g_1(\hat{h}) h_1(\hat{d}).$$

The ND algorithm finds the average normalized profile along \hat{d} and stores it in $h_1(\hat{d})$ and then stores the normalization factors in $g_1(\hat{h})$ (using the p-norm):

$$g_1(\hat{h}) = \left(\int_D |f_p|^p(\hat{h}, \hat{d}) d\hat{d} \right)^{\frac{1}{p}},$$

$$h_1(\hat{d}) = \frac{\int_H f_p(\hat{h}, \hat{d}) d\hat{h}}{g_1(\hat{h})}$$

Both functions $g_1(\hat{h})$ and $h_1(\hat{d})$ are two-dimensional and can be put into hemispherical or parabolic² texture maps for proper interpolation of unit vectors \hat{h} and \hat{d} .

Rendering

First, render the scene diffusely illuminated using $g_1(\hat{h})$ as a multiplicative texture map. Appropriate texture coordinates must be calculated at the vertices according to the parameterization of the BRDF. Without clearing the colour or depth buffers, render the scene again with no illumination and with $h_1(\hat{d})$ as a decalced texture map. For the second pass, set the depth test to "equality" and set the compositing operation to "multiply." This evaluates the following rendering equation (for one point light source) for all \underline{x} :

$$L_o(\hat{\omega}_o, \underline{x}) = \tilde{f}(\hat{\omega}_o, \hat{\omega}_i) \cos \theta_i \frac{I}{r_i^2}$$

If multitexturing support is available, this can be done in a single pass.



Figure 1: Left to right: HTSG copper, Ward's anisotropic model, and measured grey vinyl. Top: Models with BRDFs evaluated at every pixel in a raytracer. Bottom: Hardware-accelerated single-term approximations using OpenGL.

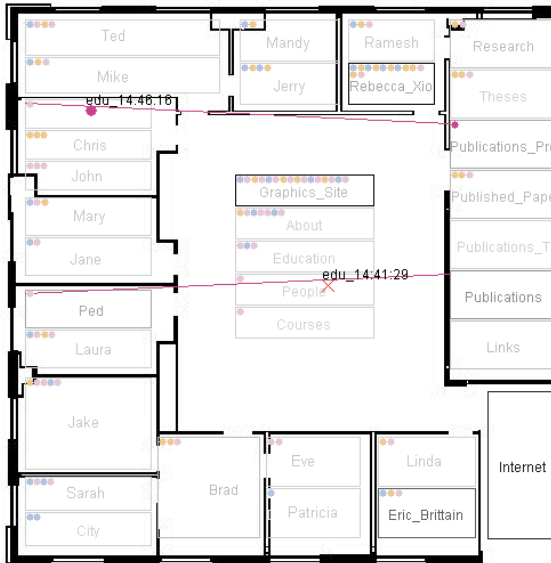


Figure 1
 Web Map visualizes the Web site for a research group. Each office contains one or more rectangular areas that represent group or individual members' home pages. Each visitor to a Web page is represented by a dot in the corresponding rectangular area.

Currently, Web users have little knowledge about the activities of fellow users. They cannot see the flow of online crowds or identify centers of online activity. LiveWeb seeks to enrich Web users' experience by visualizing the real-time activities of other users. The visualizations can help answer user questions about overall patterns and specifics such as: What are other people looking at? What is hot? Who is interested in what I am interested in?

LiveWeb visualizes two kinds of data about a Web site: its underlying structure and real-time user accesses. The visualizations first lay out the site structure, then overlay access data on top. The visual structure of a Web site can be generated in three ways:

1. It can be custom made.
2. It can be automatically generated.
3. It can emerge from users' pattern of traversal.

We have created two sample visualizations, WebMap and WebFan, using the first two methods. WebMap visualizes Website activities that can be mapped to physical structures. The custom-created layout is derived from a floor plan as shown in Figure 1. Each Web user is represented by a dot. As users move among pages, their dots move correspondingly through the WebMap. Other features include:

- Distinguishing users from different domains by color.
- Showing WebMap user's own traversal in context of others'.
- Indicating how long a user has been on a page and what other pages the user has visited.
- Accumulating user accesses over time to identify Web pages that are visited more often.
- Allowing users to navigate directly using WebMap.

WebFan visualizes Web activities using a fan-like hierarchical structure. This abstract structure allows a large set of Web pages with multiple levels to be represented at the same time for overview and comparison. Users can also interactively explore the fan structure to find out more about individual pages.

Figure 2 shows a WebFan visualization of a Web-based message board, which contains user postings and replies that each reside on a separate Web page. WebFan can also be used to visualize a generic Web site by first extracting its hierarchical structure using a Web crawler.

The authors would like to thank Judith Donath for her helpful discussions.

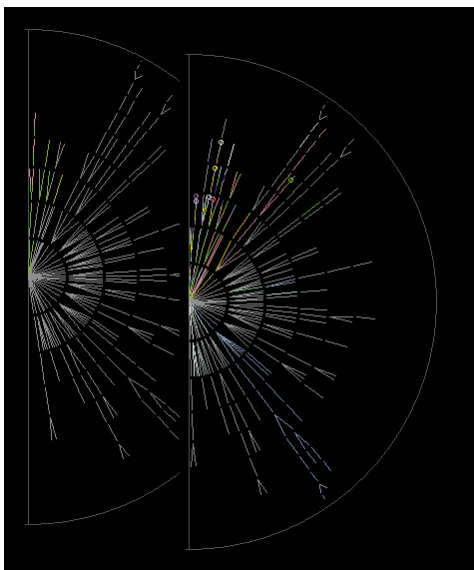


Figure 2
 WebFan visualizes activities on a Web-based message board. It uses the message-reply structure to lay out the message in a fan-like fashion. Each time a Web visitor reads a message, the line segment representing the message will change to the visitor's color. A circle on the colored section indicates that the user is currently viewing the page.

Cloth modeling has long been an active subject in computer simulation and animation. Diverse industries such as clothing and textiles, as well as film and animation, have all contributed to this growing body of knowledge. But realistic and stable computation of cloth self-intersection remains one of the challenges.

Since the early 1990s, animators have used large particle systems or mass-spring lattices to model the behavior of cloth.^{2,3,4} Whether the animations use explicit integration techniques such as Runge-Kutta or implicit techniques such as backward Euler,¹ the models are similar.

Many techniques for handling self-intersections calculate when a vertex has penetrated a triangle and then either adjust the velocity of the vertex with a rigid collision or introduce a stiff spring force to correct the penetration. To eliminate self-intersections completely, such algorithms must handle edge-edge collisions similarly, and track vertex-triangle orientations to ensure surface integrity.⁵

We propose an alternative approach. The triangular mesh in these simulations is only an approximation to the surface itself. Our forces use a different approximation that allows us to handle collisions and friction. Cloth-cloth interactions are highly complex. Colliding patches of cloth push on and slide against each other. Our method prevents self-intersections and reacts in a realistic manner as the cloth collides with itself. Our approach does not sacrifice the visual integrity of the simulation and yields realistic self-collisions in cloth.

The Model

Consider a sheet of material whose physical properties are constant along its surface. We decompose the sheet into a triangular mesh of masses and damped springs. We chose a hexagonal decomposition to avoid introducing features inherent to square grids. To introduce stiffness to the system, both for bending and twisting, we attach damped springs that span the center vertex of each hexagon.

For self-intersections, we approximate the fabric by a lattice of patches centered at each vertex v . At v , we place a repulsive force with limited range, which can apply to every other particle in the fabric except v 's neighbors. This places spheres about all vertices, which collectively cover the sheet. While these forces tend to make the sheet a bit thick for coarse grids, they work exceptionally well for fine grids ($> 30 \times 30$). We are currently working on other force configurations that will yield thinner fabric for coarse grids. Using a hierarchical decomposition of the fabric can dramatically improve the efficiency of the algorithm.

We describe the various force laws we used for repulsion, including a linear spring force and a Lennard-Jones-type 2/4 potential law, each of which we modify by cutting off their attractive components. In addition, we describe our use of adaptive step sizes to counter instabilities in the system.

Each of these force laws has its own benefits and disadvantages. The stiff spring force is simple, and its stability is easy to calculate and predict. On the other hand, a vertex with enough energy can overcome a stiff spring. In such an event, we are forced either to re-compute the entire animation with a new spring constant and step size, or to introduce a barrier beyond which the particle cannot pass (a sphere within the force sphere). In either case, we get a repulsive force for friction calculations.

A Lennard-Jones type force law gives us a repulsive force that more accurately models two pieces of cloth in contact. The force gives a little upon initial contact and then rapidly increases until the pieces of cloth can come no closer. The Lennard-Jones 6/12 potential law yields just such a force, but we found it to be too unstable for our purposes and decided to work with a similar force whose repulsive term is not quite as extreme as d^{-13} . We counter any remaining instability with an adaptive integration step size. This approach is best suited to situations where cloth-cloth collisions are rare, as when modeling clothing or cloth wrapped around stationary objects.

References

1. D. Baraff and A. Witkin. Large Steps in Cloth Simulation, Computer Graphics Proceedings, Annual Conference Series, pp. 137-146, 1996.
2. D. Hutchinson, M. Preston, W.T. Hewitt. Adaptive Refinement for Mass/Spring Simulations, Proceedings of Seventh Eurographics Workshop on Animation and Simulation, Poitiers, 1996.
3. Xavier Provot. Deformation Constraints in a Mass-Spring Model to Describe Rigid Cloth Behavior, Graphics Interface '97, pp. 147-155, 1997.
4. R. Szeliski and D. Tonnesen. Surface Modeling with Oriented Particle Systems, Computer Graphics (Proc. SIGGRAPH '92), pp. 185-194, July 1992.
5. P. Bolino, M. Courchesne, N. Magnenat-Thalmann. Versatile and Efficient Techniques for Simulating Cloth and Other Deformable Objects, Computer Graphics (Proc. SIGGRAPH '95), pp. 137-144. 1995.

Teresa Larsen
David Goodsell
The Scripps Research Institute
Tlarsenphd@home.com
www.scripps.edu/pub/olson-
web/people/larsen/HIVmodel.htm

Dru Clark
Michael Bailey
San Diego Supercomputer Center

Ernest Stewart
Matrix Enterprises



The digital model (center) provided the geometry to manufacture the physical components. Clockwise from lower left: the laminate block before liberating pieces, the individual pieces before coating and sanding, coated parts, and the laminated model, painted.

Modeling HIV

In the effort to conquer AIDS, scientists have collected more data about the Human Immunodeficiency Virus (HIV) than any other organism. Yet textbooks, journal publications, and other educational materials still use schematic representations to illustrate HIV.

In 1995, collaborating with David S. Goodsell, using in-house tools of Arthur Olson's Molecular Graphics Lab at The Scripps Research Institute, and employing sophisticated graphics tools at the Advanced Scientific Visualization Laboratory at the San Diego Supercomputer Center, Teresa Larsen and David Goodsell published a scientifically accurate model of HIV based on the data available in the literature. It incorporated the known structural features and used color, transparency, and texture mapping to illustrate the details of the model known at the time. The model renounces the traditional simplification of molecular components customarily used in schematic illustrations, such as spiral ribbons to represent the RNA and spheres for reverse transcriptase (RT) and other components. The renderings provide contextual relationships between the virus and its host in vivo and provide a visual basis for explanations of its life cycle and functional mechanisms.

Since 1995, a considerable amount of new data has come to light, and the sketch has been updated. We have now created an accurate physical model of HIV at a scale of one million times actual size. We imported each virtual component of the model as a single geometry file into the SDSC Tele-Manufacturing Facility's Laminated Object Manufacturer (LOM). As the LOM machine operates much like a pen-plotting device, proper interpretation of the 3D data requires that each piece consist of a contiguous surface of polygons. Any flaws or irregularities in the geometry can cause unpredictable results as the laser carves out each layer of the laminate to the geometry specifications.

We can now put this physical model into a researcher's or a student's hands so they can assemble and disassemble its various components and learn their structural interrelationships. The laminated model from the LOM machine, as well as digital versions of the geometry, serve as originals from which a manufacturer can create molds for mass producing the pieces. In collaboration with Matrix Enterprises, we have learned some practical limitations to maintaining the level of accuracy in the physical model portrayed in the virtual model. Consequently, the manufacturing engineers at Matrix Enterprises have provided guidance in tailoring the components to the particular requirements of the molding process.

In this physical model, we have chosen textures and colors of polymer materials to distinguish structural components. In each case, the choices also reflect durability and assembly criteria.

HIV

We define a new operator, called the morphological cross-dissolve, that blends binary or grayscale images using mathematical image morphology.² The operator gradually changes shapes in the source image until they match corresponding shapes in the target image. In the case of gray-scale images, the morphological cross-dissolve also changes gray levels while distorting shapes and achieves more natural transitions than traditional cross-dissolves. The new operator is completely automatic, requiring no user-supplied feature lines or correspondences.

To illustrate, consider the problem of gradually transforming one binary image into another, as shown in the top row of Figure 2. Each point x that is "on" in the source image but "off" in the target image must be turned off at some step in the transformation and vice-versa. By appropriately timing each transition, we create the illusion of a continuously moving boundary between the shapes.

In Figure 1, the solid and dashed curves represent the boundaries of source and target shapes, respectively; here, point x will eventually be turned off, and point y will eventually be turned on. The morphological cross-dissolve times the transition of each point using the distance to the source boundary and the distance to the target boundary. In particular, the transition time is the ratio of the former to the sum of the distances.

For discrete images, distance information is readily obtained from the *morphological dilation* operator.² The distance from the boundary of a set S to a point $x \notin S$ is the minimum number of dilations needed for S to reach x . The distance from the boundary of S to a point $x \in S$ is the minimum number of dilations needed for the complement of S to reach x .

So that expanding boundaries do not pass through all gaps between the shapes, distances in the expanding or contracting regions are constrained to lie within those regions. When the source and target shapes do not overlap, the transitions are timed so that the source shapes gradually erode away while the target shapes simultaneously undergo the reverse process.

To extend the technique to N -level gray-scale images, we partition the image G into $N-1$ binary level sets G_i , where $G_i(x)$ is "on" if and only if $G(x) > i$. We compute morphological cross-dissolves for each level set independently and recombine them.

Figure 2 compares a morphological cross-dissolve to a standard cross-dissolve. While both methods are fully automatic, only the former shows smooth transitions of boundaries rather than fades. For image pairs in which corresponding features overlap when superimposed, the results are comparable to feature-based morphs.¹

www.cs.caltech.edu/~arvo/mdissolve




Figure 1
The geometry of a morphological cross-dissolve.



Figure 2
(top) Several steps in the morphological cross-dissolve of two binary images. (middle) A morphological cross-dissolve of two gray-scale images. (bottom) An ordinary cross-dissolve of the same images. Differences are most evident on the left side of the face and the hair line. (Images from the database of faces, AT&T Laboratories, Cambridge.)

References

1. Thaddeus Beier and Shawn Neely. Feature-Based Image Metamorphosis, *Computer Graphics*, 26(2):35-42, July 1992.
2. J. Serra. *Image Analysis and Mathematical Morphology*, Academic Press, New York, 1982.

Morphological

Motivation

Rose, Bodenheimer, and Cohen describe a promising method for reusing a small number of example animations in an interactive system.¹ They use multi-dimensional interpolation of orientations using radial basis functions (RBFs) on an Euler angle parameterization. Our work is similar to this work, but our system uses quaternions as the representation of orientation, avoiding the known problems with interpolating Euler angles. Our goal was to discover a multi-variate version of Shoemaker's famous "slerp" function. (See Watt & Watt² for an introduction to quaternions and quaternion interpolation.)

Euclidean Multidimensional Interpolation

The problem of multivariate interpolation in Euclidean space consists of finding a continuous, smooth function $F(x) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ from a set of example datapoints $\{(x_i, y_i)\}$ that the function must pass through (formally, $F(x_i) = y_i$). RBFs are one technique for finding this function.

Multidimensional Quaternion Interpolation

We explored two methods for applying existing Euclidean multi-variate interpolation techniques to the (non-Euclidean) quaternion manifold: extrinsic and intrinsic. Our input examples will be of the form $\{x_i \in \mathbb{R}^n\}$ and our output (quaternion) examples of the form $\{q_i \in \mathbb{Q}\}$ where \mathbb{Q} is the unit quaternion group. As with *slerp*, the examples must all lie on the same (local) hemisphere of \mathbb{Q} . A simple preprocessing step is used to enforce this.

Euclidean Reprojection

A simple trick is to interpolate the examples as if they were Euclidean points in \mathbb{R}^4 , then project the result on \mathbb{Q} (by renormalizing). The problem with this method is that a constant speed change in the input parameters will not produce a constant speed interpolation of the reprojected points, which was the motivation for *slerp* originally.

Tangent Space Interpolation

Quaternions have an exponential form: $q = [\cos(\theta/2) \sin(\theta/2)\hat{n}] = e^{\frac{\theta}{2}\hat{n}}$ where q represents a rotation of θ around axis \hat{n} . Therefore, $\ln q = \frac{\theta}{2}\hat{n}$. The log of a quaternion lies in the tangent space anchored at the identity element. This tangent space is \mathbb{R}^3 , a Euclidean vector space. Additionally, the \ln of our examples will also lie in a local hemisphere of the tangent space, within the solid ball of radius π . (Note that tangent points need to be within $\pi/2$ of each other, however, since antipodal points on the surface of this ball are identified.)

We create our interpolation function F from the tangent space description (\ln) of the examples. We interpolate in this space and then "lift" the interpolated tangent vector back into \mathbb{Q} by exponentiating. Explicitly, the interpolated vector is $q(x) = e^{F(x, \{\ln x_i, q_i\})}$.

This function clearly interpolates examples, and the author has proven that it reduces to *slerp* in the case of one input parameter and two examples, as desired. A more technical description of this technique is available on the author's Web site.³

Results and Future Work

We have implemented both of these methods (using RBFs) to continuously interpolate between several kinematic animation examples that differ along emotional input axes, such as happiness and fatigue. The extrinsic version was incorporated in the Swamped! interactive exhibit at SIGGRAPH 98, and the intrinsic version is used in the Millennium Motel (SIGGRAPH 99). Future work will add quaternion-based inverse kinematics and constraints to this animation technique.

References

1. Rose, C., Bodenheimer, B., Cohen. M. Verbs and Adverbs: Multidimensional Motion Interpolation, IEEE Computer Graphics and Applications, Sep/Oct 1998.
2. Watt, A. and Watt, M. Advanced Animation and Rendering Techniques. ACM Press. 1992. pp. 356-368.
3. www.media.mit.edu/~aries/quats/quats.html

Multifluid Finite Volume Navier-Stokes Solutions for Realistic Fluid Animation

John A. Turner
Andre C. Mazzone
Blue Sky Studios, Inc.

It is commonly believed that solution of the full 3D Navier-Stokes equations is too computationally intensive for computer graphics applications. Previous approaches have typically used either the simplified shallow-water approximations² or, most recently, two-stage approaches involving a low-resolution 3D Navier-Stokes solution followed by a height-field solution.¹ However, recent advances in incompressible free-surface-flow algorithms and methods for the solution of linear systems of equations make high-resolution solution of the full 3D Navier-Stokes equations possible.

A computational tool, referred to here as MIZU* has been developed at Los Alamos National Laboratory for simulation of casting processes (filling, cooling, and solidification of molten fluid in molds with complex geometry**). Such simulation involves modeling physical phenomena such as unsteady, incompressible, or slightly compressible flow of multiple, immiscible fluids; interface physics (for example, surface tension); convective, diffusive, and radiative heat transfer; solidification of multi-component alloy systems; microstructural physics (for example, nucleation, dendrite growth); and material response effects (for example, stress, distortion, shrinkage, plastic flow).

MIZU solves the 3D, incompressible, variable-density Navier-Stokes equations on generalized-connectivity unstructured (GU) meshes. The use of GU meshes, which can contain hexahedra, tetrahedra, prisms, and pyramids, enables simulation of arbitrarily complex geometry, which is crucial for both mold filling and computer-graphics applications. Note that these are volumetric meshes, so the availability of tools for generating high-quality GU meshes from surface meshes created by CAD software and modeling tools such as Maya and Softimage is essential.

The flow algorithm is a 3D extension of recent advances in projection methods⁴ and is based on a colocated, cell-centered, finite volume formulation that is second-order in both time and space.*** The algorithm is implicit, so it requires the solution of linear systems of equations at each timestep. The solver library used⁵ makes use of recent advances in iterative-solution techniques, specifically preconditioned Krylov subspace methods.

Since the flow involves multiple materials, a critical aspect of the simulation involves tracking the interface between materials. Note that the flow of wine into a glass is a multimaterial problem (wine and air) and that material properties such as density vary by several orders of magnitude over an extremely small distance at the interface. In addition, the interface itself is topologically complex, and physical properties such as surface tension play a significant role in the behavior of the flow. Hence, realistic fluid simulation requires accurate modeling of these interfaces. MIZU uses a volume tracking,

or volume-of-fluid (VOF), approach. Interfaces are tracked on 3D generalized hexahedral cells and localized over one cell width at each timestep. They are assumed to be planar within each cell, yielding a globally piecewise planar approximation to the actual interface.

Once the flow conditions are computed, the results are rendered as spherical metaballs (blobs), with blob radius determined by the volume fraction of the appropriate fluid. While this smears some of the fine detail of the solution, and would likely not be appropriate for scientific visualization, it results in motion realistic enough for film and commercial applications.

The results were rendered using Blue Sky Studios, Inc.'s proprietary ray-tracing renderer, CGI Studio. Figure 1 shows a selected frame from an animation of a fluid being poured into a glass box.

As an initial test, we simulated a fluid being poured into a glass box eight centimeters on each side. Since the geometry of this situation is simple, a 32x32x32 orthogonal mesh was used, resulting in 32,768 computational cells of 0.25 cm on each side. The calculation thus requires solution of linear systems involving this number of unknowns at each timestep. While this mesh is somewhat coarse, it nevertheless yields fairly realistic results.

The simulation was mostly performed using 250 MHz R10000 processors on an SGI Origin 2000 (although MIZU provides for parallel execution in either SMP or distributed modes via MPI, these calculations were performed in serial mode). CPU time requirements were significant, but not unreasonable. Early in the simulation, when the flow is complex and difficult to compute, roughly an hour of CPU time was required for each frame of animation. Near the end of the simulation, when the flow is nearing steady-state, only a few minutes per frame were required.

The resulting animation sequences are quite realistic, particularly the degree to which the complex topology of the interface is captured. A selected frame is shown in Figure 1.

Future Improvements

Currently, blobs are placed at the centroid of each cell. More realistic results could be achieved if they were placed at the centroid of the region of the cell occupied by the fluid being rendered. Since a planar interface is reconstructed in the course of the flow calculation, this information is available.

In addition, blobs need not be spherical. Ellipsoids that "fit" the cell region in question, using cell geometry and the reconstructed fluid interface, can be used. This would further enhance the realism of the final animation.



Figure 1
Simulation time: 0.77 seconds.

References

1. N. Foster and D. Metaxas. Realistic Animation of Liquids, Proceedings GI, pp. 204-212, 1996.
2. M. Kass and G. Miller. Rapid, Stable Fluid Dynamics for Computer Graphics, Proceedings of SIGGRAPH 90, in Computer Graphics, volume 24, pp. 49-57, 1990.
3. W. J. Rider, D. B. Kothe, S. J. Mosso, J. H. Cerruti & J. I. Hochstein. Accurate Solution Algorithms for Incompressible Multiphase Fluid Flows, Technical Report AIAA 95-0699, AIAA, 1995. Presented at the 33rd Aerospace Sciences Meeting and Exhibit.
4. W. J. Rider, D. B. Kothe, E. G. Puckett & I. D. Aleinov. Accurate and Robust Methods for Variable Density Incompressible Flows with Discontinuities, In V. Venkatakrisnan, M. D. Salas & S. R. Chakravarthy, editors, Workshop on Barriers and Challenges in Computational Fluid Dynamics, pp. 213-230, Boston, MA, 1998. Kluwer Academic Publishers.
5. J. A. Turner, R. C. Ferrell & D. B. Kothe. JTpack90: A Parallel, Object-Based Fortran 90 Linear Algebra Package, Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing, Minneapolis, MN, March 14-17, 1997.

* Japanese for water.

** It must be noted that MIZU has not been released publicly and is in a state of development. It was only available to us due to the fact that one of us (Turner) was involved in its development while at LANL. More information on MIZU can be found at www.zephyr-group.com/mizu/

*** That is, if the mesh spacing is decreased by a factor of two, accuracy improves by a factor of four.³

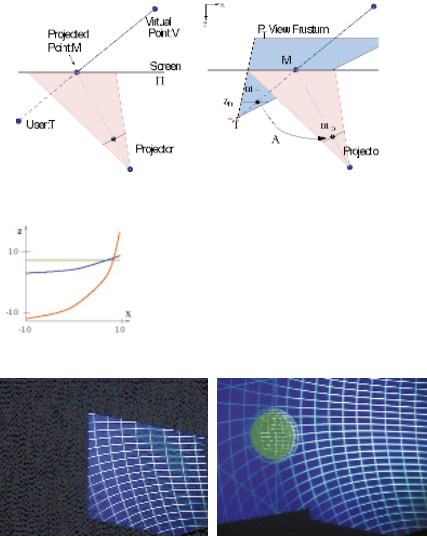


Figure 1: (a) Oblique projectors create keystone imagery. (b) The modified projection matrix achieves off-axis projection P_T and a collineation $A_{4 \times 4}$. (c) Depth buffer values along a scan line for points along constant depth. Using P_T (green). After collineation (red) and with depth-buffer approximation (blue). (d) An oblique projector (e) and its contribution to overlapped projectors.

$$A_{4 \times 4} = \begin{bmatrix} a_{11} & a_{12} & 0 & a_{13} \\ a_{21} & a_{22} & 0 & a_{23} \\ 0 & 0 & 1 & 0 \\ a_{31} & a_{32} & 0 & 1 \end{bmatrix} \quad (1)$$

$$A'_{4 \times 4} = \begin{bmatrix} a_{11} & a_{12} & 0 & a_{13} \\ a_{21} & a_{22} & 0 & a_{23} \\ 0 & 0 & 1 - |a_{31}| - |a_{32}| & 0 \\ a_{31} & a_{32} & 0 & 1 \end{bmatrix} \quad (2)$$

Projectors are typically mounted so that their optical axis is perpendicular to the planar display surface. Such configurations also are used in immersive environments to render perspective-correct imagery for a head-tracked moving user. These environments include CAVEs, PowerWalls, and ImmersaDesks (back-lit and tilted desktop workbenches). By design, typical display systems try to maintain the image plane parallel to the plane of the display surface. However, this leads to a need for constant electro-mechanical alignment and calibration of display systems.

In this sketch, we show that it is possible to allow oblique projection on planar display surfaces and still render correct imagery of 3D scenes without additional computational cost. We describe how oblique projection can replace frequent alignment with simple calibration. We use a traditional graphics pipeline with a modified projection matrix and an approximation of the depth buffer.

Oblique Projection

Consider rendering the virtual point V for the user at T using an oblique projector as shown in Figure (a). For planar display surfaces, the images m_T and m_P of the virtual point V can be computed by first finding projection of V onto the display surface, M . A simple observation is that the two images of a common virtual point are related by a collineation, which is well known to be a 3×3 matrix defined up to scale. This observation allows us to create a new projection matrix during rendering for the projector as a product of a traditional off-axis projection matrix, P_T (from the user's viewpoint) and a matrix, $A_{4 \times 4}$ (from the 3×3 collineation matrix).

Without a loss of generality, let us assume that the display plane, Π , is defined by $z = 0$. There are various ways to create P_T and $A_{4 \times 4}$. We will use a method that updates P_T as the user moves but the collineation matrix remains constant. We create an axis aligned rectangle S on Π bounding the keystone quadrilateral illuminated by the projector. Define a view frustum by first creating a pyramid with T and the four corners of S and then truncating it with a near plane, $z = T_Z - z_n$, and a far plane, $z = T_Z - z_f$. This is similar to OpenGL's *glFrustum* setup. The projection matrix for this view frustum is, $P_T = Frustum(T, S, z_n, z_f) Translate(-T)$.

Next we calculate the collineation between images of V : m_T due to P_T , and its image in the projector, m_P . If the 3D positions of points on Π illuminated by four or more pixels of the projector are known, the eight parameters of the collineation matrix, $A = [a_{11}, a_{12}, a_{13}; a_{21}, a_{22}, a_{23}; a_{31}, a_{32}, 1]$, can be easily calculated. We create a new matrix, $A_{4 \times 4}$ to transform the pixel coordinates but try to keep the depth values intact. The complete projection matrix is $A_{4 \times 4} P_T$. See equation (1).

Depth Buffer Approximation

Although the naive approach described above creates correct images of virtual 3D points, it is important to note that the traditional depth buffer cannot be effectively used for visibility and clipping. The depth values of virtual points between near and far planes due to P_T are mapped to $[-1, 1]$. Let's say, $[m_{T_X}, m_{T_Y}, m_{T_Z}]^T = P_T [V, 1]^T$ and $m_{T_Z} / m_{T_W} \in [-1, 1]$. After collineation, the new depth value is actually $m_{T_Z} / (a_{31} m_{T_X} + a_{32} m_{T_Y} + m_{T_W})$ which (i) may not be in $[-1, 1]$, resulting in undesirable clipping and (ii) is a function of pixel coordinates, changes quadratically, and hence cannot be linearly interpolated during scan conversion for visibility computation (Figure (c)). In general, we cannot achieve two hyperbolic interpolations for the depth values with a single 4×4 matrix. In other words, we must first compute an image with P_T ("divide by w' "), and then warp the resultant image. This requires a two-pass rendering method: first render the image and load it in texture memory and then achieve warping using texture mapping. However, we can achieve the rendering and warping in a single pass using an approximation of the depth buffer. Note that $m_{T_Z} / m_{T_X} / m_{T_W}$ and $m_{T_Y} / m_{T_W} \in [-1, 1]$ for points rendered inside the rectangle S . Hence $(1 - |a_{31}| - |a_{32}|) m_{T_Z} / (a_{31} m_{T_X} + a_{32} m_{T_Y} + m_{T_W})$ is guaranteed to be in $[-1, 1]$. Further, by construction of P_T , the angle between the projector's optical axis and the normal of the planar surface is the same as the angle between the optical axis and the retinal plane of the frustum for P_T . Thus, if this angle is small (i.e. $|a_{31}|$ and $|a_{32}| \ll 1$), the depth values are modified but the changes are monotonic and almost linear across the framebuffer as shown in Figure (c). See equation (2).

Applications

The modified projection matrix can be easily calculated by measuring the tracker-sensor at the four corners of the illuminated quadrilateral. The effect of quadratic changes in depth values is minimized when the projector is almost perpendicular. In CAVEs or ImmersaDesks, special effort is taken to align projector pixels to the pre-defined corners. Using the technique described, a rough positioning followed by a simple calibration is sufficient to render correct images.

Acknowledgments

I would like to thank Gary Bishop and my colleagues in the Office of the Future group at the University of North Carolina at Chapel Hill (Mike Brown, Ruigang Yang, Wei-Chao Chen, Herman Towles, Greg Welch, Brent Seales and Henry Fuchs) for helpful discussions and prototype implementation.

One approach to occlusion culling in z-buffer systems, by which we mean the culling of occluded geometry prior to rasterization, is to replace the z-buffer with a z-pyramid and perform hierarchical z-buffering^{1,2} rather than traditional z-buffer scan conversion. However, makers of z-buffer hardware have been slow to adopt this approach because it requires a major architectural revision. Here we describe a simpler and cheaper alternative, adding to a conventional z-buffer pipeline a culling stage that employs hierarchical z-buffering optimized for conservative culling (culling that sometimes fails to cull occluded geometry but never culls visible geometry).

In this architecture (Figure 1), a culling stage is inserted between the transformation and rendering stages of a conventional z-buffer pipeline. The culling stage receives transformed primitives, hierarchically tiles them into a z-pyramid,² thereby culling most occluded primitives, and passes visible and nearly visible primitives on to a conventional rendering stage. The rendering stage establishes visibility definitively at each image sample using a standard z-buffer.

One big advantage of this architecture is that it enables hierarchical z-buffering to be highly optimized for conservative culling. Toward this end, we employ a very compact z-pyramid to reduce storage cost, memory traffic, and tiling computations. Within a standard z-pyramid having 4x4 decimation, 4x4 tiles are stored as arrays of full-precision z values. Instead, we use 16-bit z values, and as diagrammed in Figure 2, at the finest pyramid level we represent 4x4 tiles as two z values and a coverage mask: a zfar value for the tile (zfar_tile), a coverage mask indicating samples covered by primitives encountered since the last update of zfar_tile, and the zfar value of the samples covered by the (zfar_mask). Since zfar_mask is at or behind the corresponding occluders, culling performed with this data structure is conservative.

This compact z-pyramid requires only about 10 percent of the storage of a standard z-pyramid and culls approximately 90 percent of occluded primitives. Simulations on a variety of simple and complex scenes show that hierarchical tiling into this z-pyramid generates only about 10 percent as many reads and writes as standard hierarchical z-buffering, and only 1-2 percent as many as traditional z-buffering. These figures assume that primitives are traversed in approximately front-to-back order, which can be accomplished by organizing the scene in bounding boxes and traversing the boxes front to back.

To illustrate the dramatic bandwidth reductions that are possible with this architecture, we rendered the scene of Figure 3, which consists of 167 million replicated polygons organized into bounding boxes. With front-to-back traversal of bounding boxes, the culling stage reduced the depth complexity of polygons processed by the rendering stage from 41.5 to 2.3, an 18-fold reduction. To reduce geometry traffic in the pipeline, we tested bounding boxes for visibility with the culling stage prior to sending their primitives through the pipeline.¹ In rendering Figure 3, the culling stage generated only .58 reads and writes of z-pyramid values per image sample, on average.

Bandwidth simulations on this and other scenes show that this occlusion-culling architecture could extend the real-time rendering performance of z-buffer hardware to much more deeply occluded scenes.




Figure 2. Side view of a finest-level 4x4 tile within the z-pyramid. To reduce memory traffic and storage, finest-level tiles are represented as two z values and a coverage mask.




Figure 3. For this scene, the culling stage reduces the depth complexity of polygons that need to be rendered from 41.5 to 2.3.

References

1. N. Greene, M. Kass, and G. Miller. Hierarchical Z-Buffer Visibility, Proceedings of ACM SIGGRAPH 93, 231-238.
2. N. Greene. Hierarchical Polygon Tiling with Coverage Masks, Proceedings of ACM SIGGRAPH 96, 65-74.




Figure 1. Block diagram of the proposed architecture. Most occluded geometry is culled with a separate culling stage that is highly optimized for conservative culling.

OpenGL Texture-Mapping With Very Large Datasets and Multi-Resolution Tiles

On OpenGL systems without the "clipmap" extension, a tiled approach can be used to provide the correct resolution image data for different parts of a 3D visualization, where very large texture maps are desired. The application, called "EarthView" is a planetary browser, which can be used to view arbitrary amounts of geo-referenced image and terrain elevation data. The proper number and size of bi-linear texture-maps that can fit into available memory are used as a resource, constantly recycled by being "sub-loaded" with new data as the user browses into new areas of the globe.



EarthView uses recursive descent, window and horizon culling, and a quad-tree structure to focus in on the part of the globe currently in view. However, all vertex and texture coordinates are not pre-stored, but rather calculated on-the-fly, based on a simple cosine table. This allows EarthView to support close-in zooms to any part of the globe, and high-resolution insets can be supplied for desired areas. Terrain elevation data are also utilized on a multi-resolution basis, which allows the terrain to "morph" as it is viewed from different positions. This allows EarthView to browse at very high frame rates.

EarthView mimics the texture-mapping calculations performed in OpenGL in order to decide which image resolution tile to use for a particular area. This creates a high-quality result similar to what could be accomplished with a single, very large, mipmapped texture, and thus avoids memory limitations. EarthView can browse very large datasets even on systems with less texture memory, and browsing of the same areas repeatedly is enhanced on systems with larger memory.

EarthView has an unusual projection for polar regions, designed to save disk space, look better, and run faster. Typical cylindrical projections stretch the polar regions, creating an imbalance between the amount of data corresponding to longitude versus latitude. Also, polar geometry is tessellated differently, avoiding the typically thin "pie-slices" at the poles. Various OpenGL and window-system features and extensions are used to enhance the view and speed up performance, including texture objects and subloads, the "clamp-to-edge" extension, the "packed pixels" extension, the "texture-LOD" extension, vertex arrays, and asynchronous I/O.

The EarthView application was designed to demonstrate the advantages of a new Intel-based hardware platform called the SGI Visual Workstations 320 and 540. A 320 system can be configured with 1Gb of system memory, and with the Integrated Visual Computing (IVC) architecture, more than 900 Mb of this can be allotted to textures, an unprecedented amount. This, combined with a high textured fill rate, a high textured polygon rate, and the ability to move data in from the disk in parallel with other operations, gives outstanding results.

For the first time, a system costing less than \$5K can browse planetary data with speed and visual realism that was previously available only to government and industry personnel on very expensive systems. As a dedicated platform, this alone would be well worth the price; when not running the browser, users still have a general-purpose, high-performance graphics workstation.

Purpose

Since its presentation in 1973, Phong shading^{1,2} has been the standard technique for rapidly rendering specular highlights. Because it supports interior highlights, Phong shading retains a high degree of visual quality even with low-polygon-count models. However, despite 25 years of improvements,^{3,4} real-time Phong shading is so costly that it's still only available on high-end hardware. This sketch presents a new implementation, "Rapid Phong," that is finally fast enough in both setup and execution to be easily implemented in real-time software and low-end hardware. By using a few matrix rotations during polygon setup, this algorithm reduces Phong's per-pixel renormalization and exponentiation to a single per-pixel linear interpolation and fast multiply.

Basic Phong Shading

Basic Phong-style shading ($\vec{R} \cdot \vec{V}$)ⁿ consists of four main steps per pixel:

1. Linearly interpolate the 3D-reflection vector (R'_x, R'_y, R'_z) across the polygon.
2. Renormalize the interpolated reflection vector \vec{R}' per pixel.
3. Calculate the highlight by dotting the reflection and view vector $\vec{R}' \cdot \vec{V}$.
4. Sharpen the specular falloff by using a specular exponent (n).

New Rapid Phong Implementation

Per triangle, we will transform the reflection vectors to eliminate the per-pixel highlight dot product, remove \vec{R}'_z interpolation and \vec{R}' renormalization using the Pythagorean theorem, simplify the specular falloff, and then roll to remove the \vec{R}'_y per-pixel interpolation.

View-Transform the Reflection: Transform the view and reflection vectors into view (eye) space, so that the new $\vec{V}' = (0,0,1)$. Then $\vec{R}' \cdot \vec{V}$ becomes $\vec{R}' \cdot (0,0,1)$, reducing the per-pixel highlight dot product to the single component \vec{R}'_z . The regular view transform can be used, since it's usually orthonormal. (Note for later that the roll around this "z" axis is arbitrary).

Remove \vec{R}'_z interpolation and \vec{R}' renormalization: Since $\vec{R}'_x{}^2 + \vec{R}'_y{}^2 + \vec{R}'_z{}^2 = 1$, only interpolate the "x" and "y" components of the transformed reflection vectors and extract \vec{R}'_z via the Pythagorean theorem. By definition, \vec{R}'_z will already be normalized. For speed, generate the \vec{R}'_z highlight via $1 - (\vec{R}'_x{}^2 + \vec{R}'_y{}^2)$, removing the square root. This also reduces the up-to-30-percent linear interpolation error to under six percent.

Simplify the specular falloff: A good approximation for the specular falloff exponentiation is to premultiply the reflection vectors by the square root of the specular exponent (n). The resulting highlight value of $1 - n * (\vec{R}'_x{}^2 + \vec{R}'_y{}^2)$ stays accurate down to 50-percent brightness and only needs a clamp to zero per pixel.


Remove \vec{R}'_y per-pixel interpolation: Since $1 - (\vec{R}'_x{}^2 + \vec{R}'_y{}^2)$ is circularly symmetric, rotate the 2D (R'_x, R'_y) reflection components into (R''_x, R''_y), aligning them with the screen's x rows and y columns. As a result, \vec{R}'_y stays constant within each row. (As noted in the view-transform step, the roll around "z" was arbitrary).

Per-Pixel Calculations:

We obtain the per pixel highlight value from $1 - (\vec{R}''_x{}^2 + \vec{R}''_y{}^2)$. Since $\vec{R}''_y{}^2$ is constant within each row, only \vec{R}''_x needs interpolating per pixel. Its squaring can be changed to an add via finite differences. The specular falloff was already included earlier, but still requires a clamp to zero. The final cost is only 3 adds and one clamp per pixel.

Finally

The reasonable setup overhead and low per-pixel cost of Rapid Phong allows it to run at speeds approaching Gouraud. Hopefully, refinements like this will finally permit the higher-quality rendering of Phong shading to become widely available, even on real-time consumer level systems.



Speed

References

1. Phong, B. T. Illumination For Computer-Generated Images, PhD Dissertation, Department of Computer Science, University of Utah, Salt Lake City, 1973.
2. Foley, J. D. and A. Van Dam. Fundamentals of Interactive Computer Graphics, Addison Wesley, Reading, MA., 1983.
3. Bishop, Gary and Weimer, David. Fast Phong shading, ACM SIGGRAPH 96 Conference Proceedings.
4. Schlick, Christophe. A Fast Alternative to Phong's Specular Model, Graphics Gems IV, Academic Press, Inc., 1994.

Physically Based Anatomic Modeling for Construction of Musculoskeletal Systems

In recent years, we have seen several anatomically inspired modeling systems, such as those created by Scheepers et al.³ and Wilhelms and van Gelder.⁴ Although these modeling systems may be suitable for defining superficial body shape, the choice of shape primitives is not sufficiently accurate for medical applications because they do not capture details of muscle fiber architecture nor allow specification of complex attachment areas of musculotendon to bone. Furthermore, these muscles are incapable of generating forces for physical simulation.

We are developing an anatomically based modeling system that integrates both the physical and geometric properties of muscle and tendon for the purpose of constructing and simulating musculoskeletal systems. Our goal is to create a muscle model that can be embedded with physical properties to enable animation of active muscle contraction and other inertial effects as the muscles move with their underlying bones. Essentially, we are incorporating low-level, physically based muscle models, such as those introduced by Chen and Zeltzer¹ with the primarily geometric, anatomically based modeling systems mentioned above.

Muscle Model

We use B-spline solids to represent muscle and tendon. The B-spline solid is the volumetric triple tensor product function of three material coordinates, u , v , and w whose domain maps to the spatial volume:

$$x(u, v, w) = \sum_{i=0}^l \sum_{j=0}^m \sum_{k=0}^n B_i^u(u) B_j^v(v) B_k^w(w) q_{ijk}$$

A 3D spatial point $x(u, v, w)$ in the solid is produced by the weighted sum of control points $q_{\{ijk\}}$ and the product of B-spline basis functions. The volumetric representation allows individual muscle fibers to be visualized as iso-parametric curves embedded in the solid. It also enables an efficient point-inclusion test that can be used for collision detection with other muscles and bone geometry.

The control points are mapped to a unique set of mass points embedded within the solid, allowing physical reaction constraints² and internal muscle forces to be applied at these mass points.

We use a Lagrangian formulation for the equations of motion of the solid to enforce global volume conservation simultaneously with the local attachment and non-interpenetration reaction constraints and the internal muscle force models.

Attachment of Musculotendon to Bone

We use these muscle models as force actuators in a simulation framework that allows muscles to be visualized and simulated with articulated skeletons made up of rigid bones and various joint constraints. To produce an initial shape for muscle, pre-built muscles from an anatomical library can be loaded into the system or the muscles can be interactively created by sketching profile curves directly onto the bone surfaces by direct manipulation.

These curves can be subsequently adjusted while maintaining bone attachment to provide high-level shape manipulators for a practitioner who may wish to specify non-trivial attachment areas or adjust the various lines of action of muscle.

Conclusion

The combination of geometric and physically based properties in a muscle model allows closer study of the interrelationships between anatomical form and function. Virtual reconstructions of both real and non-existent musculoskeletal systems can be created and simulated.

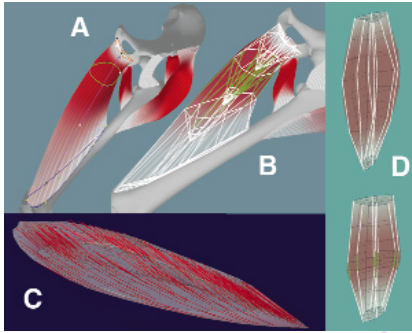


Figure 1

(a) Profile curves are used to specify muscle attachment areas and lines of action. (b) Green viscoelastic units in muscle can be selectively contracted. (c) Generated fibers show correct orientation in isolated human soleus specimen. (d) Relaxed B-spline solid (top) has selected green fibers contracted while preserving volume (bottom).

References

1. David T. Chen and David Zeltzer. Pump It Up: Computer Animation of a Biomechanically Based Model of Muscle Using the Finite Element Method, Computer Graphics (SIGGRAPH 92 Proceedings), pp. 89-98, July 1992.
2. John C. Platt and Alan H. Barr. Constraint Methods for Flexible Models, Computer Graphics (SIGGRAPH 88 Proceedings), pp. 279-288, August 1988.
3. Ferdi Scheepers, Richard E. Parent, Wayne Carlson and Stephen F. May. Anatomy-Based Modeling of The Human Musculature, Computer Graphics (SIGGRAPH 97 Proceedings), pp. 163-172, August 1997.
4. Jane Wilhelms and Allen Van Gelder. Anatomically Based Modeling, Computer Graphics (SIGGRAPH 97 Proceedings), pp. 243-252, August 1997.

Motion blur arises naturally in cinematography as a result of movement in the scene while the camera's shutter is open. In its absence (for example, when a scene is illuminated by a strobe light), movement is perceived to lack fluidity and coherence. In traditional cel animation, motion blur is simulated with streaks ("speed lines") or by instancing objects multiple times in a drawing ("vibrations"),¹ but these techniques are labor-intensive and are practical only for localized motion (not for motion arising from a camera pan, for example). Algorithms for simulating motion blur in computer-generated animation have been available for some time,² and their effectiveness in reducing "strobing" and enhancing photorealism and the sensation of motion is well-recognized. Such algorithms typically take advantage of the known velocities of objects in a 3D scene by filtering along the objects' motion paths. In contrast, velocity information is generally not directly available in 2D cel animation. Fortunately, though, "optical flow" estimation (deriving velocity information from image sequences) has received considerable attention in the computer-vision community. This sketch describes a postprocess 2D motion blur algorithm based on optical-flow estimation.



An optical-flow field is a vector field that describes the motion in an image by assigning to each pixel of the image a vector representing its velocity at a given instant in time. Ideally, each pixel's velocity reflects the velocity of the point in the 3D scene of which the pixel is a projection, but this need not be the case (consider an image of a featureless, rotating sphere). In evaluating a number of different algorithms for estimation of optical flow, Anandan's correlation-based algorithm³ was found to perform well on animation cels. Correlation-based algorithms derive velocities from the displacements of matching neighborhoods of pixels in successive images, where the quality of a match is determined by cross-correlation of the intensity distributions in the neighborhoods.

Anandan's algorithm employs a multilevel Laplacian image pyramid,⁴ in which matching takes place between two images at multiple resolutions. Results from lower-resolution matches are used as initial guesses for higher-resolution matches, which greatly reduces the search space required to recover large displacements. To single-pixel accuracy, the velocity of a pixel in the first image is its displacement from the pixel in the second image for which the sum of squared differences (SSD) in intensity between their two neighborhoods is minimized. Subpixel accuracy is achieved by minimizing a quadratic approximation to the SSD surface over a search space centered at this integer displacement. The algorithm also imposes a smoothness constraint to minimize variations in velocity both between and within pyramid levels.

Simulated motion blur is applied to an image via line-integral convolution⁵ of its optical-flow field. Satisfactory results are obtained with a simple box filter kernel, whose length and center point can be adjusted to vary properties of the effect, such as overall blurriness. (Line-integral convolution integrates along streamlines of the optical flow field, and it should be noted that these do not in general correspond to motion paths of objects in the scene, but the resulting blur is visually more appealing than that produced by, for example, a DDA line-drawing kernel.) Where possible, motion blur is applied to individual scene elements independently, as this reduces unnecessary computation and allows for greater flexibility in compositing.

The accompanying figures illustrate the process: Figure 1 is a frame from the animated feature film "The Prince of Egypt" prior to motion blurring. Figure 2 is a visualization, by line-integral convolution with a random noise field, of the computed optical-flow fields for two scene elements. Figure 3 is the final frame with motion blur.

Thanks to Lance Williams for the inspiration for this work and for his invaluable input.



References

1. F. Thomas and O. Johnston. Disney animation: The Illusion of Life, New York, Abbeville Press, 1981, pp. 116-117.
2. M. Potmesil and I. Chakravarty. Modeling Motion Blur in Computer-Generated Images, Computer Graphics (SIGGRAPH 83 Proceedings) 17, 1983, pp. 389-399.
3. P. Anandan. A Computational Framework and An Algorithm For the Measurement of Visual Motion, International Journal of Computer Vision 2, 1989, pp. 283-310.
4. P.J. Burt and E.H. Adelson. The Laplacian Pyramid as a Compact Image Code, IEEE Transactions on Communications 31, 1983, pp. 532-540.
5. B. Cabral and L. Leedom. Imaging Vector Fields Using Line Integral Convolution, Computer Graphics (SIGGRAPH 93 Proceedings) 27, 1993, pp. 263-272.

Projecting Computer Graphics on Moving Surfaces: A Simple Calibration and Tracking Method

We are currently investigating the possibilities of transforming every surface in a space into a display and, in particular, of projecting graphics on movable surfaces. In this sketch, we present a simple method to convert tracking points from a camera into positions on the image to be projected. The main advantage of our method is the simplicity of a calibration process that does not require knowledge of the intrinsic parameters (focal length, dimensions of the imaging device) of the camera or the projector.

Camera and Projector Calibration

The relationship between points observed on a planar surface from two different cameras is known to be a *homography*.¹ A homography is a 3x3 matrix that defines a linear application in the projective space that, for a given planar surface of the real world, maps all projected points in one camera's image into the other camera's image.

The fundamental observation is that from a geometric point of view, "ideal" pinhole projectors and cameras are identical (see Figure 1). Let H denote the homography that relates the image of the projector image frame to the camera image frame. This means that a 2D-point on the camera image $\bar{c} = (x_c/z_c, y_c/z_c)$ matches a 2d point $\bar{p} = (x_p/z_p, y_p/z_p)$ on the projector image as follows:

$$p = \begin{bmatrix} x_p \\ y_p \\ z_p \end{bmatrix} = Hc = H \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix}$$

A homography is completely defined if the projection of four 3D points of the world on both image planes is known. To determine the homography between a camera and a projector, we need simply to obtain the four needed points while manually aligning a projection of the surface with the real surface (see Figure 1).

The homogeneous coordinates of the four points to be projected, $p_i = (x^i_p, y^i_p, 1)$, $i=1,2,3,4$, are determined arbitrarily, making sure that the points are visible, and there is a way to move the real surface so it aligns with the projection. Then, we consider the homogeneous coordinates of the four points on the camera image as sensed by the tracking system, $c_i = (x^i_c, y^i_c, 1)$, $i=1,2,3,4$. Taking the matrices corresponding to these two sets of four points, $P=(p_1^T p_2^T p_3^T p_4^T)$ and $C=(c_1^T c_2^T c_3^T c_4^T)$, we obtain $P=HC$, whose solution is

$$H=PC^T(CC^T)^{-1}$$

During run-time, we simply take a point in the camera image $c = (x_c, y_c, 1)$, project through the homography H obtaining $p=Hc = (x_p, y_p, z_p)$, and compute the position on the projector's image plane, $(x_p/z_p, y_p/z_p)$.

Surprisingly, this calibration step is numerically stable even with only four points and can be done, in practice, in a few seconds. We believe that the stability is also related to the fact that in our experiments the projection centers of the camera and the projector are close to being aligned. Notice that there is no need to determine either the camera's intrinsic parameters or the projector's.

Tracking the Projection Surface

In our experiments, we have used plain markers on the projection surface. In particular, we employed infrared LEDs that can be easily tracked by a camera with an infrared filter. However, if we move the mask too quickly, we observe that the projected image "falls behind" the moving surface. That is, there is a "shifting" effect where the observations at discrete time t on the camera image $c(t)$ are displayed by the projector at time $t+dt$ using the estimate at time t , $p(t)=Hc(t)$. To reduce the "shifting" problem, we employ a predictive Kalman filter² that estimates the most likely position of every point at time $t+dt$, using the equations of dynamics as the underlying model of the Kalman filter, as shown in Figure 1. The parameter dt , corresponding to the average delay between sensing and displaying, is determined experimentally. The Kalman filtering approach proved to be very effective in our experiments.

Figure 1
 Calibration process and run-time system.

References

1. O. Faugeras. Three-Dimensional Computer Vision: A Geometric Viewpoint, The MIT Press, Cambridge, Massachusetts. 1993.
2. A. Gelb (ed.). Applied Optimal Estimation, The MIT Press, Cambridge, Massachusetts. 1974.

Projector

This prototype allows for face-to-face communication between remote users.¹

Unlike most conventional tele-conferencing systems, which limit the users' viewpoint to a fixed position, a mutual telexistence system should provide images of the other users that correspond to the change of a user's eye position in the computer-generated 3D space. On the other hand, with "image-based rendering," complex photo-realistic images can be effectively synthesized at arbitrary viewpoints. However, with this technique, the system must synthesize the image in real-time for smooth communication. This prototype system includes the following features to fulfill this real-time request:

1. Geometric information such as a "depth map" is needless. Only one parameter of the object's distance is required. So, the system does not include a time-consuming process such as pattern matching. Data processed by the rendering computer are reduced in advance at the stage of capturing the source image. This means that the rendering computer doesn't have to deal with bulky data of the "plenoptic function."
2. Fast rendering is enabled by using graphic hardware acceleration of texture mapping. The system is comprised of relatively low-cost general graphic hardware.

The prototype system includes a camera unit, the rendering PC, and the control PC. In the camera unit, 12 small color CCD cameras are aligned horizontally in a row on the linear actuator at intervals of 54mm. These cameras rotate around their optical axes, so the direction of their scanning lines is vertical. Moreover, all the cameras are synchronized by the same genlock signal, and a video switch is installed between the camera unit and the rendering PC. This design allows the rendering PC to selectively capture only one scanning line among 12 video streams, which reduces data and the number of video capturing devices required. The rendering PC synthesizes the images of the object from arbitrary viewpoints by texture-mapping long tile images on a transparent plane in the computer-generated 3D space. The control PC indicates the channel of this switch and controls the motion of the linear actuator.

Figure 2. shows a video sequence synthesized by the system. Moving human figures are successfully rendered in real-time.

More specific technical data is available at: www.start.u-tokyo.ac.jp/projects/mutel/




Figure.1
Left: overview of the prototype system.
Right: block diagram of the prototype system.



Figure.2 Synthesized video sequence.

Reference

1. Susumu Tachi, Taro Maeda, Yasuyuki Yanagida, Masaaki Koyanagi, and Hiroki Yokoyama: A Method of Mutual Tele-Existence in a Virtual Environment. In Proceedings of ICAT 98, pp9-18, 1996.

Quasi-Linear Z Buffer

This sketch presents a new type of depth buffer with quasi-linear mapping from eye to screen space that significantly improves depth resolution and is easy to use with standard 3D hardware and APIs.

Perspective projection of eye-object distance to depth in screen space is usually bound by normalization to [0,1] range, preservation of lines and planes, and preservation of the sign of distance change.¹ Non-linear mapping under these conditions results in a loss of roughly $\log_2(r)$ bits of the depth precision for a given ratio r of the distances to the far and near clipping planes.² It causes severe visual artifacts in open environments with large dynamic range r , typical for flight simulators and games.³ If the W value after projective transformation is proportional to the Z_e value in the eye space, linear mapping can be accomplished by storing W values in the depth buffer.⁴ However, correct implementation of the W buffer requires costly high-precision per-pixel division and isn't widely supported; it isn't suitable for scenes with very low dynamic ranges.

The main factors that affect eye-space depth resolution are mapping of the eye-space distance (Z_e) to the screen-space depth (D) and precision of the format used to store D in the depth buffer:

$$\delta Z_e(Z_e) = \frac{dZ_e}{dD} * \delta D(D)$$

Linear mapping and integer storage format keep W -buffer resolution independent from the distance to the object. We achieve a similar result by using a combination of simple mapping and format precision functions that compensate each other's non-linearity, instead of requiring both of them to be linear.

Screen depth D is computed using all mapping constraints except distance change sign; the result is complementary to the screen Z in the same range [0,1]:

$$D = 1 - Z_s = \frac{d}{f-d} * \left(\frac{f-1}{Z_e} \right)$$

where f and d are the distances to the far and near clipping planes in the eye space. Depth value is stored in the floating-point format:

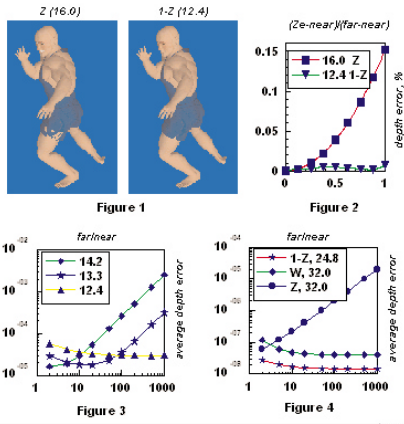
$$D = \begin{cases} 2^{e-2^n} \cdot 1.f \dots (e > 0) \\ 2^{1-2^n} \cdot 0.f \dots (e = 0) \end{cases}$$

where n is a number of bits of exponent (e) and m is a number of bits of mantissa. Depth errors due to mapping and storage format compensate each other when D is close to 0 (the object is close to the far plane), increasing overall resolution. For instance, cloth-body errors in the scene with dynamic range 1000 are visible with 16-bit screen z buffer (Figure 1, left), but not with floating-point complementary z (Figure 1, right).

The above equations define a range of ratios f/d that minimize depth error for the known format (m, n). For instance, 12.4 16-bit format ($m = 12, n = 4$), supported by many 3D accelerators, is optimal to make depth precision of the floating-point complementary z buffer almost independent from Z_e at $f/d = 100$, decreasing errors more than 10 times in comparison with screen z (Figure 2).

Average depth error in the view volume depends on f/d ratio, storage format and sampling distribution. Results for complementary z and uniform sampling (Figure 3) show that 12.4 format is the best for $f/d > 200$, while 13.3 is better for $200 > f/d > 20$. To keep quasi-linear mapping for the wide range of f/d ratios, we propose to support a set of storage formats with different exponent sizes. The best format for the current scene is derived from its dynamic range and region of interest. In many practical cases, it may be enough to use integer format ($n = 0$) for scenes with small dynamic ranges and four bits of exponent for open environments.

When per-pixel size of the depth storage increases, precision becomes limited by per-vertex input format, usually, 32-bit IEEE float. In this case, the complementary z buffer is 2.8 times better than W buffer, without an associated cost/performance penalty (Figure 4).



References

1. Newmann W.M. and Sproull R.F. Principles of Interactive Computer Graphics, 1981 New York, McGraw-Hill.
2. Kempf R., Frazier C. OpenGL Reference Manual, ver.1.1, 1997 Addison-Wesley, p. 164.
3. Swarovsky J. Extreme Detail Graphics, Proceedings of the Game Developers Conference 1999, pp. 899-904
4. Microsoft Corp. What are Depth Buffers? DirectX6 SDK, 1998, msdn.microsoft.com/library/sdkdoc/directx/imover_4xwk.htm

This sketch describes techniques for merging ray tracing with real-time hidden-surface algorithms such as the z buffer algorithm. The technique offers a smooth transition from interactive display with approximate soft shadows, reflections, and transparency to full ray tracing with all of its associated effects. This is done by using the ray tracer only for tracing shading rays at the vertices of polygons and using another hidden-surface technique for hidden-surface removal in conjunction with Gouraud interpolation.

As computer graphics have evolved, we have seen a schism emerge in rendering techniques between the so-called "real-time" techniques and the "photorealistic" techniques. The result: most interactive computer graphics have a characteristic look, with diffuse shading and highlights included because they can be done quickly and easily, but a noticeable absence of shadows and reflections. Photorealistic images, on the other hand, tend to be characterized by having a lot of reflective surfaces in addition to shadows and transparency, simply because these effects are possible. The approach we have taken integrates these techniques in a uniform way so that as faster hardware and multiprocessing become available, we can have a smooth transition from the traditional z buffer to more realistic ray tracing algorithms.

The Rendering Algorithm

If we trace rays at select points distributed across the surface of the object and interpolate the color across the object, then we may only need to trace a few hundred or thousand rays per frame. This can be done at interactive rates. In addition, we can use high-performance z buffer and Gouraud interpolation hardware to perform the tasks of hidden-surface removal and color interpolation, respectively. With this technique, we have found that on current workstations it is possible to compute shadows, reflections, and transparency in scenes with a few hundred or thousand polygons at acceptable interactive rates.

Benefits of Interpolation

One of the most striking things about the images that are produced is that the interpolation technique causes reflective objects to have the appearance of being diffuse reflectors. The same effect causes shadows and transparency to appear "soft." Previously, this effect could only be achieved through extremely expensive distributed ray tracing. Although the diffuse reflections, transparency, and shadows computed through interpolation are not physically accurate, they give a close enough illusion of the phenomena to be convincing.

Disadvantages of Interpolation

There are two major disadvantages with ray tracing at select vertices and interpolating the color in between. The first problem is that the appearance of the object is dependent upon the underlying tessellation. The second problem is that we are using a smooth interpolation to approximate the specular and transmitted components of the surface of the object, which often change very abruptly compared to the diffuse component. This leads to shading anomalies.

Comparison with Alternative Techniques

There are two primary techniques for simulating shadows with the z buffer. The simplest type of shadowing effect is ground-plane shadows. These shadows are created by projecting the object onto the ground plane and drawing dark polygons where the projections lie. These shadows have hard edges and are only cast onto the ground plane, which limits their usefulness. A more general solution is available through shadow mapping. Unfortunately, the shadow maps use a large amount of memory and are prone to artifacts. Reflections are often simulated using environment maps, images of the scene from the viewpoint of the object, which are then mapped onto the object as reflections. Although these look like reflections at first glance, they do not have the proper geometry for true reflections.

Transparency can be simulated by z buffering machines with special frame buffer hardware for "alpha-blending." However, the distortions caused by refraction of light through different media cannot be accurately simulated. Although these techniques are very useful in the hands of a skilled and experienced computer artist, they are not as general-purpose and easy-to-use as a ray tracer. In addition, after these additional features are added, the initial simplicity and speed of the z buffer algorithm is lost. These same effects can all be computed simply and precisely through ray tracing. It therefore makes sense to restrict the z buffer to what it does best (hidden-surface removal) and restrict the ray tracer to the tasks that only it can perform.

Conclusions

We have demonstrated the practicality and usefulness of the algorithms described here in a prototype system. We have found that inclusion of reflections and shadows in a real-time system not only contributes to a better understanding of spatial relationships, but also increases the gamut of possible surface appearances. While this technique may not be appropriate for every real-time application, we have found that in certain circumstances, it can be very effective.






Figure 1: The left end of the beam is fixed. The top image shows its *distorted* deformation under gravity, using linear strain. The bottom image shows the *undistorted* deformation, under the same gravitational force, using quadratic strain.




Figure 2: A flexible body collides with a rigid body.




Figure 3: A 2D example of graded mesh.

References

1. David Baraff and Andrew Witkin. Dynamic Simulation of Non-Penetrating Flexible Bodies, Computer Graphics: Proceedings of SIGGRAPH, pages 303—08, ACM, 1992.
2. Edward John Nicolson. Tactile Sensing and Control of a Planar Manipulator, PhD thesis, EECS, University of California, Berkeley, 1987.
3. D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically Deformable Models, Computer Graphics, 21, July 1987.

Global Deformation Using Nonlinear FEM

We apply the displacement-based finite element methods (FEM) to simulation of large motions and global deformations of deformable objects. A linear strain model leads to unacceptable distortion (Figure 1). To avoid this problem, we apply a quadratic strain instead (Figure 1). Essentially, this requires solving the following nonlinear system of differential equations

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{D}\dot{\mathbf{u}} + \mathbf{R}(\mathbf{u}) = \mathbf{F} \quad (1)$$

where \mathbf{u} is the nodal displacements; $\dot{\mathbf{u}}$ and $\ddot{\mathbf{u}}$, the respective velocities and accelerations; \mathbf{F} , the external force; \mathbf{M} , the mass matrix; \mathbf{D} , the damping matrix; and $\mathbf{R}(\mathbf{u})$, the *nonlinear* internal force due to deformation. For a soft material such as live tissue, the material stiffness is small. This makes explicit the time integration scheme appropriate because we can take large time steps. In particular, we apply a Newmark scheme to equation (1). We diagonalize the mass matrix \mathbf{M} and the damping matrix \mathbf{D} . This leads to a decoupled system of nonlinear equations, which requires no matrix inversions to solve.

Collision Integration Scheme

Simulating deformable object collisions using a penalty method³ requires tiny time steps to produce visually satisfactory animations. A general impulse collision¹ is considered more efficient and accurate but still requires more computation than collision-free dynamics. For deformable object collisions, the collision time can be assumed finite (unlike the instantaneous collision of rigid bodies). By recognizing this, we propose an efficient way to handle collisions.

Consider the collision between a deformable body with a stationary rigid body (figure 2). Assume at time t_n , the node p , with velocity $\hat{\mathbf{v}}(p)_n$, is colliding with a rigid surface of outward normal $\hat{\mathbf{n}}$. Then the non-penetration constraint at node p can be enforced by setting the normal component of $\hat{\mathbf{v}}(p)_{n+1}$ to zero as following:

$$\hat{\mathbf{v}}(p)_{n+1} = \hat{\mathbf{v}}(p)_n + (\hat{\mathbf{v}}(p)_n \cdot \hat{\mathbf{n}})\hat{\mathbf{n}} \quad (2)$$

If we choose $\Delta t_{n+1} = \Delta t_n$ for the Newmark scheme, we have

$$\hat{\mathbf{u}}_{n+2} \cdot \hat{\mathbf{n}} = \hat{\mathbf{u}}_n \cdot \hat{\mathbf{n}} \quad (3)$$

This shows that the non-penetration constraint is exactly enforced after two time steps, *without* solving a constrained problem. For a decoupled system, this collision handling scheme can be easily generalized to multiple point collision constraint. And unlike the coupled system, no explicit impulse calculation is necessary for frictionless collisions. When friction has to be considered, the equivalent impulse can be easily computed without matrix inversion. This collision-handling integration scheme can be easily generalized to collisions between a deformable body and a moving rigid body, and to collisions between deformable bodies. It is worth noting that a deformable object's resting on a surface can be handled exactly the same way.

Graded Mesh

While 2D FEM has achieved great success in achieving real-time performance in computer graphics applications, the computational cost is much higher for 3D applications, mainly due to the increase in the number of elements in the mesh. In a roughly uniform 2D finite element mesh, the number of elements is about $O(n^2)$, where n is the average number of elements in each principle direction. However, a similar 3D mesh would have $O(n^3)$ elements, which leads to a much larger system of equations. Nicolson² shows that the cutoff spatial frequency of an object in response to external loads decreases faster than $1/d$ in terms of the distance d away from the surface. Therefore, if we use a 3D graded mesh similar to the 2D example in Figure 3, we will lose little accuracy with respect to static forces on the surface, while reducing the complexity of the problem from $O(n^3)$ to $O(n^2)$. A similar grading property also applies to tetrahedral mesh. To maintain the geometric compatibility at the element interface, we simply set the displacement of edge constrained nodes, such as node A in Figure 3, to the midpoint of the corresponding edge. Similar constraints apply to the face-constrained nodes in a 3D mesh. On a 400MHz Pentium II PC, a uniform mesh of 1331 elements needs about 0.11 seconds per time step. The graded mesh with the same accuracy needs only 0.06 seconds per step.

*Supported by a Multi-Disciplinary Research Initiative grant for 3D Visualization, sponsored by BMD0 with support from ONR.

Computer-generated line drawings,⁸ are becoming an increasingly popular method of non-photorealistic rendering. While many factors contribute to an effective line drawing, line direction may be among the most critical for conveying surface shape. The principal directions hold particular promise in this respect.² In a recent study, human subjects consistently displayed biases toward perceiving surface contours (lines drawn on the surface) as being aligned with the principal directions.⁵ Despite the potential for principal directions and their suggestion by previous authors,^{1,2,7} there are currently no available techniques for using them to render geometrically invariant line drawings of arbitrary 3D surfaces. Markosian et al.⁴ handled arbitrary 3D surfaces, but line direction was viewpoint-dependent and lines clustered on silhouette edges with little interior curvature information.

Interrante² used a 3D principal direction texture to show the shapes of transparent isosurfaces in volume data. Stalling⁶ suggested line drawing based on 3D streamlines following the principal directions. Figure 1, derived by these methods, was a motivating factor for this research. However, this work focuses on the application to polygonal meshes instead of volume data. The major contribution of this work is to motivate the use of principal directions in line drawings and show how to apply them to arbitrary polygonally defined surfaces. The result is a geometrically invariant set of shaded strokes in 3D which can be rendered in real-time.

Estimating Principal Directions on a Polygonal Mesh

Much of the difficulty of this problem lies in accurately computing smoothly continuous estimates of the principal directions at arbitrary points on a potentially coarse polygonal mesh. The method we use was proposed by Taubin.⁶ First, vertex normals are computed. Then for each vertex v_i , we construct a symmetric matrix M_{v_i} based on the weighted sum of the directional curvatures of the edges in the neighborhood of v_i . The eigenvectors of M_{v_i} are the normal and two principal directions, which together form an orthonormal basis at v_i . To determine the principal directions, we restrict M_{v_i} to the tangent plane and then diagonalize the 2×2 submatrix of the transformed M_{v_i} . The principal curvatures can be calculated with a simple linear transformation on the corresponding eigenvalues.

Tracing Lines Over the Surface

The collection of estimated principal directions results in a vector field over the surface. We trace long curved strokes whose directions are determined by the flow of the vector field. To maintain an approximately even-line density and avoid line crossings. We extend the placement strategy of Jobard et al.³ from 2D images to the 3D surface. Line starting points are chosen randomly on the surface with the constraint that they lie at a minimum distance from existing lines. In umbilic regions where the surface shape is planar or spherical, the principal directions are undefined because there is equal curvature in all directions. In such cases, we tri-linearly interpolate. If there are no neighboring well-defined principal directions, or if there is a sudden change in surface curvature, the line is terminated.

We implemented this technique in OpenGL using shading and hidden line removal. Rendering parameters, line density, length, or waviness, can be varied for different effects. Figure 2 shows the technique applied to a coarse triangular model of a vase.

Conclusions and Future Work

We are engaged in ongoing efforts to improve the principal direction estimation algorithm and apply this technique to more complex models. Future work will incorporate methods for maintaining continuity through umbilic regions and aesthetically merging opposing lines of force. The perceptual effectiveness of principal direction lines for the understanding of surface shape will hopefully be a further area of research.

Acknowledgments

This work was supported by a University of Minnesota Grant-in-Aid of Research, Artistry, and Scholarship. Thanks to Detlev Stalling and Todd LeMoine.




Figure 1 Principal direction line drawing (without hidden surface removal) of a bone/soft tissue boundary surface in a CT volume data set.



Figure 2 Principal direction line drawing of a vase. Mesh consists of 400 triangles.

References

1. R. Coutts and D. Greenberg. Rendering With Streamlines, Visual Proc. (SIGGRAPH 97), p.188.
2. V. Interrante. Illustrating Surface Shape in Volume Data Via Principal Direction-Driven 3D Line Integral Convolution, Computer Graphics (Proc. SIGGRAPH 97), pp.109-116.
3. B. Jobard and W. Lefer. Creating Evenly-Spaced Streamlines of Arbitrary Density, Proc. of 8th Eurographics Workshop on Visualization in Scientific Computing, 1997, pp.45-55.
4. L. Markosian, M. Kowalski, S. Trychin, L. Bourdev, D. Goldstein, and J. Hughes. Real-Time Nonphotorealistic Rendering, Computer Graphics (Proc. SIGGRAPH 97), pp.415-420.
5. P. Mamassian and M. Landy. Observer Biases In the 3D Interpretation of Line Drawings, Vision Research, 38 (1998), pp. 2817-2832.
6. G. Taubin. Estimating The Tensor of Curvature of a Surface From a Polyhedral Approximation, Proc. 5th International Conference on Computer Vision (ICCV), 1995, pp.902-907.
7. G. Winkenbach and D. Salesin. Rendering Parametric Surfaces In Pen and Ink, Computer Graphics (Proc. SIGGRAPH 96), pp.496-476.
8. D. Stalling. Personal Communication.

Real-Time Translation of Human Motion from Video to Animation

Recently, considerable effort has been devoted to animating a virtual human from real human motion in video image sequences.¹ For real-time processing, however, existing systems have limitations in possible human postures and environmental/lighting conditions.

This sketch presents a robust top-down approach to translating human motion in video images to computer animation. In existing approaches, all the motion data are obtained directly from video images, frame by frame. Such a pure motion-capture approach often forces unnatural postures and/or causes sudden postural changes. In our approach, the system makes conjectures on what a real human subject is doing from the information taken from video images, and then motion generators² animate a virtual human smoothly and appropriately.

Approach

A real human subject is assumed to move on a horizontal floor, and a single video camera is placed above the floor. The camera data (for example, height, dip, etc.) are known in advance. Our prototype system is composed of two modules: a region detector and an animation generator. In the region detector, three blobs corresponding to head, trunk, and legs are tracked frame by frame. Our blob-tracking algorithm is similar to Wren et al.,³ but we have improved the blob-labeling equations based on the Maximum A Posteriori Probability (MAP) method.

In our animation generator, possible human motions are categorized, and the motions and their transitions are modeled as a state/motion transition graph. Its conditions for state transition are described with relative positions of the blobs. The animation generator first computes the horizontal position of the human body assuming that the bottom of the legs blob touches the floor. Next, from the relative positions of the blobs, the system infers the currently performed motion by making transitions on the state transition graph.

Even if the system fails to track some blobs, our system makes rough conjectures on the human posture by treating all the remaining blobs as a single blob representing the whole human body. With this feature, our virtual human is animated even if the human figure in the source images is on the order of several pixels. Our approach is thus applicable to a much larger environment than existing approaches.

More technical information is available on the Web.⁴

Results

If the input video image size is 160x120, experiments show that our system can generate about 11 frames per second on a single SGI O2 (R5000). The animated virtual human can walk, sit, and lie as the real human subject performs in an uncontrolled environment like a computer room. (See Figure 1 and additional materials on the Web.⁴)

Conclusion & Future Work

Our system can translate human motion in video images to computer animation in a more robust fashion than existing methods. To animate the human as robustly as possible, our current implementation has some frame delay between real video images and animations. This design decision is made because the system was primarily developed as a privacy-oriented surveillance system, where its observer can understand what its subject is doing without knowing who the subject is. In such a system, some frame delay would be acceptable. To apply our system to other applications (for example, interactions in virtual environments), the delay will be minimized in the near future.

Acknowledgement

This work was supported by ANIMA Electronics Co., Ltd.

References

1. K. Ebihara. Shall We Dance? SIGGRAPH 98 Conference Abstracts and Applications, 124, 1998.
2. T. Noma and N.I. Badler. A Virtual Human Presenter, Proc. IJCAI-97 Workshop on Animated Interface Agents, 45-51, 1997.
3. C.R. Wren, et al. Pfunder: Real-time Tracking of the Human Body. IEEE Trans. on PAMI, 19(7):780-785, 1997.
4. Human Motion Translator Web page. www.pluto.ai.kyutech.ac.jp/~noma/mottrans/



Figure 1: Walking sequence (top: video source; bottom: animation output)

Rendering 3D Objects into Photographs Taken by Uncalibrated Perspective Cameras

Qian Chen
Gérard Medioni
University of Southern California
qianchen@usc.edu
csuri.usc.edu/~chenqian/research.htm

One important element in virtual/augmented reality applications is the insertion of synthetic objects into real images. For this purpose, the camera pose must first be estimated. Typical approaches rely on 3D position tracking devices, precise calibration, or fiducials. Recently, affine invariant has been used to provide calibration-free augmented reality. However, the affine representation makes it difficult to describe the objects in a traditional way (using positions, distances, or angles, etc). Another aspect of the same problem, which has been largely ignored, is the mutual occlusion (or other geometrical interactions) between the virtual (inserted) and the real (existing in the scene) objects. To make this happen, accurate Euclidean models of the real objects must also be recovered.

In our research, we extended the idea of using invariance to perspective cameras, adopting a more general concept: projective reconstruction. We developed new self-calibration algorithms to transform the projective reconstruction into a Euclidean one to gain the freedom of conventional object placement while bypassing the calibration process. Rather than fiducials, image features such as corners of objects were used. As a result, object models and camera pose were obtained at the same time. Finally, special attention was paid to take care of registration errors caused by the deviation in feature detection. The overall workflow has four steps:

1. Monocular image features are detected automatically.
2. Correspondences across images are established interactively.
3. Object structure and camera pose are recovered.
4. Original images are augmented with inserted objects that may cut into the real objects.

In a projective framework, camera projection can be described by a three-stage process: projective-to-Euclidean, Euclidean-to-camera, camera-to-image. We have developed a novel projective reconstruction algorithm using matrix factorization. We have also developed a new self-calibration algorithm that separates the projective-to-Euclidean matrix from the projective reconstruction, thus giving the Euclidean reconstruction. This result is shown in Figure 1. However, using the so-far-obtained cameras for object insertion generates images with severe misalignment (left two images, Figure 2). The error comes from the deviation in feature detection. We addressed this problem by introducing a 3D affine distortion matrix right before the camera-to-image projection. This preserves the Euclidean structure of the world. The final result is shown in the right two images of Figure 2.




Figure 1
Two examples of a terminal sequence, its Euclidean reconstruction together with the recovered cameras.



Figure 2
Misalignment of real and virtual objects (left), and correct registration after error compensation (right).

Perspective Cameras

Representation of the Tactile Surface Texture of an Object Using a Force Feedback System.

The "digital archives project" is being pursued worldwide. This project preserves information on cultural properties such as national treasures, important cultural properties, and works of art as digital information. Because it is saved as digital data, the information can be used indefinitely.

Our appreciation and enjoyment of objects that have a significant 3D structure (for example, sculpture, statues of Buddha, traditional ceramic ware, etc.) can be greatly enhanced by not only being able to look at them, but also to touch and move them. Such virtual tactile feedback systems have been developed with force feedback control technology and CG display technology. But these systems have difficulty representing the physical texture of rough surfaces on an object.

In this approach, the force exerted on the fingertip by the force feedback system is calculated using the location of the finger tip relative to the polygons representing the object. The force is calculated by the normal vector of the polygon, so the same force is felt everywhere on the polygon. For example, when the system displays a "brick," we expect a rough texture, but the current system can not provide feedback for this. So we developed a method to allow the fingertip to feel rough texture using the virtual tactile feedback system.

The dynamic frictional resistance feedback to the fingertip in the system is calculated as follows:

$$F_d = -A b v$$

F_d :dynamic frictional resistance

A :constant

b :dynamic frictional resistance coefficient(0.0-1.0)

v :fingertip velocity vector projected onto the polygon surface

This equation means that if the fingertip moves faster on the object surface, more feedback force is returned. If the dynamic friction resistance coefficient is set to be large, the fingertip feels more resistance force from the system. We simulate the texture of the surface by alternating, at 100Hz, between two dynamic frictional resistance coefficients. We used a "brick" mapped with a finely textured surface. The two coefficients for this surface, $FD1$ and $FD2$, are changed by increments of 0.2, so 15 combinations are evaluated ($FD1 = 1.0$ - $FD2 = 0.8$ to 0.0 , $FD1 = 0.8$ - $FD2 = 0.6$ to 0.0 , $FD1 = 0.6$ - $FD2 = 0.4$ to 0.0 , $FD1 = 0.4$ - $FD2 = 0.2$ to 0.0 , $FD1 = 0.2$ - $FD2 = 0.0$). The results showed that when the difference of the two values is less than 0.2, the texture feels "sticky." When the difference is 0.4 to 0.6, the texture feels "rough." And when the difference is 0.6 to 0.8, the texture feels "jagged." These results show that alternately applying the two coefficients to the object can enable us to feel subtle surface texture.

The fingertip moves across the surface against the two alternating resistance forces. If the resistance force is large, it is difficult to move. If the force becomes small, it becomes easy to move. If the difference of the two alternating forces is larger, the move distance with small resistance force becomes large. So you feel a large rough texture (jagged). On the the other hand, if the difference of two forces is smaller, the move distance against smaller resistance becomes small, so you feel a small, fine, rough texture.

We used Sensable Technology Inc.'s PHANTOM as the force feedback equipment.

Force Feedback

Reference

High-Tech Visual Promotion Center Foundation
development committee, annual report 1997
(March 1998).

Shading and Shadow Casting in Image-Based Rendering Without Geometric Models

Akihiro Katayama
Yukio Sakagawa
Hiroyuki Yamamoto
Hideyuki Tamura
Mixed Reality Systems Laboratory Inc.
katayama@mr-system.co.jp
www.mr-system.co.jp/

Image-based rendering (IBR) methods such as light-field rendering¹ and the ray-space method² succeed in photo-realistic representation of complex objects without geometric models. Now the key to IBR's future is how to accomplish functionalities of conventional geometry-based approaches. We have already reported how to interact with image-based objects in the virtual environment.³ The next goal is how to make objects without polygon models react to transitional lighting conditions. This sketch describes a real-time rendering method that changes the shading of image-based objects and casts appropriate shadows according to the motion of viewpoint or objects and transitions in local lighting.

Method

A set of multiple pictures taken by a moving camera under a specific lighting condition is called an imageset. Our idea is based on a simple method of selecting the imageset with the lighting condition nearest to that of the virtual environment from the imagesets taken under various lighting conditions. Let I_c be the imageset taken under ambient light. An imageset $I(l)$ stores only the luminance component of pictures of real objects taken under a lighting condition l . While changing values of l , the image data are stored into the corresponding imageset. To render an image from a given viewpoint, corresponding pixels are extracted from $I(l)$ and I_c . Then, the chrominance part of I_c and luminance part of $I(l)$ are applied to the pixel in order to get the final shading. Note that as the technology of IBR is used in this process, an image seen from any point of view can be reproduced. The shadow cast by an object is represented by the texture mapping of the image contour, which is viewed from the light position, onto the surface the object resides on as shown in Figure 1.

Implementation

A system called CyberMirage 99 places ray-based data in the geometry space. A light source can be switched on and off, and moved. Of course, translation and rotation of objects and observers' walk-throughs are implemented. Note that all of these functions are realized in real-time. Figure 2(a) shows an image of objects lit from a certain light source. Note that the objects are appropriately shaded and cast an appropriate shadow. Figure 2(b) is the image when an observer comes near to the objects, and Figure 2(c) is the image when an observer approaches another side of the objects. Figures 2(d) and 2(e) are images when the light source is moved and the objects are rotated, respectively. These figures show the effect when the shading and cast shadow are appropriately adjusted.

Future Work

The next step of our study will explore how to expand the system so that it can accommodate multiple light sources and react to the change in intensity of light. Theoretically, this can be accomplished by preparing a lot of imagesets corresponding to the number of light sources and their intensity levels. However, it requires a new horizon of data-compression methodology in order to realize the real-time interaction.




Fig. 1 Generation of cast shadow




Fig. 2 Views of the objects under different lighting conditions

References

1. M. Levoy and P. Hanrahan. Light Field Rendering, Proc. SIGGRAPH 96, pp.31-42, 1996.
2. T. Naemura, et al. Ray-Based Creation of Photo-Realistic Virtual World, Proc. VSMM'97, pp.59-68, 1997.
3. S. Uchiyama, et al. Presentation and Interaction of Virtual 3D Objects Without Geometric Model, Proc. HCI'97, Vol.2, pp.869-872, 1997.

Shape Extraction for a Polygon Mesh

This sketch reports on our new technique for deriving geometric shape for both organic and inorganic models. This information can be used in computing data structures related to the object (for example, bounding volume hierarchy for collision detection and ray tracing, and level-of-detail for different parts of an object).

Several previous studies have reviewed generation of polygon meshes from a given point set.¹ However, given a polygon mesh, there are no well-known studies of computing a good hierarchical arrangement of polygons. A portion of our work² is the first attempt in this direction. Though the work has achieved good results in partitioning objects efficiently into smaller parts, it is sometimes affected by user parameters and modeling variations. In this sketch, we seek to devise a technique that is automated, efficient, and robust with respect to modeling variations.

Shape Extraction

Shape extraction involves two parts: computing atomic parts (each consisting of a set of polygons) that make up the input model and arrangement of these parts into a hierarchy. The proposed algorithm addresses these using a skeleton derived from the model.

Atomic Computation

A skeleton of a given model is a collection of vertices and edges resulting from a sequence of operations (such as edge contraction) that collapses polygons of the model. Each element e of an edge or vertex of the skeleton is associated with a set of polygons $T(e)$ of the given model collapsed to e . The union of all $T(e)$'s of a connected part of the skeleton is considered an atomic part of the model and is a leaf in the shape hierarchy, H .

Our previous work on using a simplified model derived from model simplification² is a variant of a skeleton. In another development,³ a series of edge-collapse operations is applied to the input model to obtain a so-called skeleton for a different purpose. Our experiment on directly applying this algorithm to compute the atomic part was not satisfactory, in particular, for organic objects. Its drawback is absorption of small (but distinct) features of a model into bigger components. Our method uses weights to control the order of collapse and makes use of previous edges to guide the collapsing.


Hierarchy Building

Each atomic part is now a leaf of the shape hierarchy H . The next step is to link these leaves into the top few levels of H . Construct G , where each vertex in G represents an atomic part and each edge connects two vertices if their associated set of polygons shares some edges or vertices. Assuming G is a single tree, hierarchy building begins at the leaf nodes of G by "folding" up until a single node is reached. At each step, create a node m in H to contain the leaf nodes and their immediate predecessors, and delete the leaf nodes from G .

The resulting H may not be binary and can be made binary with an approach² for computing bounding volume hierarchy. Also, similar ideas can be applied to connect various disconnected graphs G (each can be made as a tree by removing edges in cycles) from the skeleton.

Preliminary Results

Our preliminary results show good partitioning of geometric models, both organic and inorganic. The figure shows examples of skeletons (first row), and atomic parts (second row, where each part is shown in a different color) obtained by the current implementation. Future work includes the use of the shape information for various applications such as bounding volume hierarchy computation.



References

1. Amenta et al. A New Voronoi-Based Surface Reconstruction Algorithm, SIGGRAPH 98, pp 415-421.
2. Tan, Chong and Low. Computing Bounding Volume Hierarchy Using Simplified Models, Proc. Symp. on Interactive 3D Graphics, April 1999, pp. 63-69.
3. Deussen et. al. An Illustration Technique Using Intersections and Skeletons, Proc. Graphics Interface, June 1999.

How do we present the motion of objects in computer-generated still images? Generally, we don't. Or, if we do, we use motion blurring to simulate a real-world camera.¹ Motion blur costs a great deal of extra rendering time, and the only thing it does is blur the objects so that their contours are unrecognizable. The goal, to convey information about the movement per se, is not entirely met. If, however, we consider the application of non-photorealistic rendering techniques, it is promising to adopt successful illustrative techniques from comics to depict past and future motions of objects in a single image.

Presenting Motion

Speedlines are an important stylistic element in comics, and although they are not based on an exact physical model, their use is well known to all of us.² Other means of depicting motion in still images include motion arrows and contour repetitions. In the latter case, earlier positions of an object are drawn with less detail or slightly faded, whereas the current state of the object is presented in all detail.

Besides using those techniques in still images, their use has also been adapted in many classical animations.³

Generating Speedlines, Contours, and Arrows

For an automated generation of speedlines, arrows, and contour repetitions, models coming from typical animation systems (such as 3D Studio) provide all necessary data. After executing a rendering pipeline for line drawings³, we used 3D information gathered during the rendering in order to establish the appropriate visibility and the perspective correctness of the speedlines. Here, a 2D-based approach would be futile, since we have to keep track of the contour of every single object. We then calculate a motion path from the given keyframe data.

As these speedlines were invented by human artists, their generation requires some heuristics about where and how to draw them. In general, speedlines and contour lines:

- Are drawn in the opposite direction of the movement (thus reaching into the past).
- Start at "characteristic" points of the moving object.
- Embrace the minimum and maximum extent of the object.
- Are more or less equally sized, shaped, and directed, without intersecting each other.
- Are uniformly, but not too regularly distributed.

Candidates for starting points of speedlines are vertices of the 3D mesh which yield the outline of the object. The algorithm divides the shape into a number of stripes where each stripe can hold one speedline. If a stripe contains no such vertex, additional vertices are generated on the object's outline. If we do not restrict ourselves to the outline and also take into account inner contour lines, the result can be further improved. In addition, the speedlines should not stick to an object: the algorithm should draw them with an offset. The computation of these effects takes only a small fraction of the rendering time.

Drawing Lines with Style

The output of the renderer consists of a vector-oriented description of the image. Actual lines are drawn using line styles which, in turn, simulate hand-drawn pen strokes by the superposition of the drawing path and the chosen style deviations.⁴

Although our system supports parameter control in every detail, only very little user interaction is necessary, since nearly all speedline parameters can be computed based on the animation data. Therefore, length, width, distribution, number, etc. are determined by applying given default rules.

Some Results

The pictures show different models with speedlines applied to them. We have chosen to render quite simple models to concentrate on the effects of speedlines and contour repetitions. Also, we applied line styles only to the speed elements. Further examples including a short animation can be found at isgwww.cs.uni-magdeburg.de/~masuch/gallery.html






Figure 1

References

1. M. Potmesil, I. Chakravarty. Modeling Motion Blur in Computer Generated Images, In SIGGRAPH 83, pp. 389-399.
2. S. McCloud. Understanding Comics - The Invisible Art, Harper Collins, New York, 1993.
3. F. Thomas, O. Johnston. The Illusion of Life: Disney animation, Hyperion, 1981.
4. T. Strothotte et al. Computational Visualization: Graphics, Abstraction, and Interactivity, Springer, 1998.

Stereo Analyst: Visualizing Large Stereoscopic Imagery in Real-Time

Stereo Analyst is a new software package designed for visualizing stereoscopic imagery in real-time and for collecting 3D features in stereo. The initial targeted users for this product are those who want to visualize terrestrial image data, both in the form of aerial photography and satellite imagery. However, the tool is quite practical for visualizing close-range photography as well. Stereo Analyst includes a set of editing tools for creation of feature databases, such as road maps and 3D building outlines.

One of the primary requirements for the application is that it provide smooth real-time pan and zoom navigation. It utilizes stereo viewing hardware, such as liquid-crystal shutter glasses, in combination with OpenGL's interface for creating stereo-in-a-window on the desktop. If stereo viewing hardware is not available, the system will also render the imagery using a color anaglyph mode.

The tool allows the user to display any pair of images with overlapping areas in stereo. The user can position, rotate, and scale each image interactively relative to the other, so that a proper epipolar alignment can be achieved, and stereoscopic viewing becomes possible.


If camera position is available for the imagery, Stereo Analyst will automatically provide on-the-fly epipolar projection for the images. This is accomplished by using OpenGL's 3D transformation pipeline, coupled with hardware texture mapping, ensuring that this occurs interactively in real-time.

Another requirement for this application was that it be able to handle arbitrarily large volumes of image data. In this case, it is clearly not possible to read an entire image into main memory and then to texture memory. The problem is further complicated by the fact that with stereo imagery, there are two image sources, so twice as much texture data is needed at any given time. A texture-caching scheme is used, whereby only image data for the current area of interest is paged from disk. A multi-threaded design is used to separate slow disk access from the interactive rendering thread.

Multi-resolution decomposition of the data is used to optimize real-time performance. In this case, we want to render the data at reduced resolution when the available screen resolution is less than the available resolution of the data to be rendered from the current viewpoint. For the imagery data, we pre-process the image files to create pyramid layers, which can be thought of as a pyramid of sub-sampled versions of the imagery, with each successive level providing a power-of-two reduction factor.

Interactive fallback modes allow the application to perform well on a full range of graphics hardware. For instance, if only a small amount of texture memory is available, the system can be set to use reduced resolution data while the data are in motion, such as during panning and zooming. When the movement is halted, the full resolution is rendered, possibly at a sub-real-time redraw speed.

Stereo Analyst was developed using OpenGL, and runs on the Intel-based Windows operating systems. It is designed with a "plugin-able" architecture, meaning that it is easily extensible for domain-specific applications, both in terms of imagery sources and overlaid data collection and visualization.



Color anaglyph stereo screen capture from Stereo Analyst. The image is best viewed using a pair of red-blue glasses, in order to see the 3D stereoscopic effect. The imagery is from a pair of high-altitude aerial photographs taken near Laguna Beach, California.



Imagery

A 3D CAD system, in which users can feel as if they actually deform a real mockup and directly paint some pictures on it, has been developed using the following techniques.

One is the simple augmented-reality technique whereby a user's real hands and virtual objects are made to appear to self-occlude each other. The system uses a physical prop with the rough shape of the virtual object as an input device. A half-silvered mirror with 30-percent transmittance and a 3D tracker in the input device are used as the corresponding image of the input device is always overlaid on it.

There were many systems with this structure,^{1,2,3} but only the tangible modeling system has realized appropriate occlusions between real and virtual objects in a very simple manner. To achieve this feature, the system utilizes the fact that a human cannot distinguish a darker object or image when a brighter image or object is overlaid on it.

The input devices and the walls inside the system are coated with deep-black cloth called "Haimiron" and illuminated with spotlights. The brightness of an object covered with Haimiron is 0.6-1.2 [Cd/m²] when it is lighted up with the spotlights, whereas the brightness of human hands in the same condition is 28.6-57 [Cd/m²]. Therefore, the user's hands in the workspace can be seen through the half-silvered mirror, but the devices and the walls inside cannot. Meanwhile, the brightness of the white image (reflected on a half-silvered mirror) is 12 [Cd/m²]. The brightness of human hands overlaid with the white image is 21-30 [Cd/m²]. Therefore, hands can occlude the image to some extent when they are in front of the input device (the virtual object). Hands are automatically occluded when they are behind the input device (Figure 1).

Another technique of this system is that the input device has a force/torque sensor for directly measuring hand action for object deformation. Currently, the sensor can detect hands as they twist, stretch, shrink, bend, or shear. The actual shape of the device cannot be changed, but its corresponding image is deformed using the FFD (free-form deformation) algorithm along with the classified patterns. Users can passively feel the resistant force from the device when they hold, push, and deform it.

In one previous system, the hand device can be regarded as a virtual object itself,⁴ but such a system does not support object deformation. The tangible modeling system can be applied to direct shape modeling and 3D painting or sticker pasting for early industrial design (Figure 2).




Figure 1: Integration of user's hands and a graphic image.




Figure 2: Interactive deformation of a 3D model.

References

- Schmandt, C. M. Spatial Input/display Correspondence In a Stereoscopic Computer Graphic Work Station, *Computer Graphics*, 17(3), 253-262, 1983.
- Ullmer, B. and Ishii, H. The MetaDESK: Models and Prototypes for Tangible User Interface, *Proc. UIST'97*, 223-232, 1997
- Poston, T. and Serra, L. Dexterous Virtual Work, *Communication of the ACM*, 39(5), 37-45, 1996.
- Goble, J. C., Hinckley, K., Pausch, R., Snell, J. W., and Kassel, N. F., Two-handed Spatial Interface Tools for Neurosurgical Planning, *IEEE Computer*, 28(7), 20-26, 1995.

Tracking and Modifying Human Motion with Dynamic Simulation

Subtle details in the motion of animated, human-like characters affect the believability, aesthetics, and impact of an animation or virtual environment. However, none of the available techniques for animating characters make including these stylistic details in the generated motion easy. Motion capture produces characters with rich detail in their motion, but the data are difficult to modify for new characters and situations. Dynamic simulation generates physically correct motion for characters that can respond interactively in a changing environment. However, the controllers required for simulated characters are difficult to construct because we do not know how to specify the details of human motion procedurally.

In this sketch, we describe a technique that uses a dynamic simulation to track and modify motion capture data of human upper-body movements. By combining simulation and motion capture, we hope to retain the interactivity and realism of dynamic simulation and the subtle details of the human data while avoiding the disadvantages of each approach.

Our system uses motion-capture data as the desired inputs to a tracking controller for a dynamic simulation of the character. The simulation is created using the mass, moment of inertia, and limb lengths of the animated character. Thus, the generated motion is physically correct for that character. Driven by the tracking controller, the simulation follows the input motion with similar trajectories, thereby maintaining the style of the human data. This system was used to perform a variety of gesturing and signalling behaviors for several humanoid models.

Three additions to this basic system allow it to be used to generate a richer variety of behaviors. First, a collision handler is added to generate realistic dynamic impacts such as hands clapping in a game of patty cake. The dynamic model applies reaction forces, and the hands collide in a believable way.

Second, specialized task controllers are added to allow editing at the task or behavior level. As a simple example, consider the task of moving a staff up and down rhythmically. Kinematic differences between the motion capture subject and the animated character will cause the staff to wave around rather than move vertically. A task controller fixes this problem by setting the desired angle of the wrist to keep the hand in the correct orientation. The remaining degrees of freedom of the upper body track and maintain the style of the input motion. In a similar way, the system uses a balance controller for the degrees of freedom in the lower body and tracks motion capture data for the upper body to produce full-body motion.

Finally, by modifying the input data, our technique allows the user to edit motions at a high level while relying on the simulation to maintain physical correctness. We used a patty cake example to experiment with several such modifications. We adjusted the input data for the arms with inverse kinematics to ensure contact. Using this input, the tracking system generated physically plausible clapping sequences. We generated arbitrary clap sequences by re-ordering and time scaling segments of the input data. The system created physically realistic transitions between two distinct sequences by interpolating between two input motions and tracking the resulting data. The transitions are smooth because the resulting sequence obeys a consistent set of physical constraints for the duration of the motion.

Our preliminary results demonstrate that this approach to combining simulation and motion capture data maintains the style of the original motion while allowing the data to be modified for new characters and situations. Animations and a technical report with details on the system implementation, results, and related research can be found at: cc.gatech.edu/gvu/animation/Areas/datadriven/datadriven.



The top row shows a simulated character and the human actor whose motion is being tracked. The second row shows a variety of tracked behaviors. The third row shows an example of combining balancing and tracking on the left and simulations playing patty cake with dynamic impacts on the right.

Our new paradigm of handling digital video means moving from computer-oriented forms to user-oriented forms. In the computer-oriented approach, digital videos are represented just as a simple dataset, which means they have titles and codes. Adding semantic script information makes the datasets more user-oriented, but we feel that this is insufficient. A suitable body, a 3D representation, is needed to realize an attractive and powerful support system for human interaction.

Goals of Video Embodiment

Video embodiment is an attempt to create a body that expresses video content and structure. The most important issue is how to reflect video content and structure in the shape of the body. The shape should be compact while expressing the information that allows us to judge how suitable the video may be. The body should also simplify editing of the material.

Our Approach

This sketch presents our vision of video embodiment and introduces design activities including MovieSpiral to illustrate our key concepts:

1. Automatic feature extraction and semiautomatic structuring, to extract the important implicit features of video, and to construct structures using features and knowledge.
2. Transformation rule, to define how to transform the extracted features into an attractive and compact shape.
3. Direct and intuitive operation, to realize more direct control of editing operations with direct feedback.
4. Basic function of editing and processing, to design interaction tools that offer comprehensive and intuitive functions.

DNA-Like MovieSpiral

In nature, deoxyribonucleic acid (DNA) is a very compact form of sequential information. It consists of spirals within a super spiral, so it is spatially compact and easy to access. We shape virtual video to imitate DNA, as described below. A video is a linear sequence, but it has a hierarchical structure consisting of story, scenes, and shots. The DNA-like spiral structure is created by making first-order loops for cuts and second-order loops at scenes. The size of each shot loop represents the length of the corresponding shot, and the size of a scene loop represents scene length. Figure 1 shows the MovieSpiral of a commercial message. The total number of shots is 18, and the story consists of one scene. Another sample is a famous Japanese movie. The movie consists of 349 shots and 7 scenes.

Features

- Compact representation and montage visualization. When we see an entire MovieSpiral, we can grasp the technique the director used in making the video.
- Seamless interaction. We can obtain more detailed information. The MovieSpiral provides clues about content and structure, so it is hard to lose your way in the video information space.
- New interaction-style viewing operation. MovieSpiral is compact and expresses a lot of information that may allow us to discover attractive videos. We move around in space and watch and touch the shape to get more information.
- Identification. In fact, it seems possible to identify the style of a movie from the shape of its MovieSpiral. For example, a shape consisting of many short shots means an action scene. We may find that some directors produce works that have their own distinctive MovieSpirals.
- Manipulation based on MovieSpiral. The core of manipulating video is to interact with temporal information. MovieSpiral provides an environment in which video can be manipulated as it provides a better understanding of content and structure.



Figure 1
The MovieSpirals of a commercial message and a movie.

Virtual Car

This sketch describes development of a sales support tool that utilizes virtual reality (VR) technology. The project was commissioned by DaimlerChrysler.

Objective

The Virtual Car system presents all possible variants of the vehicle. As an alternative catalogue, it includes all extras, colors, and fabrics, and it is intended to support every customer's individual Mercedes-Benz configuration. Therefore, a very high standard of visualization of the vehicle and its features is of cardinal importance. An important aspect of the system is its employment in the showroom, where it is intended for long-term use. It must not look and perform like a computer terminal, but must incorporate a user-friendly interface. Ergonomic criteria such as comfort and usability were given special consideration.

Fundamental Concept

At first, the basic concept was: "One World - One Window." In the showroom, the real vehicle is replaced by a virtual car at 1:1 scale (Figure 1). Using a flat touchscreen mounted to a swivel arm, the user navigates through the car's interior and exterior. The visual representation is updated in real-time in coordination with the user's movements. The touchscreen enables further interaction with the virtual car.

Hardware User Interface Design

Traditional VR is still characterized by helmet and glove. In our case, the intended purpose of the tool did not allow that interface solution. It would have meant an unacceptable compromise of comfort, and it would have excluded the user from the real ambience.

Since all other commercially available input and output devices were not acceptable, a novel type of interface had to be developed for the Virtual Car application. In contrast to the classical approach in VR, the user is not immersed in the virtual environment. Instead, the virtual representation becomes immersed in the real world. Utilizing a LCD with integrated touch screen that can be freely moved in 3D space, the user experiences the car through a window into the virtual world (Figure 2). The display is framed and mounted to a swivel arm, and the 20-inch LCD (1280x1024 pixels, 16 million colors) displays almost true colour. The sensors in the axes of rotation record movements within six degrees of freedom. In order to allow operation of the touch screen the swivel arm is automatically locked in any position when the user lifts a hand to operate the touch screen.

Interactivity

Design of the integrated hardware and software systems was closely coordinated throughout the development process. All aspects of the car may be examined in detail. Using the touchscreen, the user can re-configure the car to individual specifications. If the user selects a piece of equipment, it is shown all options also available for the real car. Everything from colour, engine, wheel rims, and seats, down to the style of the gear shift knob can be selected according to the user's individual taste (Figure 3).

Software

Special authoring tools were required for compilation of the 3D model data and the essential interaction programming of the actual application. Intelligent culling tools permit interactive presentation of the model at a minimum of 12 frames/second. For each frame, one graphic pipe of an SGI Onyx2 InfiniteReality System visualizes an average of 350,000 textured polygons.

Outlook

In the next Virtual Car product cycle, the ergonomics of the integrated input and output device will be further improved. The system will also be integrated into the sales infrastructure of the manufacturer, and the process of data integration will become more cost-efficient.




Figure 1
Fundamental concept




Figure 2
Hardware user interface design



Figure 3
Car configuration




Figure 4
Employment of the system

Reference

1. "The Virtual Car Project"
(www.artcom.de/projects/vrf/).

Little attention has been given to colored pencil drawing (CPD), which has been regarded as secondary in relation to painted or finished work. However, CPD has become an artform in its own right, and it is often used for package illustration and picture books.

A CPD-like image could probably be generated by combining various functions provided by existing digital painting systems. A number of common CPD techniques exist,¹ and direct and faithful reproduction of these techniques is essential in order to represent the charm of CPD.

We have consistently taken a volume graphics² approach to modeling of CPD. Our previous CPD model consists of two sub-models.³ One describes the 3D micro-structure of paper. The other defines a method for pigment distribution, formulated as the interaction between the paper and the lead of a pencil. In this sketch, we incorporate into the model pigment redistribution sub-models for two common CPD techniques: watering and adding eraser effects.

Watering

When using water-soluble colored pencils, watering is a real CPD technique for redistributing pigment by running a wet brush over a CPD-surface. After watering, a certain amount of pigment is spread and blended, depending mainly on the uneven conditions of pigment distribution, the quality of the bonding agent in the pigment, and the quantity of water.

In this model, voxels accessible from a brush are located using volume accessibility. If a pigment voxel in the stroke region is accessible, the pigment in the voxel dissolves (becomes a "source voxel"). Pigment in a voxel can also begin to dissolve if the voxel touches a source voxel, depending on the stroke pressure and the bonding agent in the lead.

After all the source voxels are located, the brush stroke vector is used for each of the voxels as a local streamline for our 3D extended version of Cabral's line integral convolution,⁵ and the pigment movements are simulated. For each "sink voxel," in which the dissolved pigment collects, the contribution weight of the source voxel is determined according to the distance between the two voxels.

Using this sub-model, we added watering effects to a CPD volume, which was generated using our previous CPD model.

Eraser Effects

Among many CPD eraser techniques, we put a particular focus on the reproduction of the soft tint effect, because this technique has no substitute and is therefore imperative for creating CPDs with a soft impression. A soft tint is obtained by applying lightweight eraser strokes to CPD surfaces. Volume accessibility again plays a key role in evaluating which voxels are accessible from the eraser. In this case, an eraser is assumed to move only in the vertical direction. The top few layers of accessible pigment voxels are emptied by an eraser stroke. The thickness of the emptied layers is estimated according to the softness of the eraser and the pressure of the eraser stroke.




Figure 1
A volume visualizer called VolVis6 was used for volume rendering. The original pencil stroke looks somewhat scratchy, whereas pigment extended smoothly by the wet brush stroke results in the expression of a gentle feeling similar to real water-soluble CPDs.



Figure 2
A CPD volume in which an eraser effect was added around the tip of a leaf. In addition to the scratchy strokes peculiar to CPDs, the areas where thin layers of pigment remain present a lighter feeling or impression.

References

1. Martin, J. The Encyclopedia of Colored Pencil Techniques, Running Press, April 1997.
2. Kaufman, A., et al. Volume Graphics, IEEE Computer, vol. 26, no. 7, pp. 51-64, July 1993.
3. Takagi, S. and Fujishiro, I. Microscopic Structural Modeling of Colored Pencil Drawings, SIGGRAPH 97 Visual Proceedings (Technical Sketches), p. 187, August 1997.
4. Miller, G. Efficient Algorithms for Local and Global Accessibility Shading, Proc. SIGGRAPH 94, pp. 319-326, July 1994.
5. Cabral, B. and Leedom, L. C. Imaging Vector Fields Using Line Integral Convolution, Proc. SIGGRAPH 93, pp. 263-272, August 1993.
6. Avila, R., et al. Volvis: A Diversified Volume Visualization System, Proc. IEEE Visualization '94, pp. 31-8, October 1994.

One of the many challenges in modeling, animating, and rendering believable mammals in computer graphics has been the generation of realistic-looking hair and fur.

Approaches can be divided into two basic categories: modeling and rendering. In the former case, individual hair primitives are geometrically defined, whereas in the latter case, the geometry of hairs is “faked” through special rendering techniques. In this sketch, we discuss two new effects we implemented as part of our hybrid geometric/rendering hair/fur pipeline:

1. Production of a wet fur look by clumping of individual hairs.
2. A breaking up of combed hairs along fur tracks on the skin.

Clumping of Hairs

Clumping of hairs can occur when the fur gets wet. The effect is that the tips of a set of neighboring hairs (clump hairs) tend to gravitate towards the same point (clump-center hair), creating a kind of cone-shaped “super-hair,” or circular clump. We have implemented two kinds of clumping: static area clumping and animated area clumping. The former generates hair clumps in fixed predefined areas on the model, whereas the latter allows clumping areas to move on the model. In both cases, we provide parameters that can be animated to achieve various degrees of dry-to-wet fur looks.

Static Area Clumping

There are four required clumping input parameters: clump-density, clump-size, clump-percent, and clump-rate. Clump-density specifies how many clumps should be generated per square area. Clump-size defines the area of a clump in world space. Our clump method also has an optional clump-size noise parameter to produce random variations in the size of the clumps. Finally, our clump method also assigns a clump-percent and clump-rate value to each clump hair. The values for both range between 0 and 1. Clump-percent specifies the degree of clumping for a clump hair: a value of zero means that the hair is not clumped at all. It is like a “dry” hair. A value of one means that the hair is fully attracted to its clump-center hair. The tip of this hair is in the same location as the tip of the clump-center hair. Clump-rate defines how tightly a clump hair clumps with its clump-center hair. Both clump-percent and clump-rate can also be animated to provide continuous control for dry-to-wet-to-dry fur looks. This is illustrated in Figure 1, which shows two frames from an animated clump-percent and clump-rate sequence. In the top image, clump-percent is 0.9 and clump-rate is 0.3, which results in a moderate wet look. In the bottom image, clump-percent and rate are both 1.0, which produces a very wet look.

Animated Area Clumping

Animated area clumping is desirable if we want to simulate spouts of water or raindrops hitting the fur and making it increasingly wet. Animated clumping areas are defined through particles hitting surface patches. These particles originate from one or more emitters, whose attributes determine, for instance, the rate and spread of the particles. Once a particle hits a surface patch, a circular clumping area is created on the patch at that location, with clump-percent, clump-rate and radius determined by a creation expression. Each area contains several actual clumps, depending on its radius and on clump-size. Runtime expressions then define clump-percent and rate, to determine how quickly and how much the fur “gets” wet.

Breaking of Hairs

Hair breaking occurs because fur sometimes breaks along certain lines (fur-tracks) on the skin. We have identified and implemented two kinds of breaking: symmetric and one-sided. In symmetric breaking, hairs on both sides of a fur-track “break” towards that track, whereas in one-sided breaking, hairs on one side of the track break away from the track. Fur-tracks are specified as curves-on-surfaces. Each track has a radius, break-percent and break-rate for symmetric and one-sided breaking, and an additional break-vector for the latter.





Figure 1: Two frames of an animated static clumping sequence. The model is defined by 34 NURBS patches and has approximately 450,000 hairs.

Introduction

The line integral convolution (LIC) technique has been known to be an effective tool for depicting flow patterns in a given vector field.¹ There have been many extensions to make it run faster and reveal useful flow information such as velocity magnitude, motion, and direction.² There are also extensions to unsteady flows and 3D vector fields. In this sketch, a new method for highlighting flow direction in LIC images is presented. The method gives an intuitive impression of flow direction in the given vector field and reveals saddle points in the flow. The method presented is an automatic approach for highlighting flow features, which often represent the most important and interesting physical flow phenomena.

Directional LIC

A fast and simple method to depict overall flow direction in the given vector field is presented. The method, which is referred to as directional LIC, colors LIC flow texture based on its orientation and uses local streamlines computed in LIC to show the actual flow direction. The method proceeds as follows.

For each pixel in the vector field, the orientation of the flow is classified as either clockwise or counterclockwise. The direction of orientation can be determined by the curl of the local velocity vector V . Let $V=(u,v,w)$; then $\text{curl } V = [dw/dy-dv/dz \quad du/dz-dw/dx \quad dv/dx-du/dy]^T$; $\text{curl } V$ represents the axis of rotation and the amount of rotation. For 2D vector fields, $\text{curl } V = [0 \quad 0 \quad dv/dx-du/dy]^T$ and it represents a vector perpendicular to the XY plane. Directional LIC uses a color input noise, where the R, G, and B primary colors are convolved individually, and the texture color at each pixel is set according to $\text{curl } V$. If the z component of $\text{curl } V$ is positive (i.e., the flow is rotating in counter-clockwise direction), then the red primary color is set to zero; thus the texture would be colored in some variation of cyan hues. Otherwise, the blue primary color is set to zero to create yellow flow textures when $\text{curl } V$ is negative. This results in a color LIC image colored according to flow direction. However, it only shows regions with the same flow direction, not the actual particle path in each region.

A straightforward method is to release particles/dyes randomly in the input field and overlay the particle pathlines on the LIC image. Depending on the seed locations and the length of the pathlines, the underlying LIC flow texture may be obscured. A more ideal selection of seeding is one where the flow is changing from one orientation to another (neighboring pixels have opposite signs of $\text{curl } V$). Pixels in these regions are in the transient zones. Figure 1 compares directional LIC (1b) with traditional monochrome LIC (1a) using two data sets. A benefit of the directional LIC method is that it automatically highlights saddles in the flow. As shown in Figure 1b, there is one saddle point in the first data set and two saddle points in the second data set. It can also highlight flow separation and flow reattachment, since the flows are approaching from opposite directions. The transient zones are the boundaries where the texture color changes from yellow to cyan. Figure 1c shows directional LIC with pathlines released from transient zones. Particles are released from all pixels in the transient zones. The particles are then tracked during the positive convolution to form pathlines. At the end of each convolution, every other pathline is tracked continuously until there are no more pathlines. Furthermore, pathlines are terminated when they reach another transient zone. With these pathlines, the flow directions are even more clear.

Topology LIC

Another choice of seed locations for the directional flow lines are those near the critical points (points where the velocity is zero). Algorithms for computing critical points are well known in tensor field visualization. We propose to further enhance directional LIC by choosing seed points based on the critical point locations. This would give us additional insights about the flow.




Figure 1 (a) Monochrome LIC (b) Directional LIC (c) Directional LIC with Pathlines

(a)

(b)

(c)

Flow

References

1. Brian Cabral and Casey Leedom. Imaging Vector Fields Using Line Integral Convolution, in Proceedings of ACM SIGGRAPH'93, August 1993, pp. 263-270.
2. Kwan-Liu Ma, Brian Cabral, Hans-Christian Hege, Victoria Interrante, and Detlev Stalling, Texture Synthesis With Line Integral Convolution, Course notes no. 8, Aug. 1997, ACM SIGGRAPH '97.

WorldBoard: Enabling a Global Augmented Reality Infrastructure

Information in Places

WorldBoard is a family of technologies that enables association of digital objects (HTML, Java applet, image, sound files, etc.) to any physical location on the planet and to predefined tagged objects. In its simplest form, WorldBoard can be thought of as a global bulletin board containing geocoded messages or Web pages that are attached to a particular location (longitude, latitude, and altitude) or tagged object. In its ultimate incarnation, WorldBoard provides a global augmented reality infrastructure in which virtual content is perceived to be co-registered and present in the physical world.

Imagine:

- Going to a museum and seeing virtual post-it notes or 3D sounds attached to exhibits.
- Entering an office building with a virtual red carpet leading you to your eight o'clock meeting.
- Shopping for groceries and having information and nutrition analysis tools virtually attached to each product by the retailer, the manufacturer, your physician's prescribed health plan, and even Weight Watchers.
- Kids on a field trip to a wetland with virtual research and analysis tools attached to different locations.
- A construction worker looking at the ground and seeing the underground gas lines.

WorldBoard Forum


The WorldBoard concept was originally conceived, and some early prototype work was conducted, in the Apple Research Labs in 1996. Since 1997, a multidisciplinary team of researchers at Indiana University has been conceptualizing and spearheading the development of the requisite infrastructure. This infrastructure is designed to support any individual or business desiring to develop content, software, or hardware that takes advantage of location. These tools were recently released and are currently being reviewed by the WorldBoard Forum, an expanding group of university and business collaborators who are interested in the implications of WorldBoard technologies. The forum is open to anyone who wants to participate in development or discussions.



Infrastructure Tools

WorldBoard is simply an extension of the Web. Therefore, all of our tools are designed to fully leverage current and emerging standards and protocols. The following are the key components we are developing:

- Mediated Reality Markup Language (MRML), a draft standard based on XML for geocoding documents, describing how the client uses the data, and how the client renders the objects.
- WorldBoard Local Server, a database and parser that runs on a normal Web server. It parses MRML documents and notifies the search engine of key information from the documents.
- WorldBoard Search Engine, a central search engine to route users to every geocoded document on known local servers.
- WorldBoard Clients, designed to flexibly support many types of applications. Viewers can be created to render many types of data for various devices.

This infrastructure is designed to flexibly support development of any type of WorldBoard-enabled device or content. It does not dictate what kinds of I/O devices a user will have (for example, GPS, wireless LAN triangulation to determine position, or Aether wire), what the current state of the technology is, or the infrastructure that is available in a particular location. Instead, it enables developers to create the software tools and assemble the hardware in ways that make sense for the user and the content that the user will access. A primary goal of the WorldBoard Forum is to create and facilitate access to WorldBoard materials and assure that WorldBoard content is easy to create, access, and navigate. Anyone should be able to attach anything to any location or object on the planet. In keeping with this philosophy, all infrastructure components are freely available via an open-source license from the above URL.





ACM SIGGRAPH
1515 Broadway, 17th Floor
New York, New York 10036 USA
+1.212.626.0500
+1.212.764.5537 Fax



Sponsored by ACM SIGGRAPH

